University of Stavanger

Faculty of Science and Technology

MASTER'S THESIS

Study program/ Specialization: Masters of Mathematics and Physics Statistics	Spring semester, 2016. Open / ~~Restricted access~~
Writer: Caroline Johnson	…… (Writer's signature)
Faculty supervisor: Jan Terje Kvaløy External supervisor(s):	
Thesis title: Estimation Error in Statistical Process Control	
Credits (ECTS): 60	
Key words: Statistical process control, estimation error, risk-adjusted CUSUM, Helping Babies Breath, bootstrap guaranteed performance	Pages: 84…… + enclosure: …… Stavanger, 14/06/2016…. Date/year

University of Stavanger

Estimation Error in Statistical Process Control

Caroline Johnson

Department of Mathematics and Natural Science
University of Stavanger
Submission Date: June 2016
Supervisor: Jan Terje Kvaløy

Preface

I would like to thank my supervisor Jan Terje kvaløy for his continued guidance and invaluable feedback throughout the year. I would also like to acknowledge Dr. Hege Langli Ersdal and the SAFER team for the amazing work they have put into the Helping Babies Breath program and thank them for allowing me to analyse the data.

Abstract

This thesis examines estimation errors which occur in control charts with estimated parameters, the effect that these errors have on the charts and the current methods which aim to account for the errors. There is a particular focus on how these errors affect the average run length of the chart and how to improve the average run length using current methods. The best methods to account for estimation error in Shewhart control charts, from the methods discussed, are the bootstrap method and the reduction of exceedance probability method. Both aim to reduce the probability that the run length of a chart is less than a specified value by adjusting the chart's control limits. However the bootstrap method has the advantage that it can be adapted to other types of control charts including risk-adjusted versions. The bootstrap method is used to analyse the Helping Babies Breath (HBB) training program data to see if there is a clinically relevant change in the probability that an infant survives 24 hours at Haydom Lutheran Hospital, Tanzania. The use of risk-adjusted cumulative sum (CUSUM) charts, with adjusted limits to account for estimation error, shows a decrease in the probability of infants not surviving 24 hours, demonstrating the positive impact of the HBB training program.

Contents

1 Introduction **1**

2 Common Control Charts **3**
2.1 Shewhart Charts 3
2.2 CUSUM Charts 8
2.3 EWMA Charts 10
2.4 Risk Adjusted Charts 11

3 Estimation of Chart Parameters **15**
3.1 Estimation Methods 15
3.2 Examples of Estimated Control Charts 17

4 Impact of Estimation Errors **21**
4.1 Samples with Multiple Observations 21
4.2 Samples with Single Observations 23

5 Accounting for Estimation Error **27**
5.1 Corrections for Shewhart Charts 27
5.1.1 Reduction of ARL Bias 27
5.1.2 Reduction of Exceedance Probability 28
5.2 Bootstrap Method for Guaranteed Conditional Performance 29
5.3 Comparison of Methods which Account for Estimation Error 30

6 Cardiac Surgery Analysis **36**

7 Helping Babies Breath Analysis **39**
7.1 Initial Analysis 39
7.2 Risk-Adjusted CUSUM Chart 44
7.3 Further Analysis of HBB Data 49

8 Conclusion **51**
8.1 Further Work 51

A Data Sets **55**

B Impact of Estimation Error: Additional Information **56**
B.1 Samples with Multiple Observations 56
B.2 Samples with Single Observations 58

C Accounting for Estimation Error Continued **61**
C.1 Reduction of Standard Deviation Bias 61
C.2 Reduction of the Exceedance Probability 61

D Helping Babies Breath Analysis Continued **63**
D.1 HBB Regression Model Variables 63
D.2 Phase I CUSUM charts 63

E R Code **65**
E.1 Impact of Estimation Error 65
E.2 Accounting for Estimation Error 66
E.3 Cardiac Surgery Analysis 69
E.4 Helping Babies Breath 71

1. Introduction

The aim of Statistical Process Control (SPC) is to monitor the variability of some chosen quality characteristic for the output of a given process and signal when the variation is outside the limits of "normal variation". In any process there will be normal variation, or chance causes of variation. This is the natural variation, or background noise, of the process due to many small, unavoidable causes. A process is said to be in statistical control when it is operating with only chance causes of variation. There may also be other kinds of variability present in the output of a process. The sources of variability which are not part of the chance causes of variation are referred to as assignable causes of variation. For basic SPC, if assignable causes of variation are present the process is said to be out of statistical control.

Control charts are a method of indicating when a process may be out of statistical control, as well as being an effective tool in reducing the variability of the process. There are many quantities that can be monitored by control charts. The simplest is an individual measurement, where each sample is a single observation. This is useful when the process has few outputs. One example is monitoring survival rates of patients who have undergone a certain surgery. If a process produces a high volume of output then it may be better to take a sample of size n, after every time interval t. The mean of these samples can then be used to construct a control chart, as well as the sample range or standard deviation.

