Monitoring Expense Report Errors: Control Charts Under Independence and Dependence

by

Darren Williams

(Under the direction of Dr. Lynne Seymour)

Abstract

Control charts were devised to evaluate offices within the Intel Corporation. These control charts were used to determine whether or not that particular office had too many expense report errors. If so, this would indicate a need to conduct an investigation. Therefore, a set of statistical limits, as provided by the control chart, must be used to detect these occurrences. Several different control charts are discussed with this analysis, with the most emphasis placed on the EWMA control chart. Problems that then arose when the data exhibited autocorrelations were also examined. Control charts under both dependence and independence are examined and applied to the Intel data with examples and conclusions presented.

Index words: control charts, dependence, EWMA, EWMAST

Monitoring Expense Report Errors: Control Charts Under Independence and Dependence

by

Darren Williams

B. A.

Albany State University, 2002

A Thesis Submitted to the Graduate Faculty

of The University of Georgia in Partial Fulfillment

of the

Requirements for the Degree

Master of Science

Athens, Georgia

2004

Monitoring Expense Report Errors: Control Charts Under Independence and Dependence

by

Darren Williams

Approved:

Major Professor: Dr. Lynne Seymour

Committee: Dr. William P. McCormick
Dr. T. N. Sriram

Electronic Version Approved:

Maureen Grasso
Dean of the Graduate School
The University of Georgia
May 2004

Acknowledgments

First, I would like to thank God, whose many blessings has provided me with the knowledge and strength to achieve tremendous heights both academically and spiritually.

Next, I would like to thank Dr. Lynne Seymour, for all the direction, instruction and encouragement that she has given me as my major professor throughout this thesis experience. I would also like to thank Dr. T.N. Sriram and Dr. William McCormick for being helpful members in my committee.

I would also like to thank Ms. Connie Durden for her assistance in helping me putting this thesis together.

And last, but not least, I would like to thank my loving family. In particular my mother, whose hard work and encouragement helped to shape me in being the person that I am today.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The Intel Data

The data come from 126 different offices within the Intel Corporation. For each of these offices the number of expense report errors is recorded over the three-year period from 2000 to 2002. The errors are classified into two categories: travel errors and non-travel errors, with travel errors being the more severe of the two.

Table 1.1: Data Description

STE_CD	REV_NM	INTEL_WW	INTEL_MO	INTEL_QTR	INTEL_YR
A01	Travel	WW 34	8 (August)	Quarter 3	2000
A01	Travel	WW 42	10 (October)	Quarter 4	2000
A01	Travel	WW 47	11 (November)	Quarter 4	2000
A01	Travel	WW 50	12 (December)	Quarter 4	2000
A01	Travel	WW 17	4 (April)	Quarter 2	2001
A01	Travel	WW 23	6 (June)	Quarter 2	2001
A01	Travel	WW 27	7 (July)	Quarter 3	2001
A01	Travel	WW 35	9 (September)	Quarter 3	2001
A01	Travel	WW 36	9 (September)	Quarter 3	2001
A01	Travel	WW 37	9 (September)	Quarter 3	2001
A01	Non-Travel	WW 37	9 (September)	Quarter 3	2001

Note: For security purposes, this description is a representative disguise of the data.

Each office is represented by a site code under the column heading STE_CD. Column REV_NM gives the type of reported error. The time of the error is reported in four formats: INTEL_WW, which is the workweek in the particular year that the error occurred; INTEL_MO, which is the month that the error occured; INTEL_QTR, which is the quarter of the year that the error occurred; and finally INTEL_YR, which is the year that the error occurred. Because non-travel expense report errors are considered less severe than travel errors, they are given a weight of 0.5 or half that of a travel error. The cumulative number of errors per office range from only one to a maximum of 106.

The objective is to devise a statistical method for determining if an office has enough errors to warrant investigation. Therefore, there has to be a set of statistical limits that alerts the main office when a limit has been exceeded; in that case the offending office will be audited to determine the cause of the problem. This objective can be attained through the use of control charts.

1.2 Significance of Control Charts

What is quality control? Quality can be defined as fitness for use. Fitness can be divided into two categories, which are quality of design, the intentional variation of a product; and quality of conformance, how well the product conforms to the required specifications. A more modern definition describes quality as being inversely proportional to variability. Control, according to Shewhart (1931), "is a controlled phenomenon when, through the use of past experiences, we can predict, at least within limits, how the phenomenon may be expected to vary in the future. Here it is understood that prediction means that we can state, at least approximately, the probability that the observed phenomenon will fall within the given limits".

Therefore, quality control can be defined as a means of reducing the variability of a product through the use of past experiences to predict future occurrences. The purpose of quality control is to produce quality improvement, which is the reduction of variability in processes and products. It should be noted that quality improvement does not include the inherent variation of all products or processes (chance variations) but those that are not part of chance variation (*assignable causes*).

Statistical Process Control (SPC) is used to achieve process stability and improve the capability of a product or process through the reduction of variability. SPC is a collection of primarily seven problem-solving tools that can be used to achieve this goal. Of the so-called "magnificent seven" which includes the following: Histogram or stem-and-leaf display, check sheet, pareto chart, cause and effect diagram, defect concentration diagram, scatter diagram, and the control chart; the control chart is the most important and widely used.

Dr. Walter Shewhart first introduced the theory of control charts in 1931. The control chart contains a *centerline*, which represents the average value, or target value, of a quality characteristic corresponding to the *in-control* state (that is, only chance causes are present in an in-control process). The chart also contains two other horizontal lines, the *upper control limit* (UCL) and the lower control limit (LCL). They are chosen so that nearly all of the sample points fall between them if the process is in-control. However, if a point plots outside of either control limit or if the plots behave in a nonrandom manner the process is said to be out-of-control.

The major objectives of the control chart are to alert the user when a process is out-of-control so that the process may be improved. Process improvement is achieved when, if a process is deemed out-of-control, assignable causes are investigated and eliminated. These concepts and others will be examined throughout this discussion. We will begin by examining some basic concepts of the Shewhart control charts followed by more complex ones such as the EWMA and CUSUM control charts, which will be used to analyze the Intel data.

1.3 General Types of Control Charts

1.3.1 Control Charts for Variables

If the quality characteristic one is interested in controlling is quantitative, it is called a *variable*. Control charts for central tendency and variability are called *variables control charts*. Shewhart control charts (Shewhart, 1931) are used to monitor the process mean and variability and to alert the user when the process mean shifts to another level. There are two methods by which these charts are developed: the first of these is developing the control chart with known values of the process mean and standard deviation, μ and σ; and the second develops the chart through estimation of these parameters.

Control charts for $\bar{x}$, which follows the process mean, and R, which follows the process variability using the range, are developed under the assumption that the data follows a normal distribution. If μ and σ are known and come from a process that is normally distributed, and if $x_1, x_2, \ldots, x_n$ is a sample of size n, then the average of this sample is

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

where $\bar{x}$ is normally distributed with mean μ and standard deviation $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$. Therefore, if μ and σ are known, we have a 100(1-α%) chance that our sample mean will fall between these upper and lower confidence limits.

$$\begin{array}{ll} \text{Upper:} & \mu + Z_{\frac{\alpha}{2}}\dfrac{\sigma}{\sqrt{n}} \\ \text{Lower:} & \mu - Z_{\frac{\alpha}{2}}\dfrac{\sigma}{\sqrt{n}}. \end{array}$$

These upper and lower confidence limits can be used as upper and lower control limits for monitoring the process mean with the centerline being the process mean. Therefore the control chart would be of the form

$$\begin{aligned} UCL &= \mu + Z_{\frac{\alpha}{2}}\frac{\sigma}{\sqrt{n}} \\ CL &= \mu \\ LCL &= \mu - Z_{\frac{\alpha}{2}}\frac{\sigma}{\sqrt{n}} \end{aligned} \tag{1.1}$$

The assumption of normality is often not a valid one in practice. Therefore, we rely on the Central Limit Theorem that says that the distribution of averages from samples of size n taken from a stable process is approximately normal, with mean μ and standard deviation $\frac{\sigma}{\sqrt{n}}$ as $n \to \infty$. The $\frac{\alpha}{2}$ point of the normal distribution is commonly replaced with 3 in the above formula, since a standard normal random variable Z will almost always be between -3 and 3. Using these values of Z we find that $P(-3 < Z < 3) = .9973$.

In practice it is very rare that μ or σ are known. As a result, they must be estimated from preliminary samples or subgroups, typically 20 to 25, taken when the process is thought to be in control. Suppose there are m samples each containing n observations of the specified quality characteristic. Let $\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_m$ be the average of each of the samples. Then, a good estimator of the process average is the grand average

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

This value will become the centerline on the control chart.

Now an estimate of the standard deviation must be found. This is done using either the sample standard deviation method or the sample range method. We will first examine the latter. If we have a sample of size n, $x_1, x_2, \ldots, x_n$ then the range of that sample would be

$$R = x_{\max} - x_{\min}.$$

Let $R_1, R_2, \ldots, R_m$ be the ranges of m samples. Then the average range is

$$\bar{R} = \sum_{i=1}^{m} R_i.$$

The control limits on the $\bar{x}$ chart are then

$$\begin{aligned} UCL &= \bar{\bar{x}} + a\bar{R} \\ CL &= \bar{\bar{x}} \\ LCL &= \bar{\bar{x}} - a\bar{R}, \end{aligned} \tag{1.2}$$

where a is tabulated for values of m and n in Chandra (2001).

We may also monitor the process variability by plotting values of the sample range on a control chart. These calculations are:

$$\begin{aligned} UCL &= b_1\bar{R} \\ CL &= \bar{R} \\ LCL &= b_2\bar{R} \end{aligned} \tag{1.3}$$

where again, b_1 and b_2 are tabulated for various values of m and n in Chandra (2001)

As seen above, when values for the mean and standard deviation are known we can construct a control chart for the process mean. We may also construct a control chart to monitor process variance based on the sample range. To construct this chart we must first recall that $\hat{\sigma} = \frac{R}{d_n}$, where d_n is the mean of the distribution of the relative range. The standard deviation is $\sigma_R = d\sigma$, where d is the standard deviation of the distribution of the relative range. This gives us the following chart to monitor the process's standard deviation.

$$\begin{aligned} UCL &= d_n\sigma + 3d\sigma \\ CL &= d_n\sigma \\ LCL &= d_n\sigma - 3d\sigma \end{aligned} \tag{1.4}$$

Define the constants

$$\begin{aligned} D_1 &= d_n - 3d \\ D_2 &= d_n + 3d. \end{aligned}$$

Then the parameters of the R chart for a known value of the standard deviation given becomes

$$\begin{aligned} UCL &= D_2\sigma \\ CL &= d_n\sigma \\ LCL &= D_1\sigma \end{aligned} \tag{1.5}$$

where the constants D_1 and D_2 are tabulated in Ledolter and Burrill (1999).

Using the range method to estimate the standard deviation is useful primarily because of ease of calculation. However, with today's use of computers this is no longer a consideration. The sample standard deviation method, as one might expect, is preferable in estimating standard deviation. Although it can be shown that for sample sizes of $n = 4, 5$, or 6, the sample range and sample standard deviation methods are similar, Derman & Ross (1997) point out that $\frac{\bar{R}}{d_n}$ is not a reasonable estimator of the standard deviation when the underlying distribution is not normal.

The construction of control charts based on $\bar{x}$ and S, the sample standard deviation, is somewhat different than the $\bar{x}$ and R charts. First, for each sample we must calculate the sample average $\bar{x}$ and the sample standard deviation S. We know that if σ^2 is the unknown variance of a probability distribution, then the unbiased estimator of σ^2 is the sample variance

$$S^2 = \frac{\sum_{i=1}^{n} x_i^2 - \frac{(\sum_{i=1}^{n} X_i)^2}{n}}{n-1}.$$

But, S_i is not an unbiased estimator of σ. If the underlying distribution is normal, then S actually estimates $c_n\sigma$, where c_n is a constant that depends on the sample size n. Therefore, the standard deviation of S is $\sigma\sqrt{1-c_n^2}$.

To construct these control charts we will consider two cases. The first is the case when a standard value of σ is given. The chart will be of the form

$$\begin{aligned} UCL &= c_n\sigma + 3\sigma\sqrt{1-c_n^2} \\ CL &= c_n\sigma \\ LCL &= c_n\sigma - 3\sigma\sqrt{1-c_n^2}. \end{aligned} \tag{1.6}$$

Define the constants

$$\begin{aligned} B_1 &= c_n - 3\sqrt{1-c_n^2} \\ B_2 &= c_n + 3\sqrt{1-c_n^2}. \end{aligned}$$

Therefore, the parameters of the S chart with a standard value for σ given becomes

$$\begin{aligned} UCL &= B_1\sigma \\ CL &= c_n\sigma \\ LCL &= B_2\sigma \end{aligned} \tag{1.7}$$

where B_1 and B_2 are tabulated for various sample sizes in Chandra (2001).

Equation 1.1 would still be used to monitor the process mean.

The second case is when there is no known value for σ. It must then be estimated by analyzing past data. If m samples are given, each of size n, then S_i is the standard deviation of the ith sample. The average of these m samples is

$$\bar{S} = \frac{1}{m}\sum_{i=1}^{m} S_i.$$

Now, we have $\frac{\bar{S}}{c_n}$ as an unbiased estimator of σ. The chart will be of the form

$$\begin{aligned} UCL &= \bar{S} + 3\frac{\bar{S}}{c_n}\sqrt{1-c_n^2} \\ CL &= \bar{S} \\ LCL &= \bar{S} - 3\frac{\bar{S}}{c_n}\sqrt{1-c_n^2} \end{aligned} \tag{1.8}$$

We can define these constants

$$\begin{aligned} B_3 &= 1 - \frac{3}{c_n}\sqrt{1-c_n^2} \\ B_4 &= 1 + \frac{3}{c_n}\sqrt{1-c_n^2} \end{aligned}$$

Then for the S chart we have

$$\begin{aligned} UCL &= B_4\bar{S} \\ CL &= \bar{S} \\ LCL &= B_3\bar{S}. \end{aligned} \tag{1.9}$$

where the constants B_3 and B_4 are tabulated in Chandra (2001). The corresponding $\bar{x}$ chart is of the form

$$\begin{aligned} UCL &= \bar{\bar{x}} + a\bar{S} \\ CL &= \bar{\bar{x}} \\ LCL &= \bar{\bar{x}} - a\bar{S} \end{aligned} \tag{1.10}$$

where the constant $a = \frac{c_n}{\sqrt{n}}$. The origins of the constant involved may be found in Bowker and Lieberman (1972).

Also, some have suggested using $\sqrt{\frac{\sum_{i-1}^{k} S_i^2}{k}}$, the average standard deviation of k subgroups, as a better estimator of σ. For a discussion of this method see Derman and Ross (1997).

The interpretation of the preceding control charts is done primarily by determining if one or more of the points fall outside of the control limits. This is not the only indication of an out-of-control process. If the pattern of the points exhibits any nonrandom or systematic behavior, the process is said to be out-of-control. For a better discussion of the interpretation of these patterns see the Western Electric Statistical Quality Control Handbook (1956).

It should also be noted that the control chart that monitors the standard deviation, R or S control charts, should be examined before we examine the corresponding $\bar{x}$ control chart since the R or S control chart must first be in control for the $\bar{x}$ control chart to have any meaning. Therefore, if the R or S control chart is out-of-control, those assignable causes must be eliminated before going on to the $\bar{x}$ control chart.

The ability of the $\bar{x}$ control chart to detect a shift in the process mean can be described by their operating characteristic (OC) curves. If a standard value for σ is known and constant, the OC curves can be constructed for the $\bar{x}$ chart. Suppose the mean shifts from some in-control value μ to another value $\mu + k\sigma$. Then, the *β-risk* or the probability of not detecting this shift on the first subsequent sample is

$$\beta = P\lfloor LCL \leq \bar{x} \leq UCL | \mu = \mu_0 + k\sigma \rfloor.$$

Given that $\bar{x} \sim N(\mu, \frac{\sigma}{\sqrt{n}})$,

$$\beta = \Phi(L - k\sqrt{n}) - \Phi(-L - k\sqrt{n})$$

where $L = 3$, the usual 3-sigma limits and Φ denotes the standard normal cumulative distribution function. The probability of detecting the shift on the rth sample is

$$\beta^{r-1}(1 - \beta).$$

The expected number of samples required before the shift is detected is the Average Run Length (ARL), which equals $\frac{1}{1-\beta}$. Therefore the out-of-control ARL, denoted ARL_1, is of the form

$$ARL_1 = \frac{1}{1-\beta}$$

which gives us the number of samples required in order for a shift in the mean to be detected given that the system is out-of-control. Likewise, we can also calculate the ARL for the in-control process; this ARL is of the form

$$ARL_0 = \frac{1}{\alpha}.$$

This gives us the number of samples, on average, in which an out-of-control signal occurs even when the process is really in-control.