When a monitored quality characteristic is qualitative or each output has many quality characteristics which need to be monitored, such that each output can either be conforming or nonconforming, a control chart of the number or proportion of nonconforming outputs can be constructed. When monitoring the number of nonconforming units, each sample must have the same size. When monitoring the proportion of nonconforming units then the sample size can vary, however large variations in the sample size could have an effect on the control chart. For units that have input variables which affect the quality characteristic, a risk-adjusted chart can be constructed which takes the assignable variation from the input variables into account.

A typical control chart contains a centre line, a lower control limit and an upper control limit. The centre line is the long run average value of the quality characteristic corresponding to the in-control state, for example the mean of the underlying process distribution. In some applications it may instead be the target value of the quality control statistic. The upper and lower control limits are most commonly a distance of L from the centre line, where L is often expressed in standard deviation (SD) units. If the value of a sample plots within the upper and lower control limits it is assumed that the process is in statistical control. However, if the value of a sample exceeds the control limits, it is assumed that the process is out of statistical control. When this occurs, the process is investigated to see if an assignable cause can be found to account for the large variability.

The centre line and control limits are calculated using the underlying distribution of the quality characteristic, relating to the output when the process is in-control. When the underlying distribution is not known, it is instead estimated from historical process data. An estimated distribution is instead used to calculate the centre line and control limits. This can result in

the control chart resembling an in-control process when the process is actually out-of-control or vice versa.

The use of control charts in SPC carries the problem of unknown parameter estimates which are used to calculate the control limits. The use of estimated parameters results in errors when calculating the centre line and control limits, affecting the accuracy of the charts. Many have acknowledged this estimation error and demonstrated the effect it has upon control charts, however few have looked at how to account for the estimation error.

This thesis will look at the basic control charts, how they are constructed and the errors which occur when estimating the chart parameters. The effects of these estimation errors will be outlined and a comparison of a few methods which have been suggested to account for the errors are discussed. The most applicable method will then be applied to a medical data set with the aim to see if the Helping Babies Breath resuscitation training program has been successful in decreasing infant mortality rates in Haydom Lutheran Hospital, Tanzania.

2. Common Control Charts

The three most common control charts are Shewhart, Cumulative Sum (CUSUM) and Exponentially Weighted Moving Average (EWMA). Each has a different method in calculating the centre line and/or the control limits. These methods will be discussed for the case in which it is assumed that the underlying distribution is known, along with the mean and SD.

Performance measures such as Average Run Length (ARL) and hitting probabilities can be used to assess the control charts. The ARL of a chart is the average number of samples taken until the chart signals that the process is out-of-control. Two values that are of interest when looking at the ARL are the in-control ARL and the out-of-control ARL. The in-control ARL, ARL_0, is the average number of samples until a false out-of-control signal is given when a process is in-control. The out-of-control ARL, ARL_1, is the average number of samples it will take for the chart to give an out-of-control signal once the process is out-of-control (Montgomery, 2013).

The hitting probability of a control chart is the probability that the chart will give an out-of-control signal within T samples (Gandy and Kvaløy, 2013). For Shewhart charts the ARL and hitting probability are easy to calculate. However for the CUSUM and EWMA charts the easiest method to calculate the ARL and hitting probability is through simulation. Both the ARL and hitting probability are affected by the control limits as well as the size, and therefore standard deviation, of each sample. For a specified ARL or hitting probability to be achieved, the control limits and/or sample size have to be adjusted, to achieve the specified performance measure (Montgomery, 2013).

2.1 Shewhart Charts

Most common Shewhart charts are $\bar{x}$ and R charts. The $\bar{x}$ chart, or the control chart for the mean, monitors the mean or average value of the chosen quality characteristic. At time i the sample mean, $\bar{x_i}$, is plotted on the control chart. The sample mean is calculated using

$$\bar{x_i} = \frac{x_{i1} + x_{i2} + \ldots + x_{in}}{n}$$

where x_{ij} is the j^{th} observation of sample i with $j = 1, 2, \ldots, n$ and n is the size of the sample. If the sample mean is within the control limits, the process is in-control. However, if a sample plots outside the control limits, then it is out-of-control and the cause of the variability must be investigated (Saleh et al., 2015).

For the $\bar{x}$ chart the centre line will be the in-control mean of the monitored quality characteristic, and is denoted as μ_0. The Upper Control Limit (UCL) and Lower Control Limit (LCL) are

$$UCL = \mu_0 + L\frac{\sigma_0}{\sqrt{n}} \tag{2.1}$$

and

$$LCL = \mu_0 - L\frac{\sigma_0}{\sqrt{n}} \tag{2.2}$$

where μ_0 is the in-control mean, σ_0 is the in-control SD and L, the critical value, is some constant which is usually set so as to only detect large changes in the process (Grigg and Farewell, 2004). Slight adjustments to the value of L will affect the ARL and the hitting probability of the Shewhart chart.