Although the $\bar{x}$, R, and S control charts are the most common control charts applied to variable data, there are others that deserve attention. The first of these is the S^2 control chart, which is based directly on the sample variance. The S^2 control chart has the form

$$\begin{aligned} UCL &= \frac{\overline{S^2}}{n-1}\chi^2_{\frac{\alpha}{2}n-1} \\ CL &= \overline{S^2} \\ LCL &= \frac{\overline{S^2}}{n-1}\chi^2_{1-(\frac{\alpha}{2}),n-1} \end{aligned} \tag{1.11}$$

where $\chi_{\frac{\alpha}{2}^2}$ and $\chi^2_{1-(\frac{\alpha}{2})n-1}$ are the upper and lower $\frac{\alpha}{2}$ percentage points of the chi-square distribution with $n-1$ degrees of freedom. The control limits are determined using the fact that the distribution of $(n-1)S_i^2/\sigma^2$ is chi-squared with $n-1$ degrees of freedom. This fact is also the reason behind $E[S_i] = c_n\sigma$ from the previous section and that $\frac{\bar{S}}{c_n}$ is an unbiased estimator of σ. A known value σ^2 could replace the $\bar{S}^2$ in Equation 1.11 if available. One should also note that Equation 1.11 could be used to get the false alarm probability of an S chart in agreement with that of a $\bar{x}$ chart. Recall that the probability of a false alarm (the probability of having a value outside the control limits when the process is in control) for an $\bar{x}$ chart with 3-sigma limits is .0027. The same cannot be said of the S chart. Even if the data comes from a normal distribution the sample standard deviation will not be normally distributed.

For instance, it can be shown that for a subgroup size of $n = 5$ the probability of a false alarm equals .0040. This implies that our ARL_0, or average number of samples for an out-of-control signal to occur even if the process is in control, for the S-chart with $n = 5$ is $\frac{1}{.0040} = 250$ compared to the value of $\frac{1}{.0027} = 370$ one might have expected with the $\bar{x}$ chart. Using Equation 1.11 we can choose the in-control ARL, 370 say, and then choose the control limits to meet this goal. For $0 < \alpha < 1$, let $\chi^2_{n,\alpha}$ be such that $P\{\chi^2_n > \chi_{n,\alpha^2}\} = \alpha$. Therefore, if we want the false alarm probability to be .0027, or equal probabilities of .00135 of having an S-value above and below the control limits, we can use the fact that

$$P\{\chi^2_n > \chi_{n,.00135}\} = .00135$$

and

$$P\{\chi^2_n < \chi^2_{n,.99865}\} = .00135.$$

If the system is in control throughout the process with probability of .0027 it will be true that

$$\chi^2_{n-1,.99865} < (n-1)\frac{S_i^2}{\sigma^2} < \chi_{n-1,.00135}$$

Therefore our limits on the S^2 chart would now become

$$\begin{aligned} LCL &= \frac{\overline{S^2}}{n-1}\chi^2_{n-1,.99865} \\ UCL &= \frac{\overline{S^2}}{n-1}\chi^2_{n-1,.00135} \end{aligned} \tag{1.12}$$

The same rules apply for interpretation of the S^2 as for the previous charts.

The second control chart we will examine in this section is an individual measurement control chart. This chart deals with situations where the sample size used to monitor the process is $n = 1$. These charts are primarily used in situations where the repeated measurements on the process differ only because of analysis error, such as those of the chemical industry, and in industries where every unit manufactured are inspected. This procedure uses the moving range of two successive observations to estimate the process variability. The individual measurement chart is of the form

$$\begin{aligned} UCL &= \bar{x} + \frac{3}{d_n}\overline{mr} \\ CL &= \bar{x} \\ LCL &= \bar{x} - \frac{3}{d_n}\overline{mr} \end{aligned} \tag{1.13}$$

where $mr_i = |x_i - x_{i-1}|$ and d are tabulated in Montgomery (1997) for different values of n. Note that n denotes the number of observations for which the moving average is calculated. A control chart on the moving range may also be constructed. This chart is of the form

$$\begin{aligned} UCL &= D_4\overline{mr} \\ CL &= \overline{mr} \\ LCL &= D_3\overline{mr} \end{aligned} \tag{1.14}$$

where D_3 and D_4 are tabulated in Montgomery (1997) for different values of n. The interpretation of the individual measurement chart is interpreted in the same manner as an ordinary $\bar{x}$ chart. The interpretation of the moving range chart is more difficult to interpret because the moving ranges are correlated, and this correlation may exhibit a pattern of runs or cycles on the chart. The individual measurements on the $\bar{x}$ chart are assumed to be uncorrelated, but any pattern on this chart should still be investigated.

The variable control charts that we have discussed, $\bar{x}, R, S, S^2$ and individual measurement, were not used in the analysis of the Intel data. The primary reason for this is because these charts require a target value μ and σ and if those are not available, they may be estimated from a process that was thought to be in control. In this case, target values are unknown, and there is no past in-control process from which these parameters could be estimated. The Individual Measurement control chart could not be used for this reason also, in addition to the fact that this chart.s performance is greatly affected if the process is not normally distributed.

1.3.2 Control Charts for Attributes

There are often occasions where the quality characteristic of interest cannot be represented numerically. Instead the quality characteristic is qualitative, and can be described as either conforming or nonconforming. Characteristics of this type are called attributes, and control charts may be derived for attribute data.

The first attribute chart to be examined is the control chart for the fraction nonconforming, or the p chart. When a system is in control, the probability that each item processed will be independently defective is p. If we let X denote the number of defectives items in a random sample of n items then, assuming the system is in control, X will be a binomial random variable with parameters n and p:

$$P\{X = x\} = \begin{pmatrix} n \\ x \end{pmatrix} p^x (1-p)^{n-x} \quad x = 0, 1, \ldots, n$$

with

$$E[X] = np \text{ and } Var(X) = np(1-p).$$

The sample fraction nonconforming (the ratio of the number of nonconforming units to the sample size n) in a sample is therefore,

$$\hat{p} = \frac{X}{n}.$$

The distribution of $\hat{p}$ can be obtained from the binomial. The mean and variance are

$$\mu = p$$

and

$$\sigma^2_{\hat{p}} = \frac{pq}{n}$$

where $q = 1 - p$.

If a standard value is known or specified we can use the preceding equations to construct a control chart. Therefore, the p chart for fraction nonconforming is of the form

$$UCL \; = \; p + 3\sqrt{\frac{pq}{n}}$$

$$\begin{aligned} CL &= p \\ LCL &= p - 3\sqrt{\frac{pq}{n}}. \end{aligned} \tag{1.15}$$

The operation of the chart involves plotting the $\hat{p}$ statistic on the control chart. If a standard value is not given for p, then it must be estimated from observed data. We first select m preliminary samples, each of size n. If there are X nonconforming units in sample i, then the fraction nonconforming in the ith sample is $p_i = \frac{X_i}{n}$ $i = 1, \ldots, m$, and $\bar{p} = \frac{\sum_{i=1}^{m} \hat{p}_i}{m}$. This $\bar{p}$ is used to estimate the unknown p. We now have a p chart of the form

$$\begin{aligned} UCL &= \bar{p} + \sqrt{\frac{\bar{p}\bar{q}}{n}} \\ CL &= \bar{p} \\ LCL &= \bar{p} - \sqrt{\frac{\bar{p}\bar{q}}{n}}. \end{aligned} \tag{1.16}$$

These control limits should be regarded as trial control limits. The trial control limits should be used to determine if the process was in control when the preliminary data was collected. If there are any points that exceed the control limits, these points should be investigated. If assignable causes are discovered, these points should be discarded and new control limits should be calculated. One should note that these trial control limits are not necessary if a standard value is given for p.

In many situations it is not convenient to consider subgroups of fixed size. Therefore, our control chart will be based on the number of nonconforming units as opposed to the fraction nonconforming. This chart is often called an np control chart. Suppose that when in control, each item will be defective independently with probability p. Then, if the subgroup is of size n, the np control chart is of the form

$$\begin{aligned} UCL &= np + 3\sqrt{np(1-p)} \\ CL &= np \\ LCL &= np - 3\sqrt{np(1-p)} \end{aligned} \tag{1.17}$$

If p is unknown it can be estimated by

$$\bar{p} = \frac{\sum_{i=1}^{m} \hat{p}_i}{m}$$

where m is the number of samples.

In some situations with the fraction nonconforming chart, there may be a variable sample size. Or, in general, a different number of units may be produced in each period. There are three main methods used to handle this situation; Constructing control limits based on the specific sample size, average sample size, and standardized control charts. For more on these methods see Montgomery (1997).

The O-C curves for the fraction nonconforming chart is found by calculating the probability of a Type II error, or probability of incorrectly declaring the process in control. These probabilities may be calculated from

$$\begin{aligned} \beta &= P\{\hat{p} < UCL|p\} - P\{\hat{p} \leq LCL|p\} \\ &= P\{D < nUCL|p\} - P\{D \leq nLCL|p\} \end{aligned}$$

This is obtained from the cumulative binomial distribution. The ARLs for the fraction nonconforming chart may also be calculated either directly from the binomial distribution or from an OC curve. The formulas

$$ARL_0 = \frac{1}{\alpha}$$

and

$$ARL_1 = \frac{1}{1-\beta}$$

are the same for any Shewhart control chart.

There are situations where there are nonconformities in a unit of product, but there still may not be enough nonconformities to be classified as nonconforming or inadequate. In these situations we would construct a control chart on the number of nonconformities per unit, rather than the fraction nonconforming. These are sometimes called c charts. If the probability of a nonconformity is small and constant, the number of opportunities for nonconformities is infinitely large, and the inspection unit is the same for each sample, then the data values will be approximately Poisson with mean $= \lambda$. As long as the departures from these conditions are not severe, the Poisson model is still appropriate. If X_i is the ith data value then the mean and variance equal λ when the process is in control. Also, when in control, each data value will be within 3 standard deviations with high probability (if $\lambda \geq 10$, then by the Central Limit Theorem the probability will be very close to .0027). Therefore, the control limit would be

$$\begin{aligned} UCL &= \lambda + 3\sqrt{\lambda} \\ CL &= \lambda \\ LCL &= \lambda - 3\sqrt{\lambda}. \end{aligned} \tag{1.18}$$

Note, that if these calculations lead to a negative value for LCL, then set LCL=0. When no standard is given we should estimate λ by

$$\bar{\lambda} = \sum_{i=1}^{m} \lambda_i / m$$

where m is the number of samples. The chart will now be of the form

$$\begin{aligned} UCL &= \bar{\lambda} + 3\sqrt{\bar{\lambda}} \\ CL &= \bar{\lambda} \\ LCL &= \bar{\lambda} - 3\sqrt{\bar{\lambda}}. \end{aligned} \tag{1.19}$$

As stated before, if there is no standard given for λ, then the control limits should be seen as trial control limits. Any point falling outside these limits should be examined. It must then be decided if the system is temporarily out of control or if no statistical control has been established. In the latter case, the control limits should be recalculated. If the mean number of defects per unit is small when the process is in control, then the units should be combined and the data should consist of the number of defects in a given number of units. The sum of these Poisson random variables remains a Poisson random variable with a larger mean

$n\lambda$. It is often useful to combine these units when the mean number of defects per unit is less than 25 when the system is in control. For a discussion on the advantages of combining these units see Derman and Ross (1997).

In products such as trucks, computers, and radios, there are often many different types of nonconformities. All nonconformities are not equally important. A product with several minor nonconformities may not be classified as nonconforming, but a product with one serious nonconformity may be deemed nonconforming. Therefore, there must be a system that classifies nonconformities by severity and assigns weights accordingly. One of the more common methods to deal with this product is to comprise a demerit scheme based on four classes of nonconformities, with class A being very serious, class B being serious, class C being moderately serious, and class D being minor. If each class is independent then a Poisson distribution models the occurrence of nonconformities in each class. Therefore, the number of demerits in the inspection unit is

$$d_i = 100C_{iA} + 50C_{iB} + 100c_{iC} + C_{iD}$$

Where c_{iA}, c_{iB}, c_{iC} and c_{iD} correspond to the number of Class A, Class B, Class C, and Class D nonconformities, respectively, in the ith inspection unit. Though this is a common weighting method, it is not the only method. Any reasonable set of weights will work for one's specific problem.

The attribute control charts (p, np ,c charts) were also not used in the analysis of the data because in order to construct these charts one needs to know the total number of items in the population. In this case, the population is the total number of possible errors, not the number of expense reports (which may contain multiple errors). Clearly, such information is unavailable, but these charts would be ideal if it were.

1.3.3 Other Control Charts

The Shewhart control charts and the CUSUM and EWMA charts to be discussed in Chapter 2 are the most popular and widely used. But, there have been recent developments added to improve or modify these charts to fit other situations of interest. Among these are control charts for short production runs and multivariate control charts.

The control chart for short production runs can be applied using the standard $\bar{x}$ and R charts with some modifications. These charts use as their basis the deviation from the target value as the variable on the control chart. These charts rely on assumptions such as: the process standard deviation is approximately the same for all parts; and the process will work best when the sample size is constant for all part numbers. When the first assumption is violated one should use standardized $\bar{x}$ and R charts. Attribute control charts for short production runs are handled similarly. Again, standardized control charts are used instead of the common control charts for attributes. For more information on control charts for short production runs see Farnum (1992).

There are situations in which a product has many different features that, when put together, constitutes a good working product. In such cases it is possible to monitor these processes simultaneously. Monitoring the processes simultaneously is favored over monitoring the processes independently because independent monitoring can be misleading, since the probabilities of the Type I and Type II errors are distorted. In general, this distortion increases as the number of quality characteristics being monitored increases. If there are d independent quality characteristics and the probability of a Type I error equals α on each chart, then the true probability of a Type I error for the joint procedure is

$$\alpha_1 = 1 - (1 - \alpha)^d$$

and the probability that all d means simultaneously plot inside the control limits given that the process is in control is

$$(1 - \alpha)^d.$$

Clearly this distortion can be severe.

This difficulty was handled with the introduction of the multivariate control charts which were first introduced by H. Hotelling (1947). Some of the more useful multivariate control charts are the Chi-square and the Hotelling T^2 control charts. The Chi-square control chart assumes that two quality characteristics have a bivariate normal distribution and uses other parameters to create a chi-square distribution, from which the control chart is constructed. A discussion on the Chi-square control procedure can be seen in Nelson (1987). The Hotelling T^2 control charts are just a variation of the Chi-square procedure. Other multivariate methods involving the EWMA and the CUSUM are also in wide use. See Crosier (1988), Lowry et al. (1992), and Alt (1985) for discussions on these charts.

One of the disadvantages of the Shewhart control charts is that they only consider information from the last plot and not information from subsequent plots. Thus, these charts ignore information given by the entire sequence of points. As a result, Shewhart charts are less sensitive to small process shifts of $\leq 1.5\sigma$, say. In the next section we will consider control charts which consider all points and not just points in relation to those around them. The two charts we will use as an alternative to the Shewhart chart are the cumulative-sum (CUSUM) and the exponentially weighted moving-average (EWMA) control chart.

Chapter 2

Control Charts for the Intel Data under Independence

2.1 Exponentially Weighted Moving Average Control Chart

Earlier we pointed out that the Shewhart control charts were not effective in detecting small shifts in the mean. One control chart commonly used to detect small shifts in the process mean is the exponentially weighted moving average (EWMA) control chart.

S.W. Roberts first introduced the EWMA control chart in 1959. The EWMA chart reacts to small shifts in the mean by using all of the information in the data set, as opposed to only the preceding data value. The EWMA statistic for the ith observation x_i, is

$$Z_i = \lambda x_i + (1 - \lambda) Z_{i-1} \tag{2.1}$$

where λ is called the smoothing constant and lies between 0 and 1 and $z_0 = \mu_0$, which is the in-control mean (or the estimated mean, sometimes the average $\bar{x}$ of preliminary data). The statistic plotted on this chart, Z_i, is a weighted average of all previous observations. To see this, write z_i as

$$\begin{aligned}
Z_i &= \lambda X_i + (1-\lambda)[\lambda X_{i-1} + (1-\lambda) Z_{i-2}] \\
&= \lambda X_i + \lambda(1-\lambda) X_{i-1} + (1-\lambda)^2 Z_{i-2} \\
&= \lambda X_i + \lambda) X_{i-1} + (1-\lambda)^2 [\lambda X_{i-2} + (1-\lambda) \\
&= \lambda X_i + \lambda(1-\lambda) X_{i-1} + \lambda(1-\lambda)^2 X_{i-2} + (1-\lambda)^3 Z_{i-3} \\
&\vdots \\
&= (1-\lambda)^i Z_0 + W_1 X_1 + W_2 X_2 + \ldots + W_{i-1} X_{i-1} + W_1 X_1
\end{aligned}$$

Or in general

$$Z_i = \lambda \sum_{j=0}^{i-1} (1-\lambda)^j x_{i-j} + (1-\lambda)^i Z_0 \tag{2.2}$$

where the weight given to the jth observation x_j is $\lambda(1-\lambda)^j$.

Since $1-\lambda < 1$, the weights decrease exponentially with the age of the observation. If λ is small, then the weights decrease very slowly and the EWMA statistic is really an average of all the previous observations. For example if $\lambda = 0.3$, then the weight given to the current sample is 0.3 and the preceding weights are .21, .147, .1029, and so on. Note that these weights sum to one since

$$\begin{aligned}\lambda \sum_{j=0}^{i-1} (1-\lambda)^j &= \lambda[\frac{1-(1-\lambda)^i}{1-(1-\lambda)}] \\ &= 1-(1-\lambda)^i \end{aligned} \tag{2.3}$$

The EWMA chart is ideal for individual observations because it is very robust to the normality assumption. This robustness is due to the fact that z_t is the weighted average of past and current observations. For more on the robustness of the EWMA control chart to the normality assumption see Borror, et. al (1999).

If the observations x_i are independent and normally distributed with mean μ and variance σ^2, then z_i is the sum of independent normal variables and is itself normally distributed. Thus the mean of z_i is

$$\begin{aligned}E[Z_i] &= \mu[\lambda + \lambda(1-\lambda) + \lambda(1-\lambda)^2 + \ldots + \lambda(1-\lambda)^{i-1} + (1-\lambda)^i] \\ &= \mu\end{aligned}$$

and the variance of z_i is

$$\begin{aligned}Var[Z_i] &= \lambda^2\sigma^2[1 + (1-\lambda)^2 + (1-\lambda)^4 + \ldots + (1-\lambda)^{2(i-1)}] \\ &= \frac{\lambda\sigma^2[1-(1-\lambda)^{2i}]}{2-\lambda}\end{aligned}$$

Plotting z_i versus the sample number i (or time) and using the results for the expected mean and variance , the EWMA control chart is of the form

$$\begin{aligned} UCL &= \mu_0 + c\sigma\sqrt{\frac{\lambda}{(2-\lambda)}[1-(1-\lambda)^{2i}]} \\ CL &= \mu \\ LCL &= \mu_0 - c\sigma\sqrt{\frac{\lambda}{(2-\lambda)}[1-(1-\lambda)^{2i}]} \end{aligned} \tag{2.4}$$

where c is the width of the control limits. How to choose c will be discussed shortly. If no standard value is given for the mean so that it must be estimated by past data, replace μ_0 with $\bar{x}$. One should also note that if i is moderately large,

$$(1-\lambda)^{2i} \approx 0$$

and the quantity $[1-(1-\lambda)^{2i}]$ approaches unity as i gets larger. Therefore, after several time periods, the upper and lower control limits will be steady at

$$\begin{aligned} UCL &= \mu_0 + C\sigma\sqrt{\frac{\lambda}{(2-\lambda)}} \\ LCL &= \mu_0 - C\sigma\sqrt{\frac{\lambda}{(2-\lambda)}} \end{aligned} \tag{2.5}$$

The construction of the EWMA control chart requires that the smoothing constant λ and the control coefficient c be specified. This is accomplished using the ARL.

The ARL of the EWMA for independent sample averages can be expressed as the solution to an integral equation. If the EWMA starts at u then the ARL is given by

$$L(u) = 1 + \frac{1}{\lambda}\int_{LCL}^{UCL} L(y) f\{[y-(1-\lambda)\mu]/\lambda\}dy \tag{2.6}$$

where $f(x)$ is the $N(\mu, \sigma^2/n)$ density function, μ is the process mean, σ is the nominal process standard deviation, and n is the sample size for each sample. Crowder (1989) suggests using the ARL in the following steps in constructing the EWMA control chart.

1) Choose the smallest acceptable ARL for the case in which the process shift is zero (fixing the Type I error);

2) Decide the magnitude of the shift in the process to be detected, and choose λ to produces a minimum ARL for that shift size;

3) Once λ is chosen, find the control limit constant c which satisfies Step 1;

4) Perform a sensitivity analysis by comparing out-of-control ARLs for the optimal (λ, c) combination to other combinations of (λ, c) producing the same in-control ARL. From these choices pick the combination with the most desirable performance overall in terms of Type II errors.

In general, it is a good idea to use smaller values of λ to detect smaller shifts. Values of λ between 0.05 and 0.25 are the most commonly used in practice. Higher values of c work well with higher values of λ and lower values of c work well with low values of λ. Tables and graphs of the ARL for different values of λ and c can be seen in Lucas and Saccucci (1990) and Crowder (1989).

The EWMA can be altered in a manner that will allow for the monitoring of the process standard deviation. MacGregor and Harris (1993) developed a chart to monitor the process standard deviation, called the EWMS (Exponentially Weighted Mean Square). Assuming x_i is normally distributed with mean μ standard deviation σ, the EWMS is given by

$$\begin{aligned} S_n^2 &= \sum_{k=1}^{n} \lambda(1-\lambda)^{-k}[Y_k - \mu]^2 + (1-\lambda)^n S_0^2 \\ &= (1-\lambda)S_{n-1}^2 + r[Y_n - \mu]^2 \end{aligned} \tag{2.7}$$

where as before S_n^2/σ^2 is a weighted sum of chi-squared random variables, and is approximately distributed as $\chi^2(\nu)/\nu$ where the number of degrees of freedom ν depends upon λ, the correlation of the Y_k's, and the degrees of freedom associated with S_0^2. Because $E[\chi^2(\nu)] = \nu$ it follows that S_n^2 is an unbiased estimator of σ^2. If the observations are independent the degrees of freedom are given by $\nu = (2-\lambda)/\lambda$. Now if σ_0 is set equal to a target standard deviation, one can plot S_n on an EWRMS (Exponentially weighted root mean square) control chart with control limits given by

$$\begin{aligned} UCL &= \sigma_0 \sqrt{\frac{X^2 v/\alpha/2}{V}} \\ LCL &= \sigma_0 \sqrt{\frac{X^2 v/(1-\frac{\alpha}{2})}{V}}. \end{aligned} \tag{2.8}$$

MacGregor and Harris (1993) point out that the EWMS responds to changes in both the mean and standard deviation. They suggest that it is sometimes more useful to compute the EWMV (Exponentially Weighted Moving Variance). The EWMV is obtained by replacing μ in Equation 2-7 with a mean estimate, $\hat{\mu}_n$, at each point in time. In that case,

$$S_n^2 = (1-\lambda)S_{n-1}^2 + \lambda[Y_n - \hat{\mu}_n]^2 \tag{2.9}$$

The upper and lower control limits are of the form

$$\begin{aligned} UCL &= \sigma_0 + C_1\sigma_0 \\ LCL &= \sigma_0 - C_2\sigma_0 \end{aligned} \tag{2.10}$$

where c_1 and c_2 are the critical values calculated from the Johnson curve approximation or the $g\chi^2(\nu)$ approximation (Macgregor and Harris (1993)). They also deal with situations in which the observations are not independent, but are instead autocorrelated.

2.1.1 Cumulative-Sum Control Charts

Cumulative-sum charts were first introduced by E.S. Page in 1954. These charts plot the cumulative sums of the deviations of the sample value from a target value μ_0 by directly incorporating all the information in the sequence of sample values. This quantity is of the form

$$C_i = \sum_{j=1}^{i}(\bar{x}_j - \mu_0) \tag{2.11}$$

where $\bar{x}_j$ is the average of the jth sample.

If the process remains in control about a target value μ_0 then the cumulative sum of Equation 2.11 is a random walk with mean zero. But, if a shift from μ_0 to some value

$\mu_1 = \mu_0 + k$ occurs, then an *upward drift* is said to be present in the cumulative sum. Likewise if a shift from μ_0 to some value $\mu_1 = \mu_0 - k$ occurs, then we say that a *downward drift* will be present in the cumulative sum. Any trend upward or downward is a possible sign of an out-of-control system. There are two primary methods by which the CUSUM may be represented: the tabular CUSUM and the the V-mask CUSUM.

The tabular CUSUM accumulates deviations from μ_0 that are above and below target statistics C_i^+ and C_i^-, respectively. These statistics are

$$\begin{aligned} C_i^+ &= \max[0, X_i - (\mu_0 + K) + C_{i-1}^+] \\ C_i^- &= \max[0, (\mu_0 - K) - X_i + C_{i-1}^-] \end{aligned} \tag{2.12}$$

with C_i^+ and C_i^- initialized to 0. The value K is referred to as the *reference* or *slack value.* This value is often chosen halfway between the target value μ_0 and an out-of-control value μ_1 that one is interested in detecting. If the shift is expressed in standard deviations, such as $\mu_1 = \mu_0 + k\sigma$, then K can be expressed as

$$K = \frac{|\mu_1 - \mu_0|}{2} \quad \text{or} \quad \frac{K}{2}\sigma$$

This is done primarily because of the similarities between the CUSUM chart and the sequential probability ratio test (see Johnson (1961)). The system is considered out of control anytime either of the statistics C_i^+ or C_i^- exceeds the control limits of H, where H is chosen to be 5σ.

A useful feature of the CUSUM chart is that it indicates when the shift probably occurred. Let N^+ denote the number of consecutive times $C_i^+ > 0$ and let N^- denote the number of consecutive times $C_i^- > 0$. Then the system is said to have gone out of control at time N, where

$$N = \min\{N^-, N^+\}.$$

That is, if the system exceeds the control limit at some period i then the process probably went out of control at period $i - N$. We can also provide an estimate of the new mean μ_1, when the process has shifted. This calculation is

$$\hat{\mu} = \begin{cases} \mu_0 + K + \frac{C_i^+}{N^+}, & \text{if} \quad C_i^+ > H \\ \mu_0 - K - \frac{C_i^-}{N^-}, & \text{if} \quad C_i^- > H \end{cases} \tag{2.13}$$

It should also be noted that the run test or pattern test cannot be applied to the CUSUM chart because the plotted values are not independent.

The optimal design of the tabular CUSUM chart suggested by Gan (1991) requires that we specify the reference value K and the decision interval H. If $K = k\sigma, k$ should be chosen as to minimize the ARL_1 at the selected shift for a fixed value ARL_0. Usually k is chosen to be $\frac{1}{2}\delta$, where δ is the size of the shift in standard deviations. Once a value for k is chosen then we must choose h, where $H = h\sigma$,to give us the desired ARL_0. From there the ARL_1 for the (k, h) combinations are compared to other (k, h) combinations that produce the same ARL_1.

If one wants to detect a shift of one standard deviation, then values of $h = 5$ or $h = 4$ and $k = 0.5$ are recommended. One can also choose h and k to have an ARL_0 equal to that of Shewhart charts with three-sigma limits. For ARL values for different (h, k) combinations see Hawkins (1992) and Woodall and Adams (1993).

There are many methods by which the ARL of a CUSUM chart may be calculated. For a discussion of some of these procedures see Siegmund (1985) and Fellner (1990). Tables are also given for certain combinations of h and k in Lucas (1976). All of the calculations suggested by these authors are computationally intensive. Therefore, a very useful method suggested by Hawkins (1992) uses an ARL approximation equation. He argued that for the equation to provide a good fit to the ARL data, a transformation had to be done. The transformation he used was the inverse normal transformation, and the resulting approximation equation is

$$Y_{hk} = \alpha_h + \beta_k + \epsilon_h n_k + \epsilon_h^* n_k^* \tag{2.14}$$

where the coefficients $\alpha, \beta, \xi, \eta, \xi^*$ and η^* were fitted by weighted least squares. Therefore, to calculate the estimated ARL_0, the necessary values of $\beta_k, \eta_k, \eta_k^*, \alpha_h, \xi_h$, and ξ_h^* are obtained for the selected h and k. These values can be used to compute Equation 1-26. Finally the ARL_0 is computed by

$$ARL_0 = \frac{1}{\phi(-Y_{nk})} \tag{2.15}$$

where ϕ stands for the standard normal density function. His approximations were shown to be typically within 3% of the actual ARL when the process was in control and out-of-control. Note that he also gives the values with the head start feature, which will be discussed later. It should also be noted that the formula for finding the ARL here gives us the one-sided ARL. To get the two-sided ARL, divide the ARL by two. Though this method is useful, most statistical computer packages compute the ARLs for you.

The CUSUM chart may be adjusted for certain situations. One such adjustment involves using rational subgroups instead of the usual individual observation. This would be done by simply replacing x_i with $\bar{x}$ and σ with $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$ in Equation 1.24. This is done in situations where it is more beneficial to take more than one observation at a given time and get the average of those observations to represent the quality characteristic.

The use of subgroups instead of individual measurements does not usually work as well with CUSUM charts as it does with Shewhart charts. Another adjustment to the CUSUM chart is to use a combined CUSUM-Shewhart procedure to detect larger shifts. This process requires that the Shewhart control limits be located approximately 3.5σ from the centerline. This improves the ability of the chart to detect larger shifts while only slightly decreasing the ARL_0 (see Lucas (1982)). Finally, one may use a procedure introduced by Lucas and Crosier (1982) that gives the CUSUM chart a head start, or a fast initial response (FIR). The process involves setting the starting values C_0^+ and C_0^- to $H/2$ instead of zero. If the process starts in-control at the target value, the FIR feature has no effect - the CUSUMs will quickly decline to zero. But, if the process does not start in control the FIR allows us to detect it more quickly. For the ARLs of CUSUM charts with the head start feature see Lucas (1985).

So far in our discussion we have been focused on two-sided CUSUMS, but it is sometimes useful to only use one-sided CUSUMS. The calculation for this one-sided CUSUM is the same depending on whether one is interested in the mean shifting above or below some target value. However, the ARL calculations are not the same. The methods stated earlier may be applied to the one-sided ARLs also. In fact, the approximation method described by Hawkins (1992) actually gives us the ARL in terms of a one-sided CUSUM. As stated above, to get the two-sided CUSUM we just divide by two.

The second method used to represent a CUSUM is the V-Mask procedure. The V-mask is a graphical tool applied to successive values of the CUSUM statistic

$$C_i = \sum_{j=1}^{i} y_i = y_i + C_{i-1} \tag{2.16}$$

where y_j is the standardized observation $y_j = (x_i - \mu_0)/\sigma$. To determine if a process is out-of-control with the V-mask procedure one must first place the V-mask on the cumulative sum control chart with the point O on the last value of C_i and the line OP (line from the origin to the vertex) parallel to the horizontal axis. If any of the cumulative sums lie outside the arms of the mask, the process is considered to be out-of-control. The V-mask's performance depends on the lead distance d and the angle θ shown in Figure 2.1. Equivalence of the tabular CUSUM and the V-mask procedure can be achieved if

Figure 2.1: V-Mask

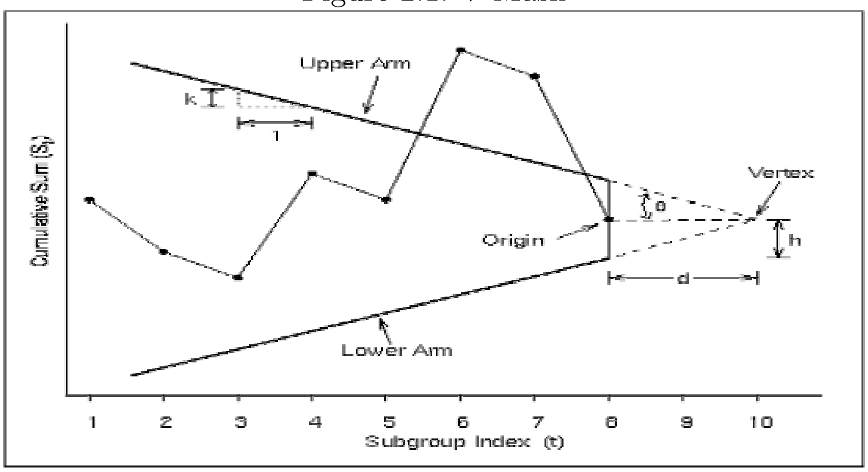

$$k = A \tan\theta$$

and

$$h = Ad\tan(\theta) = dk$$

Tabular CUSUM chart are usually preferred to the V-mask procedure primarily because interpretation of the V-mask procedure is difficult and one can only do a 2-sided scheme with the V-mask procedure. For more on the V-mask procedure see Lucas (1976).

CUSUM control charts for monitoring the process variability are also of interest. Hawkins (1993) suggests creating a standardized quantity

$$V_i = \frac{\sqrt{|y_i|} - 0.822}{0.349} \tag{2.17}$$

This quantity is derived first by letting x_i be normally distributed with mean μ_0 and standard deviation σ. The value of x_i is then standardized giving

$$y_i = \frac{(X_i - \mu_0)}{\sigma}$$

Given that the in-control distribution of ν_i is approximately normal with mean zero and standard deviation one, the two one-sided standard deviation CUSUMs are of the form

$$\begin{aligned} S_i^+ &= \max[0, V_i - K + S_{i-1}^+] \\ S_i^- &= \max[0, K - V_i + S_{i-1}^-] \end{aligned} \tag{2.18}$$

where S_0^+ and S_0^1 are set to 0 and h and k are selected as previously discussed for monitoring the process mean.

Chang and Gan (1995) suggest another scheme to monitor process variability. There CUSUM chart is based on the logarithmic transformation of the sample variance $\log(S^2)$. They argued that because $\log(S^2)$ is approximately normal the chart parameters could be obtained using tables constructed for CUSUM charts for monitoring the mean. Their control chart for monitoring the standard deviation is of the form

$$\begin{aligned} C_t &= \max\{0, C_{t-1} + y_t - K_C\} \\ D_t &= \min\{0, D_{t-1} + y_t + K_D\} \end{aligned} \tag{2.19}$$

where k_c and k_d are constants $y_t = \log(S_t^2)$, $C_0 = u$ for $0 \leq u < h_c$, and $D_0 = \nu$ for $-h_D < \nu \leq 0$. The one-sided upper CUSUM chart detects increases in σ^2 and issues an out-of-control signal whenever $C_t \geq h_c$. The one-sided lower CUSUM chart detects increases in σ^2 an issues a signal when $D_t \leq -h_D$. Chang and Gan (1995) also provide tables for the ARLs and compare them to the ARL of Shewhart and EWMA charts from the same type of standard deviation monitoring process.

2.2 Application

The major goal in the analysis of the Intel data described in Chapter 1 is to identify which of the 125 companies had an excessive amount of errors. Those companies with an excessive amount of errors should be investigated. The errors were examined by month. Because the data were compiled over three years, January 2000 is the 1st month, January 2001 is the 13th month, and January 2002 is the 25th month. Therefore, the EWMA control charts are based on 36 months. As mentioned earlier the EWMA and the CUSUM control charts are very insensitive to departures from normality, but for the sake of putting all numbers on equal footing the total number of errors in each month is divided by the number of days in that month to give an average monthly number of errors. For example, the total number of errors in January is divided by 31, February 2000 by 29, February 2001 by 28, March by 31, and so on. The mean and standard deviation were calculated using all 125 companies. The results are that the mean, or target value, equals 0.081646 and the standard deviation equals .047578. Therefore, z_0 in Equation 2.1 is initialized at 0.081646.

The parameters λ and c must be chosen before the analysis stage begins. As noted before, they should be chosen in order to achieve a desired ARL. Letting $\lambda = 0.1$ and $c = 2.7$. gives an in-control ARL of approximately 500 and an out-of-control ARL of 10.3.

(Note that all analysis was done with SAS statistical software.)

The first office to be examined is company 'AS1'. This company has the most expense report errors of all the companies with 106 over three years. The EWMA chart is given in Figure 2.2

Figure 2.2: Office 'AS1' EWMA

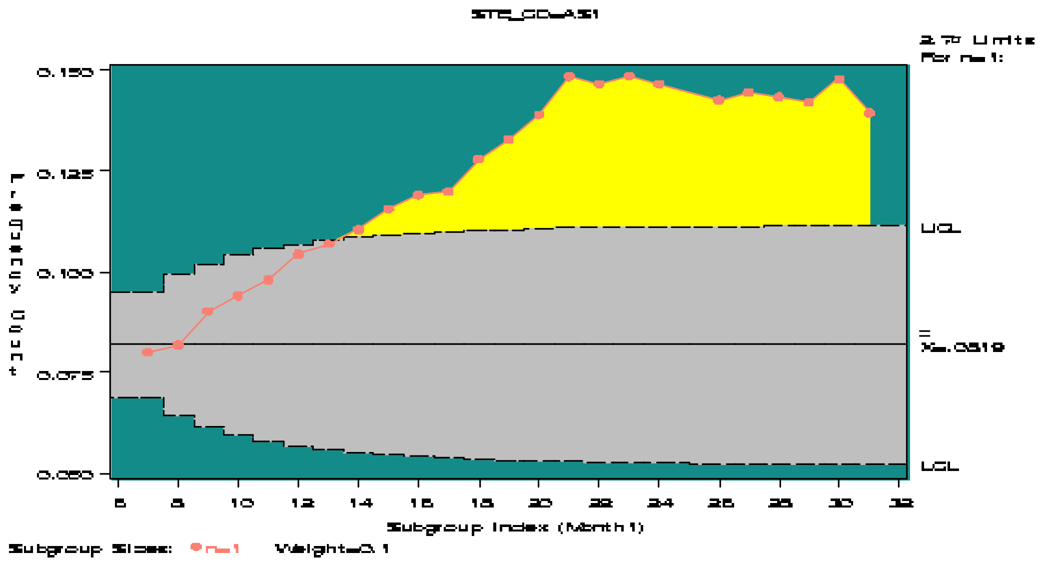

From this illustration it can be seen that this process is clearly out-of-control. It also shows a continuous upward trend until near October 2001 where it then levels off. The process seems to go out-of control around the 14th and 15th months. This company should definitely be investigated for assignable causes.

The next chart to be examined is that of company 'CR01'. This office has the second most expense report errors with 103. This chart is given in Figure 2.3

As was the case with office 'AS1' this process is obviously out-of-control. It also shows an upward trend before leveling off near the 22nd month. It seems to go out-of-control near months 12 and 13. Again, this company should definitely be investigated.

Figure 2.3: Office ‘CR01’ EWMA

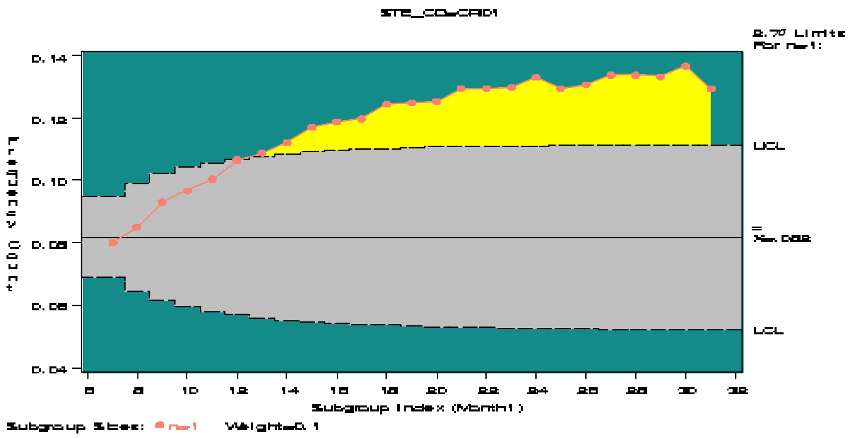

The next office examined is ‘WMNB’, which has 95 errors. It’s plot can be seen in Figure 2.4.

Figure 2.4: Office ‘WMNB’ EWMA

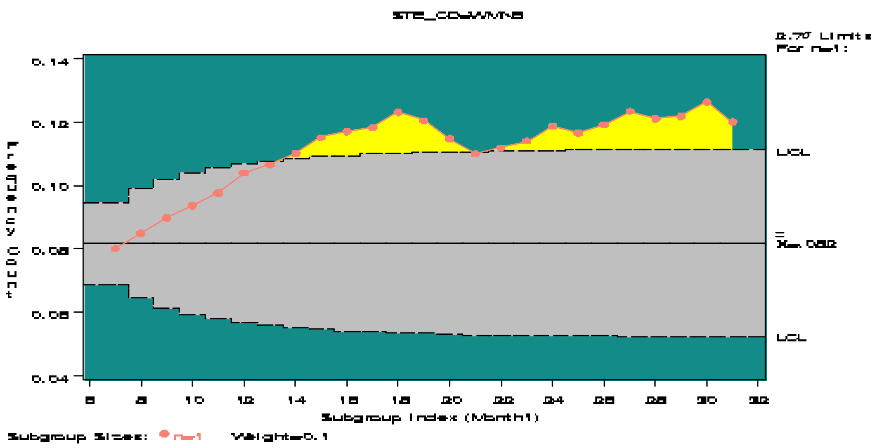

Again we see that this process is deemed out-of-of control.

The first three offices examined had an excessive amount of expense report errors. Therefore, the fact that they show an out of control process should come as no surprise, so now those offices with low to intermediate numbers errors will be examined.

The first of these we will examine is company 'AHE'. This company has 86.5 total expense report errors. Recall that a non-travel expense error is not as serious as a travel expense error, so it is counted as a half-error. This graph is given below in Figure 2.5

It can be seen that this process is also out-of control. It differs from the first two in that it shows fluctuations for a period of time followed by a trend upward. The process is signaled to be out-of-control near months 22 and 24. Again, an investigation is needed.

Figure 2.5: Office 'AHE' EWMA

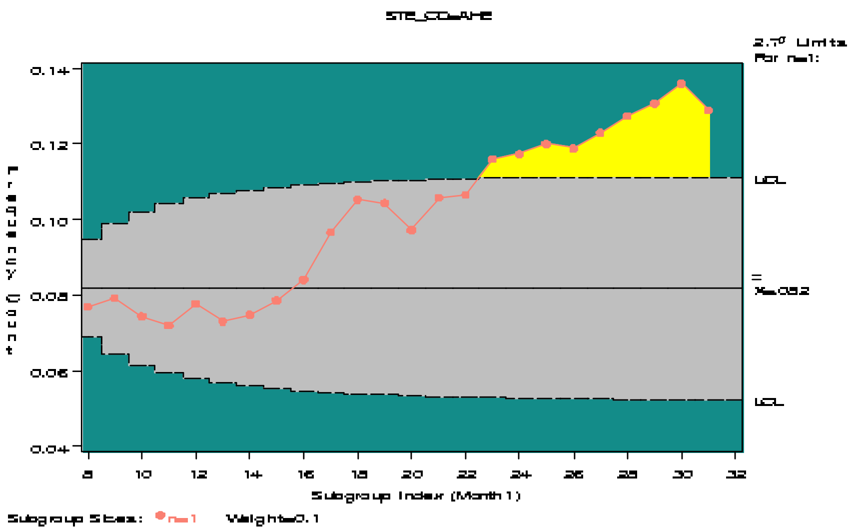

The next office is that of office 'T95-1' which has a total of 53 errors. It's graph is presented in Figure 2.6

This process is clearly in-control. All the values plot well within the control limits. There is no reason to investigate this company.

Office 'F18' with 34 errors fits into that low end of the intermediate category. This number of errors is expected to be acceptable. Examining Figure 2.7 we see that it is indeed an in-control process.

The offices seen so far that have out-of-control processes have multiple points that plot outside the control limits. But, for a process to be labeled out-of control only one point must plot outside these limits. This can be seen in Figure 2.7. This office, FSG', has a total of 52 errors.

Figure 2.6: Office ‘T95-1’ EWMA

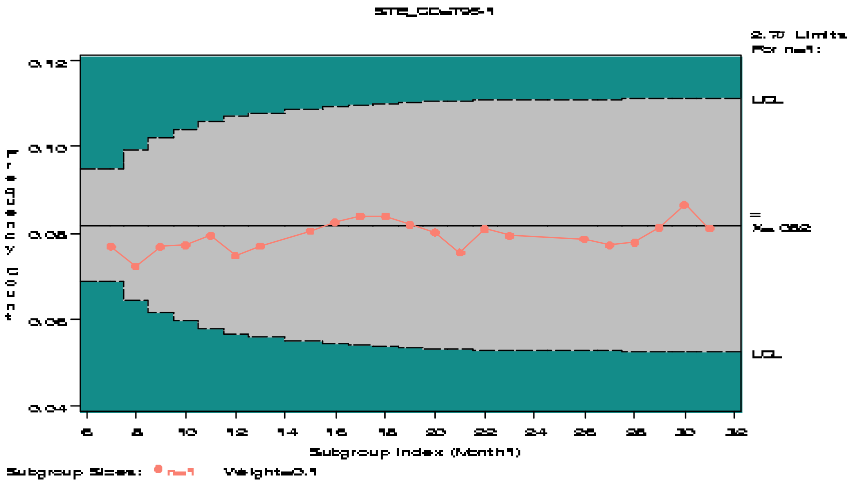

Though so far only points that exceed the upper control limit have been seen, it is possible for a point to exceed the lower control limit as it is in Figure 2.9. Figure 2.9 represents company ‘WEUZ’ which has 17 errors. Why is it a bad thing to exceed the lower control limits? The smallest amount of errors might be ideal! If there is a normal amount or slightly below normal amount of errors everything is fine. But, a very small number of errors should throw up an alarm: the data may be mistakenly reported, or, more seriously the data could be fraudently reported to deceive, making the office seem to be in compliance when it is truly not.

As mentioned earlier the EWMA control chart and the CUSUM control have very similar detection capabilities. Because the data will be analyzed using both CUSUM and EWMA, the chosen parameters should give approximately equal ARLs. To obtain an ARL for the CUSUM chart equivalent to that of the EWMA with $\lambda = 0.1$ and $L = 2.7$, the tabular CUSUM parameters must be set to $h = 5$ and $k = 0.5$. These parameters give an approximate out-of-control ARL of 500 and an approximate in control ARL of 10.3.

Figure 2.7: Office 'F-18' EWMA

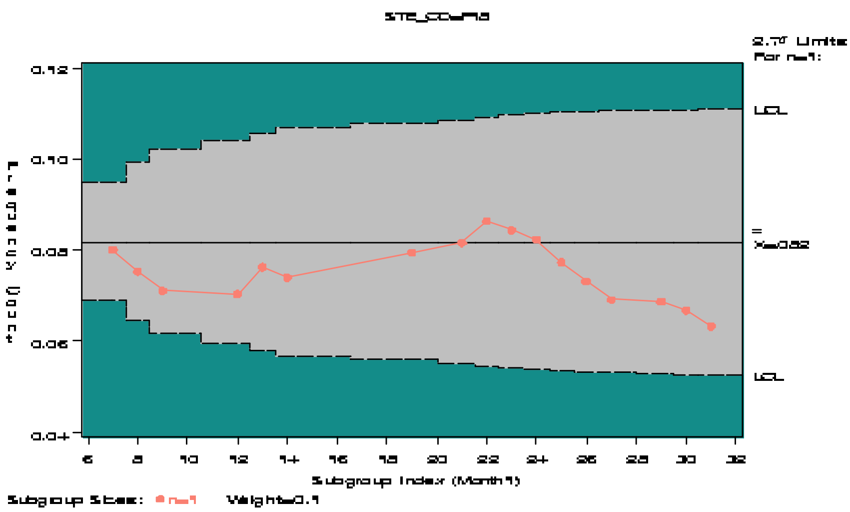

The tabular CUSUM in Table 2.1 is for company 'AS1' which was analyzed with the EWMA chart in Figure 2.2. Figure 2.2 shows that this office exceeded the upper control limit at around the 14th or 15th months and continues out-of-control from there on. Table 2.1 also shows that the upper control limit was exceeded at the 15th month and continues from there on. These two charts, as expected, are in agreement about the need for this company to be investigated.

Table 2.2 shows office 'WMNB'. This table shows that the process exceeded the upper control limit around month 14 or 15. This is also in agreement with Figure 2.4.

Finally, in Table 2.3 office 'T95-1' is seen to be in control, which is in agreement with the EWMA chart from Figure 2.6. Table 2.3 also doesn't show any points exceeding the upper or lower control limits.

Figure 2.8: Office ‘F-18’ EWMA

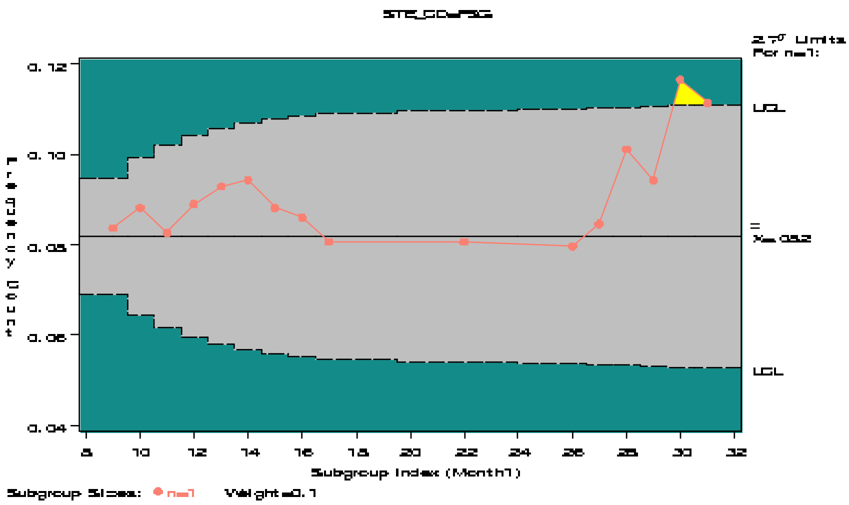

As noted earlier, when a process is out-of-control, the tabular CUSUM could indicate when the shift of the mean occurred that caused the process to go out of control. Recall, to accomplish this one must first examine the first period for which the upper or lower cusum decision interval was exceeded. Then one must examine the counters, N^+ or N^-, which record the consecutive periods since the CUSUM C_i^+ or C_i^- were above the value of zero. To see this, look at Table 2.1. It showed that the first period in which the upper decision interval was exceeded was period 9 (month 15). From Table 2.4, we see prior to period 9, the number of consecutive periods that the upper CUSUM was above zero is 7. The conclusion is then that the process was last in control at period $9 - 7 = 2$ (month 8), thus the shift likely occurred between months 8 and 9.

2.3 Alternative Methods of Analysis

As noted earlier the Intel data counted the number of expense report errors in a given workweek. Those errors were also represented with their respective month and quarter. The data could have been analyzed using any of the three time variables. Quarter was not chosen

Table 2.1: Office 'AS1'

Month1	Subgroup Sample Size	Individual Value	Cusum	Decision Interval	Decision Interval Exceeded
7	1	0.06451613	0.000000	5.0000	
8	1	0.09677419	0.000000	5.0000	
9	1	0.16666667	1.278810	5.0000	
10	1	0.12903226	1.768229	5.0000	
11	1	0.13333333	2.347865	5.0000	
12	1	0.16129032	3.513905	5.0000	
13	1	0.12903226	4.003324	5.0000	
14	1	0.14285714	4.782724	5.0000	
15	1	0.16129032	5.948764	5.0000	Upper
16	1	0.15000000	6.877986	5.0000	Upper
17	1	0.12903226	7.367406	5.0000	Upper
18	1	0.20000000	9.345390	5.0000	Upper
19	1	0.17741935	10.849740	5.0000	Upper
20	1	0.19354839	12.692399	5.0000	Upper
21	1	0.23333333	15.369557	5.0000	Upper
22	1	0.12903226	15.858977	5.0000	Upper
23	1	0.16666667	17.137786	5.0000	Upper
24	1	0.12903226	17.627206	5.0000	Upper
26	1	0.10714286	17.657491	5.0000	Upper
27	1	0.16129032	18.823530	5.0000	Upper
28	1	0.13333333	19.403166	5.0000	Upper
29	1	0.12903226	19.892586	5.0000	Upper
30	1	0.20000000	21.870569	5.0000	Upper
31	1	0.06451613	21.006749	5.0000	Upper

Table 2.2: Office 'WMNB'

Month1	Subgroup Sample Size	Individual Value	Cusum	Decision Interval	Decision Interval Exceeded
7	1	0.06451613	0.000000	5.0000	
8	1	0.12903226	0.489420	5.0000	
9	1	0.13333333	1.069055	5.0000	
10	1	0.12903226	1.558475	5.0000	
11	1	0.13333333	2.138111	5.0000	
12	1	0.16129032	3.304150	5.0000	
13	1	0.12903226	3.793570	5.0000	
14	1	0.14285714	4.572970	5.0000	
15	1	0.16129032	5.739009	5.0000	Upper
16	1	0.13333333	6.318645	5.0000	Upper
17	1	0.12903226	6.808065	5.0000	Upper
18	1	0.16666667	8.086874	5.0000	Upper
19	1	0.09677419	7.899674	5.0000	Upper
20	1	0.06451613	7.035853	5.0000	Upper
21	1	0.06666667	6.217141	5.0000	Upper
22	1	0.12903226	6.706561	5.0000	Upper
23	1	0.13333333	7.286196	5.0000	Upper
24	1	0.16129032	8.452236	5.0000	Upper
25	1	0.09677419	8.265035	5.0000	Upper
26	1	0.14285714	9.044435	5.0000	Upper
27	1	0.16129032	10.210475	5.0000	Upper
28	1	0.10000000	10.090936	5.0000	Upper
29	1	0.12903226	10.580356	5.0000	Upper
30	1	0.16666667	11.859166	5.0000	Upper
31	1	0.06451613	10.995345	5.0000	Upper

Table 2.3: Office 'T95-1'

Subgroup Month1	Sample Size	Individual Value	Cusum	Decision Decision Interval	Decision Interval Exceeded
7	1	0.03225806	0.00000000	5.0000	
8	1	0.03225806	0.00000000	5.0000	
9	1	0.11666667	0.23004861	5.0000	
10	1	0.08064516	0.00000000	5.0000	
11	1	0.10000000	0.00000000	5.0000	
12	1	0.03225806	0.00000000	5.0000	
13	1	0.09677419	0.00000000	5.0000	
15	1	0.11290323	0.15110961	5.0000	
16	1	0.10000000	0.03157119	5.0000	
17	1	0.09677419	0.00000000	5.0000	
18	1	0.08333333	0.00000000	5.0000	
19	1	0.06451613	0.00000000	5.0000	
20	1	0.06451613	0.00000000	5.0000	
21	1	0.03333333	0.00000000	5.0000	
22	1	0.12903226	0.48941964	5.0000	
23	1	0.06666667	0.00000000	5.0000	
26	1	0.07142857	0.00000000	5.0000	
27	1	0.06451613	0.00000000	5.0000	
28	1	0.08333333	0.00000000	5.0000	
29	1	0.11290323	0.15110961	5.0000	
30	1	0.13333333	0.73074525	5.0000	
31	1	0.03225806	0.00000000	5.0000	

Table 2.4: Office ‘AS1’ Counter

Subgroup Month1	Sample Size	Individual Value	Number of Upper Cusum	Consecutive Upper Sums >0
7	1	0.06451613	0.000000	0
8	1	0.09677419	0.000000	0
9	1	0.16666667	1.278810	1
10	1	0.12903226	1.768229	2
11	1	0.13333333	2.347865	3
12	1	0.16129032	3.513905	4
13	1	0.12903226	4.003324	5
14	1	0.14285714	4.782724	6
15	1	0.16129032	5.948764	7
16	1	0.15000000	6.877986	8
17	1	0.12903226	7.367406	9
18	1	0.20000000	9.345390	10
19	1	0.17741935	10.849740	11
20	1	0.19354839	12.692399	12
21	1	0.23333333	15.369557	13
22	1	0.12903226	15.858977	14
23	1	0.16666667	17.137786	15
24	1	0.12903226	17.627206	16
26	1	0.10714286	17.657491	17
27	1	0.16129032	18.823530	18
28	1	0.13333333	19.403166	19
29	1	0.12903226	19.892586	20
30	1	0.20000000	21.870569	21
31	1	0.06451613	21.006749	22

Figure 2.9: Office ‘WEUZ’ EWMA

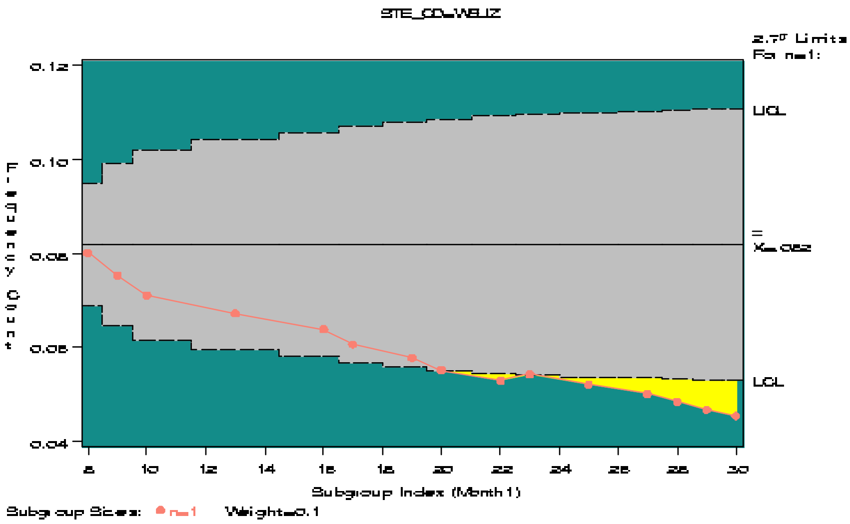

because with only three years of data there would be only 12 time periods, which would not be enough for a good analysis. Month was chosen over workweek because the performance of some of the charts were less than satisfying when workweek was used. Now some of these and the reasons behind these results will be examined.

Using workweek as the time variable, $\mu = 1.04629$ and $\sigma = 0.200818$. To outline a problem with using workweek, office ‘SB2’ will be examined. This office only has 6.5 total errors, yet from Figure 2.10 the process is clearly out-of-control. With only 10 errors one would surely expect for this process to be in control. The reason behind this paradox lies in the way that the EWMA statistic is calculated. Recall that the EWMA statistic is

$$z_i = \lambda x_i + (1 - \lambda) z_{i-1}.$$

A closer look shows that the errors first occurred at workweek 52, during which there were 2 errors. Therefore, the calculations for the EWMA statistic would be

$$\begin{aligned} z_i &= 0.1(2) + (1 - 0.1)(1.04629) \\ &= 1.14 \end{aligned}$$

Also, the upper control limit, with $\lambda = 0.1$ and $L = 2.7$, would be 1.1005. Therefore, the first point 1.14 would exceed the upper control limit. Figure 2.11 shows office 'SB2' analyzed by month.

On the other hand, there are also cases where a process will not give an out-of-control signal even though there are a lot of errors present. For example, Figure 2.12 shows 'WMNB'. Office 'WMNB' has 95 total errors, and from Figure 2.4 shows that the process is out of control. But, if we do the analysis based on workweek, in Figure 2.12, we now will conclude that the process is in control. This is because even though this company has a lot of errors, none of these errors are occurring in the same week. This will cause the EWMA statistic to never go out-of-control. Doing the analysis in months instead of workweeks alleviated this problem by pooling the errors and increasing the sensitivity of the EWMA chart.

Figure 2.10: Office 'SB2' by Workweek

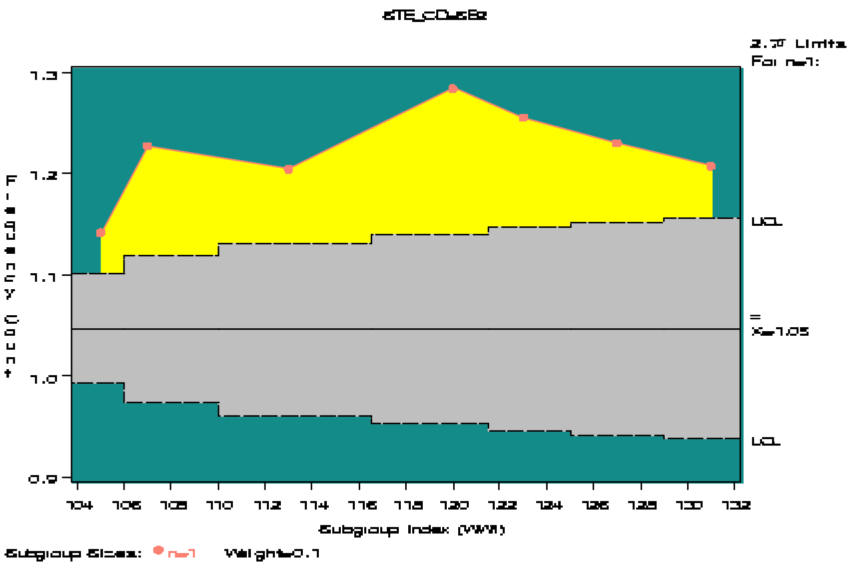

Instead of using the number of errors for each month the average number of errors per day of that month was used. This didn't change my analysis at all. If you look at Figure 2.13, which is for office 'T95-1', there is no difference between this chart and its counterpart Figure 2.6. The same is evident for Figure 2.14, office 'AS1', and its counterpart Figure 2.2.

The mean is 2.489 and the standard deviation is 1.4469 for the undivided errors per month. Since months are of different lengths, the average eliminates any possible “false” seasonality.

Figure 2.11: Office ‘SB2’ by Month

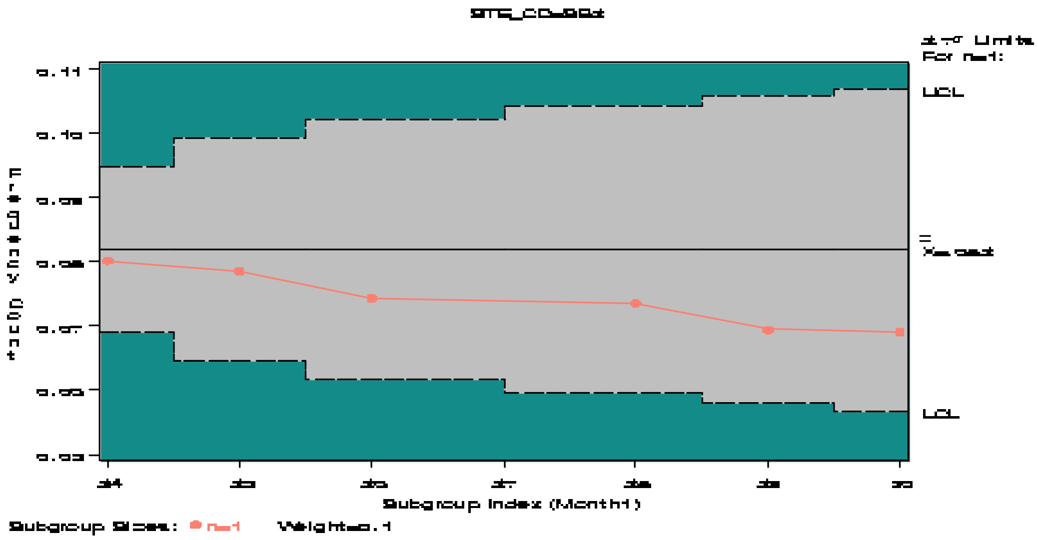

Figure 2.12: Office ‘WMNB’ EWMA

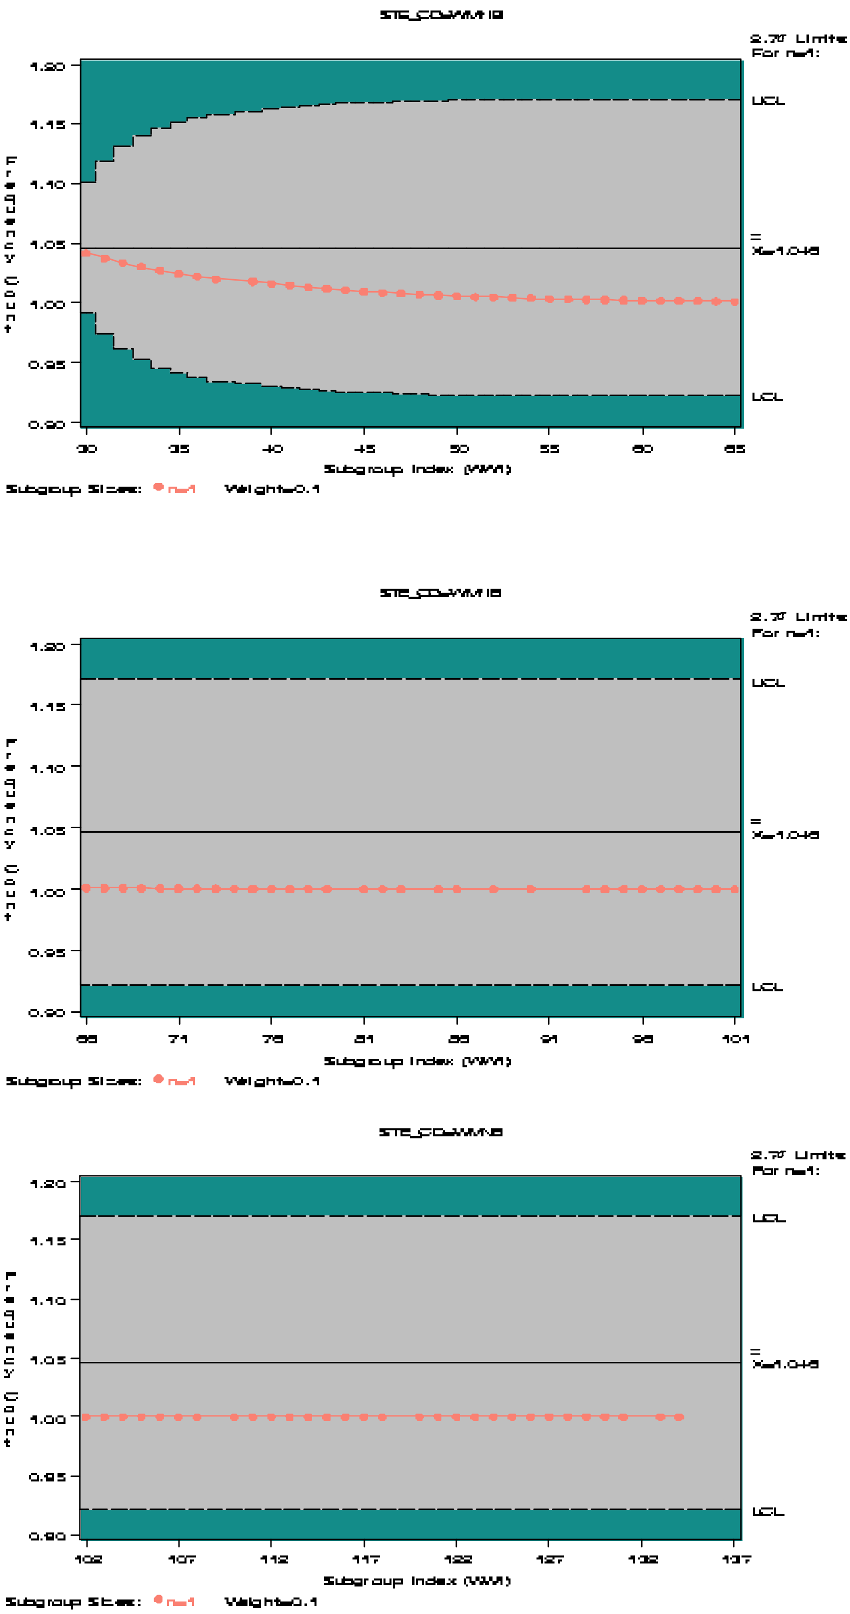

Figure 2.13: Office ‘T95-1 with total errors per month

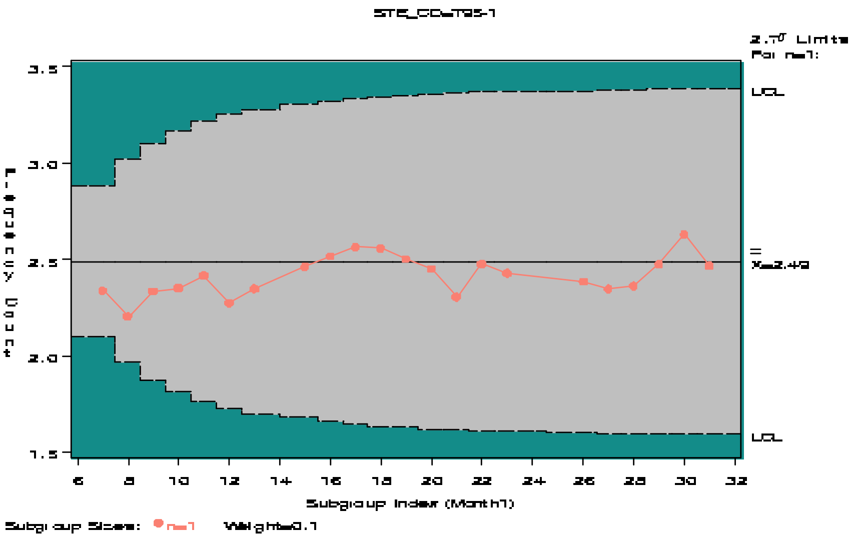

Figure 2.14: Office ‘AS1’ with total errors per month

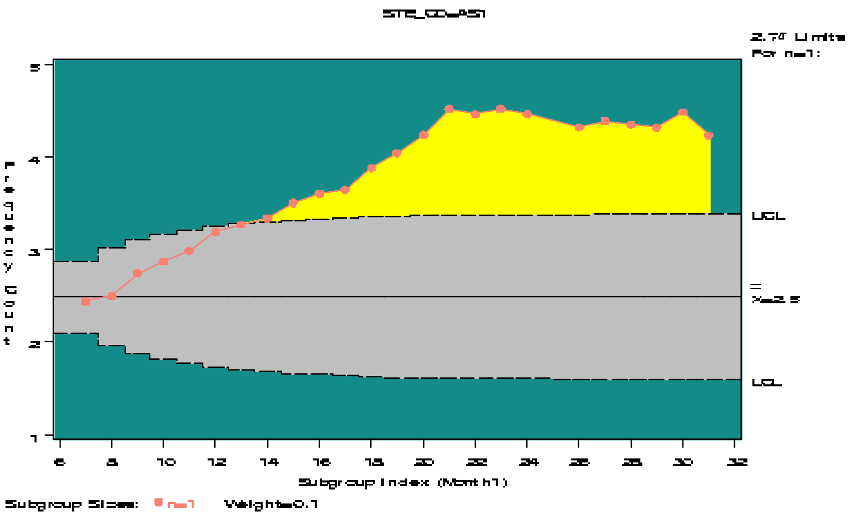

Chapter 3

Control Charts for the Intel Data Under Autocorrelation

A fundamental assumption of the control charts that we have discussed so far is that the observations are all independent. However, in practice, this assumption is often violated, and the resulting dependence is called autocorrelation. An investigation of this data showed that there are 18 companies that exhibit some form of autocorrelation.

In general, autocorrelation means that there are carryover effects from past observations on the present or future observations. The primary effect that autocorrelation has on control charts is that it causes an increase in the frequency of false alarms. Many methods have been suggested to deal with the situation of autocorrelation. Some have suggested that autocorrelation has no effect on the performance of a control chart unless the autocorrelation is very large- for example, over 0.8 (Wheeler, 1991). On the other hand, some suggest that even a small amount of autocorrelation will have a profound effect on control chart performance. In section 3.1 we will examine both arguments and present three different views on dealing with autocorrelation.

3.1 Remedies for Autocorrelated Data

3.1.1 Residual Control Charts

The first and probably most widely used method to deal with autocorrelation is constructing a residual control chart. Residual control charts involves fitting an appropriate times series model to the observations and then applying control charts to the residuals of the model. The residual can be thought of as the actual data value minus the fitted value or predicted value. There will always be some amount of error when fitting a time series model. The general time series model used is the Box-Jenkins (1976) autoregressive integrated moving average model (ARIMA). This model is of the form

$$\varphi_p(B) \bigtriangledown^d X_t = \theta_q(B) a_t \tag{3.1}$$

where $\varphi_p(B) = (1 - \varphi_1 B - \varphi_2 B^2 - \varphi_3 B^3 - \ldots - \varphi_p B^p)$ is an autoregressive polynomial of order p, $\theta_q(B) = (1 - \theta_1 B - \theta_2 B^2 - \theta_3 B^3 - \ldots - \theta_q B^q)$ is a moving average polynomial of order q, $\bigtriangledown$ is the backward difference operator, B is the backshift operator, and $a_t \overset{iid}{\sim} N(\mu, \sigma^2)$. If the model is an appropriately fitted ARIMA model then the residuals a_t will behave like independent and identically distributed random variables (Box and Jenkins, 1976). Therefore, control charts can be applied to these residuals. Of course, the simplicity of the usual Shewhart, EWMA, or CUSUM control chart is lost when using residual control charts. This is because fitting an appropriate time series model to the data requires a lot more statistical knowledge than do the common charts. Though one must posses a lot more statistical knowledge to implement the residual chart the process can be somewhat simplified because of a special case of the ARIMA model. This special case, the ARIMA(0,1,1) or integrated moving average, provides a good approximation for many applications. This last point will be discussed in greater detail in subsequent sections especially in it's relation to the EWMA statistic.

3.1.2 EWMA centerline control chart

Another method suggested to deal with the presence of autocorrelation in process was suggested by Montgomery and Mastrangelo (1991). They used the EWMA statistic as an approximation to the exact residual model. A closer look at the EWMA statistic reveals that it is based on a special case of the Box and Jenkins (1976) ARIMA models. This model is the integrated moving average IMA(1,1), or ARIMA(0,1,1).

$$X_t = X_{t-1} - \theta a_{t-1} + a_t \quad \text{or} \quad (1 - B) Z_t = \theta B_{at} \tag{3.2}$$

Box and Jenkins (1976) use the fact that $z_t = \hat{z}_{t-1}(1) + a_t$ and describe the forecasting of an IMA(1,1) process at time $t + l$ with the equation

$$Z_t = \lambda(\lambda_t) + (1 - \lambda)Z_{t-1} \tag{3.3}$$

where $\lambda = 1 - \theta$. Because the previous forecast $\hat{z}_{t-1}(1)$ falls short of the actual value by a_t, it is adjusted by λa_t. They describe λ as a measure of the proportion of any given shock a_t, which is permanently absorbed by the level of the process. They also describe the forecasting of the IMA (1,1) with the equation

$$\hat{Z}_t(l) = \lambda Z_t + (1 - \lambda)\hat{Z}_{t-1}(l) \tag{3.4}$$

which implies that the new forecast is a linear interpolation at λ. They point out that if λ is very small, we shall be relying primarily on a weighted average of past data and heavily discounting the new observation z_t. They go on to show that for a IMA(1,1), the forecast for all future time is an exponentially weighted moving average of current and past z's. As we can see the EWMA statistic is based on the IMA(1,1) model. Thus, a EWMA with $\lambda = 1-\theta$ is the optimal one-step ahead forecast for the process.

What if the process is not exactly IMA(1,1), but is instead modeled by some other ARIMA model? If the observations from the process are positively autocorrelated and the process mean doesn't drift too quickly, the EWMA statistic with an appropriate value for λ will provide a good one-step-ahead predictor. Just as the case with the residuals, if the process is modeled correctly the one-step-ahead prediction errors should be independently and identically distributed with mean zero and some standard deviation σ_p. These results are used by Montgomery and Mastrangelo (1991) as the basis for a control procedure based on the EWMA statistic that is an approximation of the exact ARIMA model approach from section 4.2.1. This process would consist of plotting these one-step-ahead EWMA prediction errors on a control chart. This chart would be accompanied by a EWMA centerline control chart of the original observations on which the forecast is superimposed. These two charts would provide information about the state of statistical control and process dynamics to be visualized.

Before this process is begun, a value for λ and an estimate of σ_c must be specified. The value of λ should be selected to minimize the sum of squares of the one-step ahead prediction errors. Montogomery and Mastrangelo (1991) suggest estimating σ_c in one of three ways:

1. Divide the sum of squared prediction errors for the optimal λ by n
2. Use the mean absolute deviation or MAD; or
3. Directly calculate a smoothed variance.

The control chart would be

$$\begin{aligned} UCL_{t+1} &= Z_t + \mu_{\alpha/2}\sigma_P \\ CL &= Z_t \\ LCL_{t+1} &= Z_t - \mu_{\alpha/2}\sigma_P. \end{aligned} \tag{3.5}$$

3.1.3 CUSUM

Autocorrelation in process data also has an effect on the performance of the CUSUM control chart. The two general approaches used to combat this problem are plotting the residuals from an adequately fit time series model and plotting the original observations on a standard control chart and adjusting the control limits and process parameter estimation to account for autocorrelation. These approaches will now be examined with an example.

Lu and Reynolds (2001) studied the properties of these two procedures. Their investigation is done in the case of processes that can be modeled as an AR(1) process plus an additional random error or equivalently an ARMA(1, 1) process. There model which is used to model observations from an autocorrelated process is

$$x_k = \mu_k + \epsilon_k, \quad k = 1, 2, \ldots \tag{3.6}$$

Where x_k is the observation taken at sampling time k and μ_k is the random process mean at sampling time k. It is assumed that μ_k can be described as the AR(1) process

$$\mu_k = (1-\phi)\epsilon + \phi\mu_{k-1} + \alpha_l k, \quad k = 1, 2, \ldots \tag{3.7}$$

where ξ is the process mean, α_k is random error, and φ is the autoregressive parameter with $-1 < \varphi < 1$. The α_k's are assumed to be independent normal random variables with mean 0 and variance σ_α^2 and independent of the ϵ_k's. It is also assumed that the starting value μ_0 follows a normal distribution with mean ξ and variance $\sigma_X^2 = \sigma_\alpha^2/(1-\varphi^2)$, which means that the distribution of X_k is constant with mean ξ and variance $\sigma_X^2 = \sigma_\mu^2 + \sigma_\epsilon^2$. The quantity μ_k, mean at time k, is different than the overall process mean $\xi = E(\mu_k)$. The proportion of the process variance due to μ_k is defined as

$$\psi = \frac{\sigma_\mu^2}{\sigma_x^2} = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2} \tag{3.8}$$

which means the proportion of the variance due to ϵ_k is $1-\psi$. The correlation between two adjacent observations is

$$\rho = \varphi\psi$$

By studying the performance and ARL properties of the two approaches Lu and Reynolds concluded that for moderate levels of autocorrelation, both types of CUSUM charts require about the same amount of time to detect shifts in the process mean. However, for higher levels of autocorrelation, the two types of CUSUM charts perform similarly for small shifts, but the residual CUSUM is somewhat better than the CUSUM of observation for large shifts. Therefore, they concluded that it is satisfactory to use a chart based on the original observations, rather than the residuals, unless the level of autocorrelation is relatively high. The ease of interpretation of the original observation control chart makes it the preference of practitioners. But, when using this chart one should account for the autocorrelation both in estimating the process standard deviation and in determining the control limits. For higher levels of autocorrelation it is also advisable to increase k, the slack value, in the CUSUM chart of the observation. This in turn makes it easier to detect small shifts in the process mean in the presence of autocorrelation.

The CUSUMR (Residual CUSUM) of the observation used, by Lu and Reynolds (2001), is similar to the one used for independent observations in Equation 1-24, except that the observation X_k is replaced by the residual e_k, and the in-control mean μ_0 is no longer needed because the in-control mean of the residuals is zero. Therefore, the two CUSUM control statistics are

$$\begin{aligned} CR_k^+ &= \max\{0, CR_k^+ + (e_k - k\sigma_\gamma)\} \\ CR_k^- &= \min\{0, CR_k^- + (e - k\sigma_\gamma)\} \end{aligned} \tag{3.9}$$

where $k\sigma\gamma$ is the reference value. A signal is given if CR_k^+ falls above an upper control limit $h\sigma_\gamma$ or if CR_k^- falls below a lower control limit - $h\sigma_\gamma$. The choice of k will not necessarily be the same as in the case of independent observations because the mean of the residuals is not constant after the shift of the process mean. Tabulated values of k are presented in Lu and Reynolds (2001). These tabulated values are all calculated using various values of ψ and φ with a desired shift in the mean δ (all values are tabulated to have an approximate in-control ARL of 370.4). Their research showed that for relatively low levels of autocorrelation a moderate value of k around 0.5 would give reasonably good performance over a large range of shifts for the CUSUM of observations and the CUSUMR. For relatively high levels of autocorrelation one should use a relatively large value of k, 1.0 say, in the CUSUM of the observations. However, a moderate value of k such as 0.5 is still a good value for a CUSUM of the residual. An example is presented in Table 3.1.

Office ‘A06’ was also examined in Section 3.1.3; this application is in agreement with that chart.

The second approach of Lu and Reynoldss (2001) for dealing with autocorrelation involves adjusting the control limits and the parameter estimation techniques. However, these methods can not easily be implemented on the Intel data. For example, to implement

Table 3.1: Office 'A06'

Month1	Subgroup Sample Size	Individual Value	Cusum	Decision Interval	Decision Interval Exceeded
8	1	0.03225806	0.2105120	5.0000	
10	1	0.02121721	0.1377388	5.0000	
11	1	-.00742150	0.0000000	5.0000	
12	1	0.03644057	0.3178265	5.0000	
13	1	-.02208172	0.0000000	5.0000	
14	1	0.00854439	0.0000000	5.0000	
15	1	0.00217711	0.0000000	5.0000	
16	1	0.02283011	0.0000000	5.0000	
17	1	0.02240280	0.0000000	5.0000	
18	1	0.07899355	1.4096465	5.0000	
19	1	-.08004415	0.0000000	5.0000	
20	1	0.13412043	2.8240865	5.0000	
21	1	0.03081721	2.9976292	5.0000	
22	1	-.07586164	0.4340165	5.0000	
23	1	0.07899355	1.8436630	5.0000	
24	1	-.03165706	0.4142472	5.0000	
25	1	0.06110753	1.3649759	5.0000	
26	1	-.03725099	0.0000000	5.0000	
27	1	0.06887036	1.1499065	5.0000	
28	1	0.02465377	1.1653083	5.0000	
29	1	0.04898810	1.8050780	5.0000	
30	1	0.08081721	3.2615157	5.0000	
31	1	-.10184594	0.0312004	5.0000	

this method one must estimate certain parameters to account for the process.s autocorrelation, such as the process standard deviation. However, their adjustment to the standard deviation to account for autocorrelation involves the range method for standard deviation estimation. In this analysis, the standard deviation was estimated directly, and they presented no methods for this case. Another complication to implementing this method is that it suggests increasing the control limit to account for autocorrelation. But, if the control limits are increased this will also effect the ARL properties. This will primarily increase the out-of-control ARL, which is not an ideal situation. Because of these complications of dealing with autocorrelation and the problem of plotting residuals (see Section 3.1.4 and Chapter 4), the CUSUM chart was eliminated as the main control scheme.

3.1.4 EWMAST

The final method discussed was introduced by Zhang (1998). He recommends his chart as an alternative to the residual control chart for autocorrelated data. He argued that even though a residual control chart could be applied to any autocorrelated data even if the data are from a nonstationary process, the nonstationary chart's signal of an out-of-control condition only means that the process has some system deviations because a nonstationary process does not have a constant mean or variance. It was also noted that the detection capabilities of the residual control chart were poor and that the properties of the residual charts are different than those of traditional control charts. For more discussion on this see Wardell, et.al (1992).

As an alternative he suggests using an EWMAST control chart, which is an exponentially weighted moving average control chart for a stationary or weakly stationary process. When the process is positively autocorrelated, the control limits of this chart are wider than the usual EWMA control chart. Given that the EWMA of a stationary process is asymptotically a stationary process, and that the covariance between z_t and z_{t+1} converges when t is large

(Zhang 1998), there exists an integer M such that

$$\begin{aligned} Cov[Z_t, Z_{t+r}] &= [\lambda/(2-\lambda)] \times \sigma_x^2 \{\sum_{k=0}^{M} \rho(k+\tau)(1-\lambda)^k [1-(1-\lambda)^{2(M-k)}] \\ &+ \sum_{k=1}^{\gamma} \rho(k-\tau)(1-\lambda)^k [1-(1-\lambda)^{2M}] + \sum_{k=\tau+1}^{M+\tau} \rho(k-\tau)(1-\lambda)^k \\ &\times [1-(1-\lambda)^{2(M+\tau-k)}]\} \end{aligned} \quad (3.10)$$

Thus, the approximate variance of z_t is

$$\sigma_Z^2 = [\lambda/(2-\lambda)]\sigma_x^2 \times \{1 + 2\sum_{k=1}^{M} \rho(k)(1-\lambda)^k \times [1-(1-\lambda)^{2(M-k)}]\} \quad (3.11)$$

where σ_x^2 is the target or estimated process variance, and $\rho(k)$ is the autocorrelation of X_t at lag k. Therefore, the EWMAST chart is of the form

$$\begin{aligned} UCL &= \mu + L\sigma_Z \\ CL &= \mu \\ LCL &= \mu - L\sigma_Z \end{aligned} \quad (3.12)$$

Note, the regular EWMA chart is a special case of the EWMAST chart when $\rho(k) = 0$ for $k \geq 1$. If this is the case, the term in the brackets of Equation 3-13 equals 1, which in turn reduces Equation 3.13 to Equation 2.5, the steady state variance of the usual EWMA control chart.

This procedure is simpler to implement than a residual chart because no time series modeling is needed. It also does not suffer from the same detection capability problems that the residual chart has. Zhang used the ARL to conclude that the performance of the EWMAST chart is better than the residual control chart, the EWMA centerline control chart, $\bar{x}$ chart, and the regular EWMA chart when the process autocorrelation is not very positively strong and the mean shifts are small to medium.

3.2 Application

The first office exhibiting autocorrelation to be examined is office ‘A06’. This office has a total of 63.5 expense report errors. Figure 3.1 shows the usual EWMA control chart for this company. It shows an out-of-control process with a gradual increase in the latter months. But, if one looks at Table 3.2, one sees that there is some correlation between the observations. The test for white noise, or lack of correlation, is an approximate test, testing the null hypothesis that the autocorrelations up to a given time lag are 0 against the alternative hypothesis that some of the autocorrelations up to a given lag are significantly different than 0. If the p-value is greater than our significance level, 0.05, the null hypothesis is accepted and the conclusion is that the observations are uncorrelated. But, if the p-value is less than 0.05 the null hypothesis is rejected and the conclusion is that the observations are correlated. From Table 3.2 it can be seen that up to lag 6 the p-value is greater than .05 (.0763), but from lags 13 to 18 the p-value is not greater than 0.05 (0.0095). Therefore, there is statistical evidence that some of the observations are autocorrelated. Figure 3.2 shows a plot of the 1st differenced autocorrelation and partial autocorrelation functions. These indicate that the process could possibly be modeled by an ARIMA(0,1,1), which we mentioned earlier is usually a good approximation for many applications.

Figure 3.1: Office 'A06'

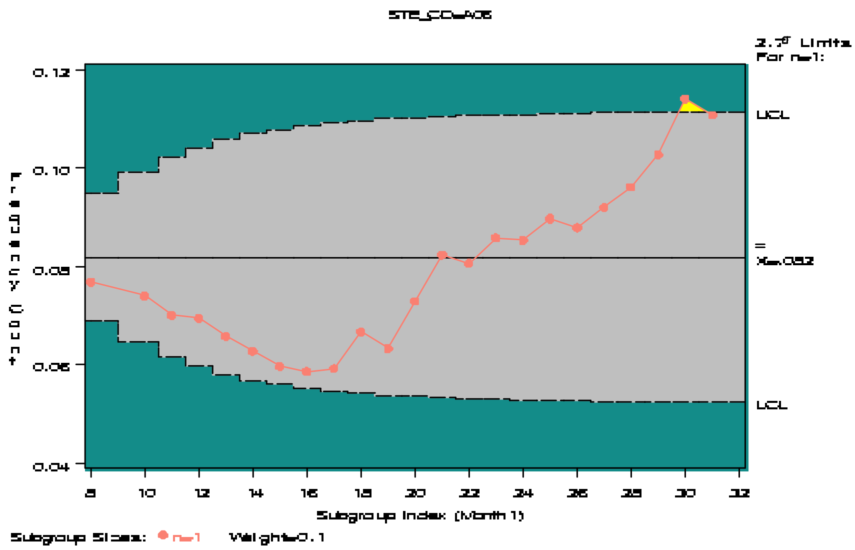

Table 3.3 and Table 3.4 test for independence of the residuals and for significance of the moving average parameter, respectively. If the model is a good fit then the residuals should

Table 3.2: Autocorrelation Check for White Noise

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations				
6	11.42	6 0.0763	0.338	0.398	0.303	0.123	0.177	-0.005
12	15.34	12 0.2235	0.101	-0.050	0.100	-0.028	-0.228	-0.099
18	35.00	18 0.0095	-0.163	-0.262	-0.216	-0.230	-0.176	-0.164

behave like independent and identically distributed random variables. It is easily seen that this is the case because the null hypothesis that the correlations are 0 is not rejected, and thus the residuals are independent.

Figure 3.2: Office 'A06' Autocorrelation Function

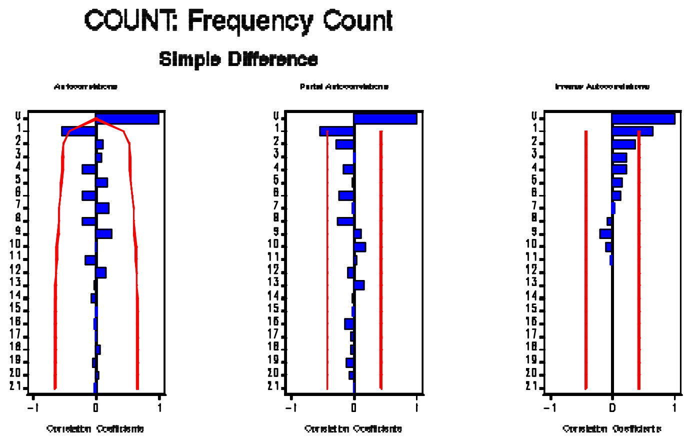

A control chart may now be applied to the residuals. The first chart, given in Figure 3.3 is a chart of the residuals applied to a standard $\bar{x}$ control chart, followed by a chart of the forecasts. Figure 3.4 is a chart of the residuals applied to a EWMA control chart.

The process now seems to be in control on both charts. Should one now conclude that the process is really in control and the out-of-control condition of Figure 3.1 using the regular EWMA control chart was a false alarm? The answer to this question lies in the design of these residual charts. The residual control charts require the use of the mean and standard

Table 3.3: Office 'A06' Autocorrelation Check for White Noise of Residuals

To Lag	Chi-Square	DF	Pr > ChiSq	Autocorrelations					
6	2.39	5	0.7927	-0.007	0.149	-0.068	0.136	-0.185	-0.046
12	3.67	11	0.9786	-0.046	-0.051	-0.074	0.069	0.097	-0.064
18	16.96	17	0.4573	0.116	0.123	0.140	-0.097	0.280	-0.082

Table 3.4: Office 'A06' Parameter estimate

Parameter	Estimate	Standard Error	t Value	Approx Pr > $\|t\|$	Lag
MA1,1	0.56201	0.19581	2.87	0.0092	1

deviation of each office and not the overall (target) values used in the general EWMA. This means that the residuals cannot be compared to the standard values. If the residuals were plotted with the target mean of 0.081861 and target standard deviation of 0.04768, the graph would never show an out-of-control condition because the residual values are much smaller than the actual values. Unless there is a way of determining a target residual value, the use of residual control charts in this project should be examined very carefully.

If the view given by Wheeler (1991) that the autocorrelations will only have an impact if the autocorrelations are greater than 0.8 is accepted, the conclusion is that there is no effect due to the autocorrelation. Table 3.5 shows the autocorrelation plot of office 'A06'.
Clearly, none of the autocorrelations are near 0.8. In fact, the greatest autocorrelation is 0.39815. This problem will be revisited later in relation to the Intel data.

Some examples of the EWMAST chart applied to some of the autocorrelated offices are presented next. The first example is that of office 'A06', which was analyzed earlier with a

Table 3.5: Office 'A06' Autocorrelations

```
Lag   Covariance    Correlation    -1 9 8 7 6 5 4 3 2 1 0 1 2 3 4 5 6 7 8 9 1
  0    0.0028717        1.00000    |                    |********************|
  1   0.00096971        0.33768    |            .       |*******.            |
  2    0.0011434        0.39815    |           .        |********.           |
  3   0.00087057        0.30315    |          .         |******   .          |
  4   0.00035210        0.12261    |         .          |**        .         |
  5   0.00050744        0.17670    |         .          |****      .         |
  6   -0.0000135        -.00470    |         .          |          .         |
  7   0.00028927        0.10073    |         .          |**        .         |
  8   -0.0001424        -.04959    |         .         *|          .         |
  9   0.00028766        0.10017    |         .          |**        .         |
 10   -0.0000793        -.02762    |         .         *|          .         |
 11   -0.0006560        -.22844    |         .     *****|          .         |
 12   -0.0002852        -.09930    |        .         **|           .        |
 13   -0.0004683        -.16308    |        .        ***|           .        |
 14   -0.0007535        -.26240    |        .      *****|           .        |
 15   -0.0006198        -.21582    |        .       ****|           .        |
 16   -0.0006610        -.23017    |       .       *****|            .       |
 17   -0.0005042        -.17558    |       .        ****|            .       |
 18   -0.0004698        -.16359    |       .         ***|            .       |
 19   -0.0005410        -.18838    |       .        ****|            .       |
 20   -0.0003861        -.13444    |       .         ***|            .       |
 21   -0.0003016        -.10502    |       .          **|            .       |
 22   0.00002563        0.00892    |      .             |             .      |
```

Figure 3.3: Office ‘A06’ Residual Chart

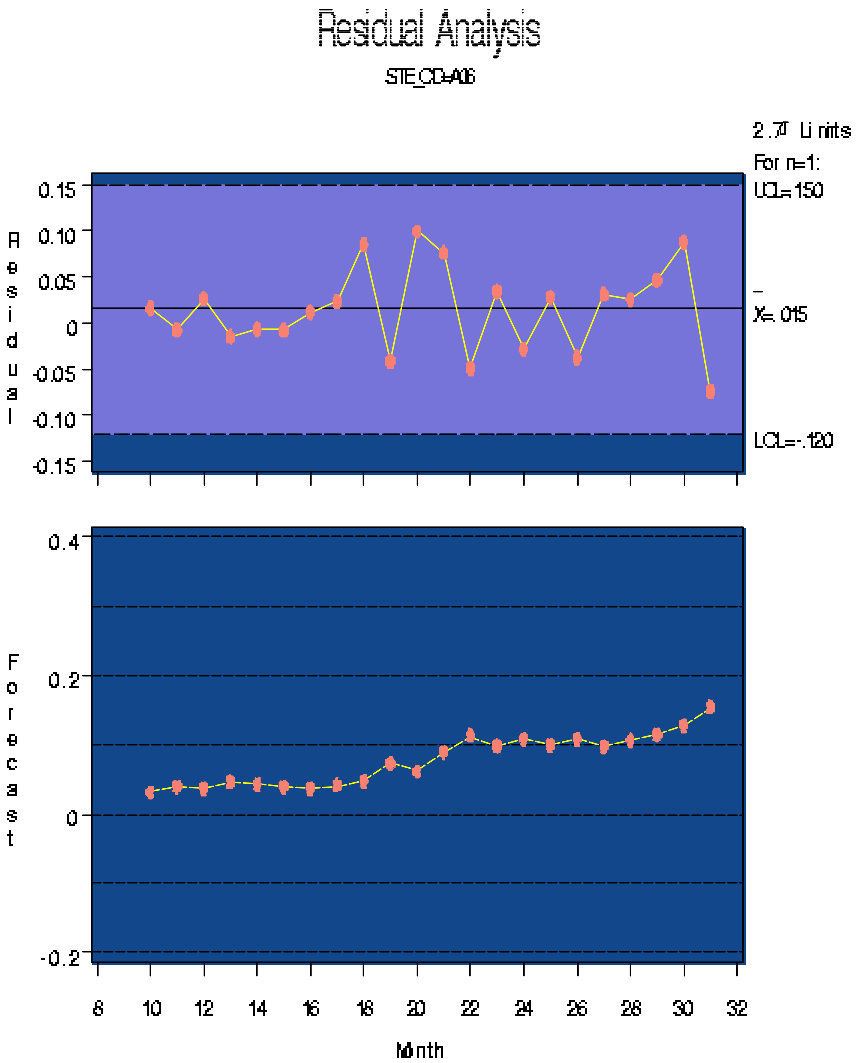

residual chart. Figure 3.5 is the EWMAST chart for this company. It shows an in-control process, which is the same conclusion reached by the residual chart.

Figure 3.6 shows the EWMA chart for office ‘CR01’, which has a total of 103 errors. Figure 3.7 shows the corresponding EWMAST chart for office ‘CR01’. As expected, the process is out-of-control. Examining these charts in comparison to the regular EWMA control chart it is apparent that the limits are wider to allow for autocorrelation. In fact, from Equation 2-4

Figure 3.4: Office 'A06' EWMA Residual

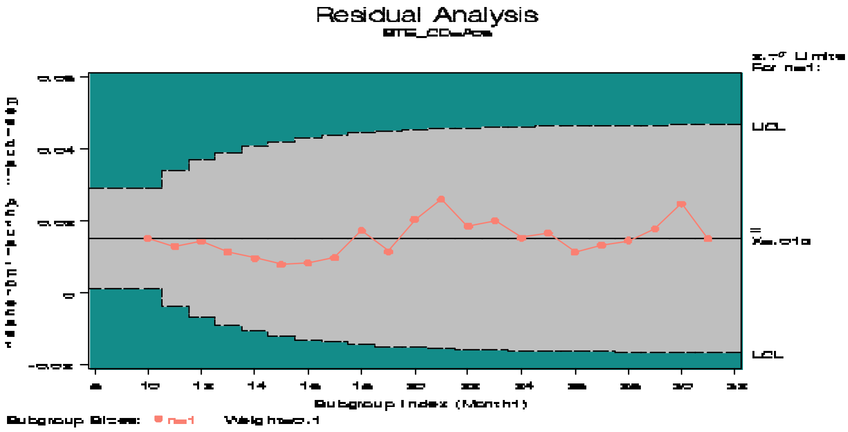

Figure 3.5: Office 'A06' EWMAST

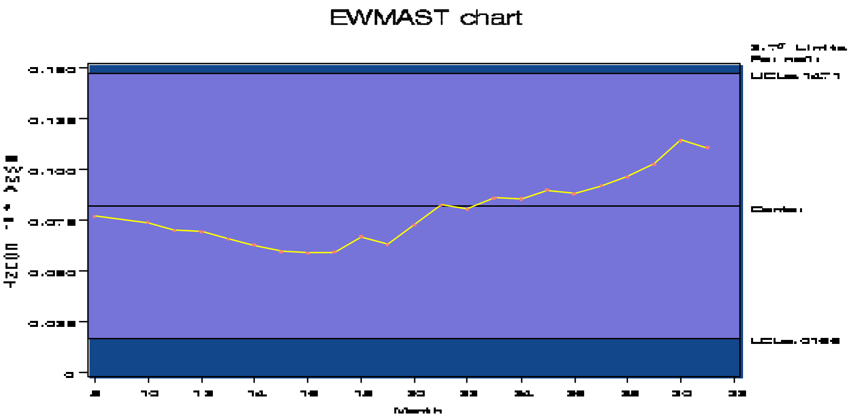

the control limits are

$$
\begin{aligned}
UCL &= .08186 + 2.7 * .04768 * \sqrt{\frac{0.1}{2-0.1}} = 0.1114 \\
LCL &= .08186 - 2.7 * .04768 * \sqrt{\frac{0.1}{2-0.1}} = 0.0523
\end{aligned}
$$

which is thinner than the EWMAST's limits of UCL=0.1294 and LCL=0.0343. In the case of Figure 3.7, the wider limits of the EWMAST chart decreases the number of false alarms.

Figure 3.6: Office 'CR01' EWMA

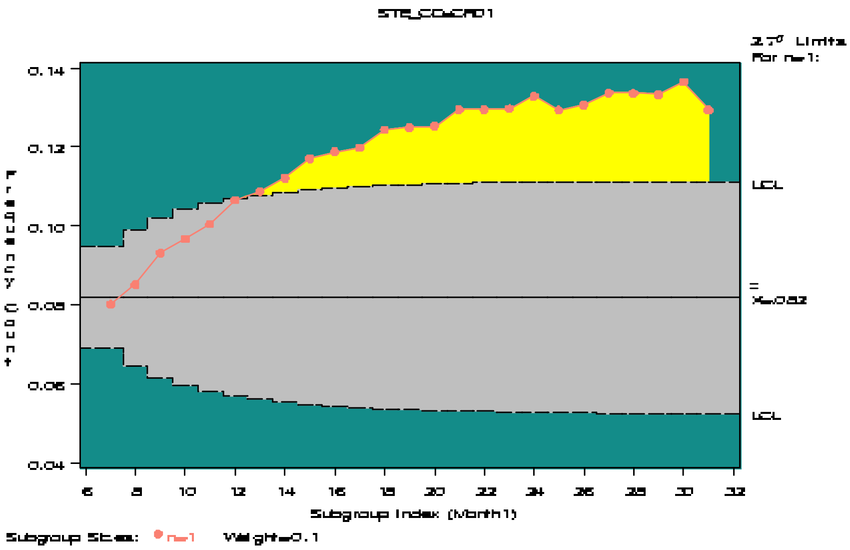

Figure 3.7: 'CR01' EWMAST

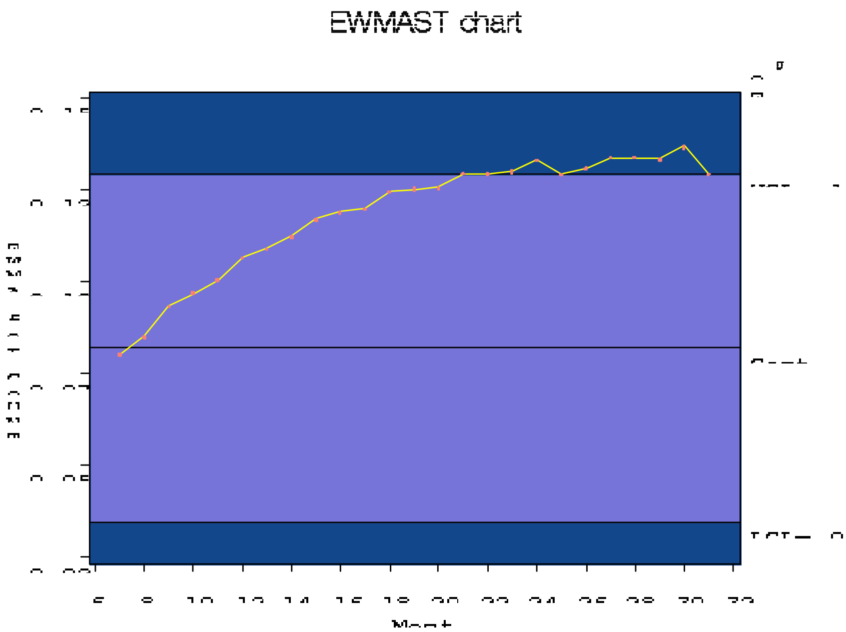

Chapter 4

Conclusion

Throughout this discussion many forms of process control has been presented in the form of control charts. Because of the nature of the Intel data only a small number of these charts were chosen for demonstration. Of those, the EWMA control chart was chosen as the focus because of its sensitivity, robustness to the normality assumption, and for its forecasting ability.

This thesis also dealt with the problem of autocorrelation in process data in the EWMA control chart using the EWMAST chart. But, autocorrelation of process data is a problem that has an effect on every control chart. There are many papers that deal with this problem, for example, see (Balkin and Lin (2001) for a discussion of the effect of autocorrelation on the Shewhart charts, and Johnson and Bagshaw (1974) for the effect of autocorrelation on CUSUM control charts.

Autocorrelation is not the only problem that can arise when dealing with control charts. Many of the control charts assume normality, and if this assumption is violated it can sometimes have a negative effect on the performance of a control chart. The popular $\bar{x}$ chart has been shown by many, such as Schilling and Nelson (1976), to be robust to departures from the normality assumption as long as the population isn't extremely nonnormal. The EWMA control chart was shown by Border, et. al (1999) also to be robust to departures from the normality assumption. They also discuss how greatly the ARL of the Shewhart control chart for individuals is affected by nonnormality and how little the ARL of the EWMA control chart is affected.

Earlier, it was noted from Zhang (1998) that if a process is stationary or weakly stationary with not very strong positive autocorrelation and small to medium shifts of the mean then an EWMAST control chart should be used, but if the process is stationary or near nonstationary with strong positive autocorrelations then a residual control chart is preferable. Zhang's method was used for this analysis instead of residuals because having to fit a *correct* time series model is a drawback especially for those not trained in time series analysis. When using the approximate time series model ARIMA(0,1,1) there are two problems that may arise: first, this approximate model will not leave independent residuals on many processes; and second, if the plotted residuals come very close to the control limits one will not know if that point would exceed that limit if the correct model had been used. Also, the lack of a common mean and standard deviation to analyze each office gives some doubt about its conclusions. One problem that could be encountered with the EWMAST chart is the interpretation of the terms "weakly" stationary and "almost" nonstationary. These terms could mean different things to different people. What one thinks is stationary or weakly stationary another may think is stationary or almost stationary. Therefore, careful consideration should be exercised before declaring the stationarity of a process.

Some believe that if the autocorrelation is not very high then there will be no effect on the performance of the control chart. In this analysis, none of the autocorrelations were very high. When the lag-one autocorrelations were moderate- 0.5887, or 0.644451, say- the autocorrelations of the subsequent lags usually die off very quickly. When examing the usual EWMA chart, the residual chart, and the EWMAST chart there seems to be very little difference in the detection performance. Only those processes that were borderline out-of-control in the first place showed a difference. Though some of the autocorrelated offices were shown to be in control the question still should be asked: Why are the errors dependent? These offices should be examined also to determine the origin of this condition.

Table 4.1: Complete Results of all Offices

STE_CD	Total Errors	Analysis Type	Autocorrelated	Control Limit	Decision Exceed
A01	20.5	EWMA	no	lower	audit
A02	71	EWMAST	yes	upper	audit
A03	29	EWMA	no	none	o.k.
A04	30.5	EWMA	no	none	o.k.
A05	17	EWMA	no	none	o.k.
A06	63.5	EWMAST	yes	none	borderline
A11	84.5	EWMA	no	upper	audit
A13	76	EWMA	no	upper	audit
A17	2	EWMA	no	none	o.k.
A22	1	EWMA	no	none	o.k.
A23	79.5	EWMAST	yes	upper	audit
A24	24.5	EWMA	no	none	o.k.
A25	25.5	EWMAST	yes	none	o.k.
A27	21	EWMA	no	none	o.k.
A3B	30	EWMA	no	none	o.k.
A3P	12	EWMA	no	none	o.k.
ACS	19	EWMA	no	none	o.k.
ADJ	37	EWMA	no	none	o.k.
AGB	23	EWMA	no	none	o.k.
AHE	86.5	EWMA	no	upper	audit
AK4	90	EWMA	no	upper	audit
AME	17	EWMA	no	none	o.k.
AMH	14	EWMA	no	none	o.k.
API	39	EWMA	no	none	o.k.
APL	87	EWMA	no	upper	audit
AS1	106	EWMA	no	upper	audit
ASE	68.5	EWMAST	yes	none	borderline
AT1	11	EWMA	no	none	o.k.
ATW	30.5	EWMA	no	none	o.k.
CN01	26	EWMA	no	none	o.k.
CR01	103	EWMAST	yes	upper	audit
D1C	1	EWMA	no	none	o.k.
D1F	11	EWMA	no	none	o.k.
D2F	20.5	EWMA	no	none	o.k.
DBF	20.5	EWMA	no	none	o.k.
F07	2	EWMA	no	none	o.k.
F08	9	EWMA	no	none	o.k.
F11	25.5	EWMA	no	none	o.k.
F12	41.5	EWMA	no	none	o.k.
F14	49.5	EWMA	no	none	o.k.

F17	2	EWMA	no	none	o.k.
F18	34	EWMA	no	none	o.k.
F22	12	EWMA	no	none	o.k.
F23	5.5	EWMA	no	none	o.k.
FET	13	EWMA	no	none	o.k.
FNC	24.5	EWMA	no	none	o.k.
FSG	52	EWMA	no	upper	o.k.
FX6	25	EWMA	no	none	o.k.
HK01	1	EWMA	no	none	o.k.
IR01	1	EWMA	no	none	o.k.
MHT	23	EWMA	no	none	o.k.
NL01	4	EWMA	no	none	o.k.
NL02	23	EWMAST	yes	none	o.k.
NL03	5	EWMA	no	none	o.k.
NULL	6	EWMA	no	none	o.k.
PC2	26.5	EWMA	no	none	o.k.
PC4	3	EWMA	no	none	o.k.
PC5	1	EWMA	no	none	o.k.
PC6	6	EWMA	no	none	o.k.
S03	4	EWMA	no	none	o.k.
S1C	2	EWMA	no	none	o.k.
S23	4.5	EWMA	no	none	o.k.
SB2	10	EWMA	no	none	o.k.
SD1	10.5	EWMA	no	none	o.k.
SD2	24	EWMA	no	none	o.k.
SDB	3	EWMA	no	none	o.k.
SF7	2	EWMA	no	none	o.k.
SF8	24	EWMA	no	none	o.k.
SG01	35	EWMA	no	none	o.k.
SHK	20	EWMA	no	none	o.k.
ST1	3	EWMA	no	none	o.k.
SX1	15.5	EWMA	no	none	o.k.
SX2	5	EWMA	no	none	o.k.
SX4	12	EWMA	no	none	o.k.
SX6	3.5	EWMA	no	none	o.k.
SX7	2	EWMA	no	none	o.k.
SX8	18.5	EWMA	no	none	o.k.
T11-1	2	EWMA	no	none	o.k.
T11-2	6.5	EWMA	no	none	o.k.
T17-1	40	EWMAST	yes	none	o.k.
T17-2	2	EWMA	no	none	o.k.
T27-1	4	EWMA	no	none	o.k.
T27-2	4.5	EWMA	no	none	o.k.
T3-1	52	EWMA	no	none	o.k.
T31-1	87	EWMA	no	upper	audit
T31-2	65	EWMA	no	none	borderline
T3-2	33	EWMA	no	none	o.k.

T33-1	30	EWMAST	yes	lower	audit
T33-2	57	EWMAST	yes	none	o.k.
T5-1	15.5	EWMA	no	none	borderline
T5-2	57	EWMAST	yes	none	o.k.
T66-1	86	EWMA	no	upper	audit
T66-2	53.5	EWMA	no	none	o.k.
T9-1	17	EWMA	no	none	o.k.
T9-2	47	EWMAST	yes	none	o.k.
T93-1	86.5	EWMAST	yes	upper	audit
T93-2	65.5	EWMA	no	upper	audit
T94-1	18.5	EWMA	no	none	o.k.
T94-2	48.5	EWMA	no	none	o.k.
T95-1	53	EWMA	no	none	o.k.
T95-2	39.5	EWMA	no	none	o.k.
TB2	2	EWMA	no	none	o.k.
TD1-1	13	EWMA	no	none	o.k.
THY	1	EWMA	no	none	o.k.
TT1-1	6	EWMA	no	none	o.k.
TT1-2	3	EWMA	no	none	o.k.
TW01	10	EWMA	no	none	o.k.
US03	65	EWMAST	yes	none	borderline
US05	26	EWMA	no	none	o.k.
US06	8	EWMA	no	none	o.k.
US10	51	EWMA	no	upper	audit
US12	1	EWMA	no	none	o.k.
WCHA	103	EWMAST	yes	upper	audit
WCHB	6	EWMA	no	none	o.k.
WEUA	20	EWMA	no	none	o.k.
WEUZ	17	EWMA	no	none	o.k.
WJPA	17	EWMA	no	none	o.k.
WJPB	5	EWMA	no	none	o.k.
WJPZ	7	EWMA	no	none	o.k.
WMNA	5	EWMA	no	none	o.k.
WMNB	95	EWMA	no	upper	audit
WMYA	98	EWMAST	yes	upper	audit
WORA	84	EWMAST	yes	upper	audit
WPNA	2	EWMA	no	none	o.k.
WPNB	77	EWMAST	yes	upper	audit

Bibliography

Alwan, L. and Roberts, H. (1988). "Time-Series Modeling for Statistical Process Control," *Journal of Business and Economic Statistics* 6.

Atienza, O., Ang. B. W., and Tang, L. C. (2002). "A Cusum Scheme for Autocorrelated Observations" *Journal of Quality Technology* 34.

Alt, F. (1985). "Multivariate Quality Control," *Encyclopedia of Statistical Sciences* 6.

Borrow, C., Montgomery, D., and Runger, G. (1999). "Robustness of the EWMA Control Chart to Non-normality," *Journal of Quality Technology* 31.

Box, G. and Jenkins, G. (1976). Time Series Analysis: Forecasting and Control

Chandra, M. (2001). Statistical Quality Control.

Chang, T. and Gan, F. (1995). "A Cumulative Sum Control Chart for Monitoring Process Variance," *Journal of Quality Technology* 27.

Derman, C. and Ross, S. (1997). Statistical Aspects of Quality Control.

Crosier, R. (1988). "Multivariate Generalizations of Cumulative Sum Quality Control Schemes," *Technometrics* 30, 291-303.

Crowder, S. (1989). "Design of Exponentially Weighted Moving Average Schemes," *Journal of Quality Technology* 21.

Faltin, F. and Woodall, W. (1991). "Discussion," *Journal of Quality Technology*, Vol. 23, No. 3.

Farnum, N. (1992). “Control Charts for Short Runs: Nonconstant Process and Measurement Error,” *Journal of Quality Technology* 24, 138-144.

Fellner, W. (1990). “Algorithm AS 258: Average Run Length for Cumulative Sum Schemes,” *Applied Statistics* 39, pp. 402-412.

Gan, F. (1991). “An Optimal Design of CUSUM Quality Control Charts,” *Journal of Quality Technology* 23.

Hawkins, D. (1992). “A Fast, Accurate Approximation of Average Run Lengths of CUSUM control charts,” *Journal of Quality Technology* 25.

Hawkins, D. (1993a). “Cumulative Sum Control Charting: An Underutilized SPC Tool,” *Quality Engineering* 5.

Hotelling, H. (1947). “Multivariate Quality Control,” *Techniques of Statistical Analysis.*

Iwersen, J. (1997) “Statistical Process Control for Autocorrelated Processes: A Case Study,” *Statistical Quality Control* 5.

Johnson, N. (1961). “A Simple Theoretical Approach to Cumulative Sum Control Charts,” *Journal of the American Statistical Association* 54.

Johnson, R. and Bagshaw, M. (1974). “The Effect of Serial Correlation on the Performance of CUSUM Tests,” *Technometics*, Vol. 16, No. 1.

Lowry, C., Woodall, W., Champ, C., Rigdon, S. (1992). “Multivariate Exponentially Weighted Moving Average Control Chart,” *Technometrics* 30, 291-303.

Lu, C. and Reynolds, M. (2001). “Cusum Charts for Monitoring an autocorrelated Process,” *Journal of Quality Technology* 33.

Lucas, J. (1976). “The Design and Use of Cumulative Sum Quality Control Schemes,” *Journal of Quality Technology* 8.

Lucas, J. (1976). "The Design and Use of V-Mask Control Schemes". *Journal of Quality Technology* 8, 1-12.

Lucas, J. (1982) "Combined Shewhart-CUSUM Quality Control Schemes," *Journal of Quality Technology* 14.

Lucas, J. (1985). "Counted Data CUSUM.s," *Technometrics* 27.

Lucas, J. and Crosier, R. (1982). "Fast Initial Response for CUSUM Quality Control Schemes," *Technometrics* 24.

Lucas, J. and Saccucci, M. (1990). "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancement," *Technometrics* 32.

Macgregor, J. and Harris, T. (1993). "The Exponentially Weighted Moving Variance," *Journal of Quality Technology* 25.

Montgomery, D. and Mastrangelo, C. (1991). "Some Statistical Process Control Methods for Autocorrelated Data," *Journal of Quality Technology* 23.

Nelson, L. (1987). "Chi-Square Control Chart for Several Proportions," *Journal of Quality Technology* 19, 229-231.

Ryan, T. (1991). "Discussion," *Journal of Quality Technology* 23.

Schmid, W. (1997). "On EWMA charts for Time Series," *Statistical Quality Control* 5.

Schrock, E. (1950) Quality Control and Statistical Methods.

Schilling, E. and Nelson, P. (1976). "The effect of Nonnormality on the Control Limits On xbar Control Charts," *Journal of Quality Technology* 8.

Shewhart, W. (1931) Economic Control of Quality of Manufactured Product.

Shewhart, W. (1939) Statistical Method from the Viewpoint of Quality Control.

Wardell, W., Moskowitz, H., and Plante, R. (1992) “Control Charts in the Presence of Data Correlation,” *Management Science* 38, pp. 1084-1105.

Wheeler, D. J. (1991b). “Shewhart.s Chart: Myths, Facts.” *American Society for Quality Control*.

Woodall, W., and Adams, B. (1993). “The Statistical Design of CUSUM Charts,” *Quality Engineer* 5.

Western Electric (1956). Statistical Quality Control Handbook, Western Electric Corporation, Indianpolis, Ind.

Zhang, N. (1998) “A Statistical Method from the Viewpoint of Quality Control,” *Technometrics*, Vol. 40, No. 1.