R charts, or the control chart for the range, monitors the range of each sample when the sample size, n, is greater than 1 (Montgomery, 2013). At time i the range of sample, R_i, is plotted on the control chart. The centre line and control limits of the R chart are dependent on the size of each sample, n, where $n \geq 2$ and, if applicable, the specified ARL. The centre line is given by $d_2\sigma_0$, the $UCL = D_2\sigma_0$ and $LCL = D_1\sigma_0$, where D_1, D_2 and d_2 are some constants dependent on the sample size n. A table displaying the different values of each D_1, D_2 and d_2 can be found in *Appendix VI* of Montgomery (2013).

Basic Example

A Shewhart chart of 25 samples, each of size $n = 5$ has been constructed, where the 125 observations are independent variables drawn from the standard normal distribution. Therefore the mean and SD are $\mu_0 = 0$ and $\sigma_0 = 1$ respectively. The control limits have been calculated using $L = 3$, as this is common practise, giving $ARL_0 = 370$.

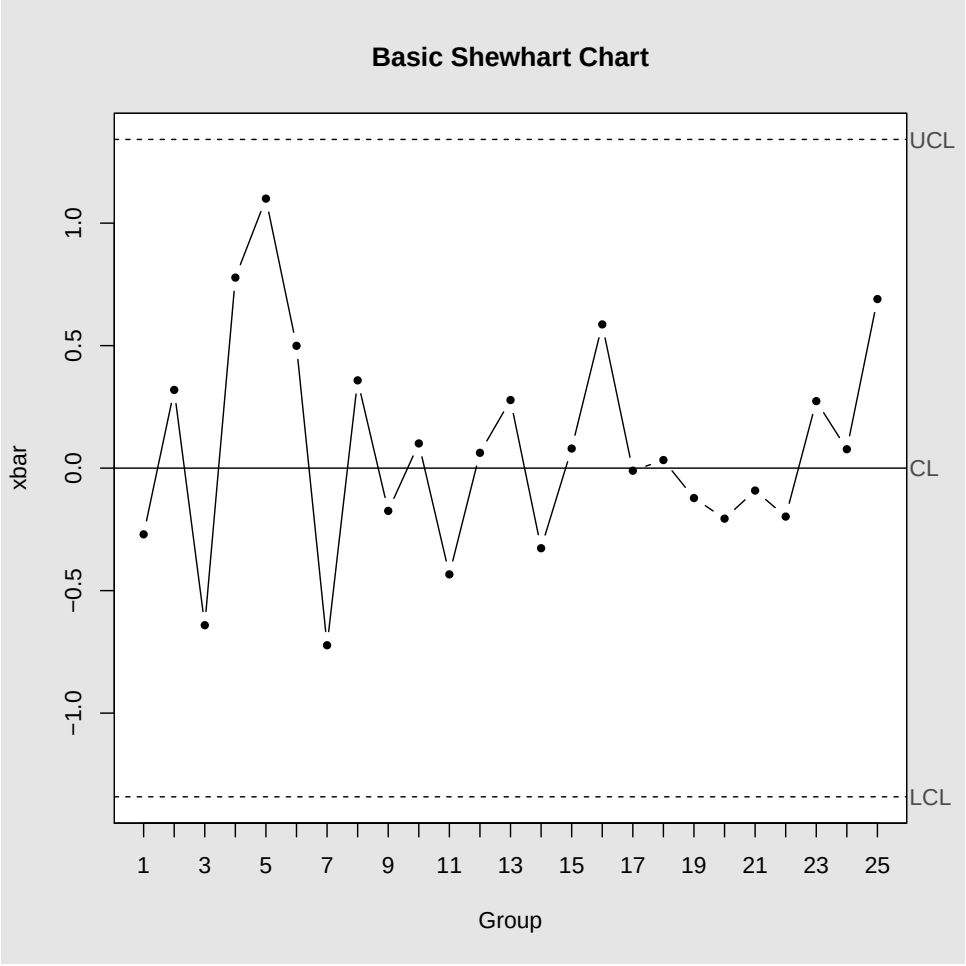

Figure 2.1: Example of a Shewhart control chart for standard normal distribution

An example of a basic Shewhart chart which is in control is shown in Figure 2.1. The

observations used for the data can be found in Table A.1 in Appendix A. The centre line is equal to $\mu_0 = 0$ and the control limits are $LCL = -1.34$ and $UCL = 1.34$. None of the 25 samples plot outside the control limits suggesting that the process is in control.

Performance Measures

The ARL of the Shewhart $\bar{x}$ chart is easy to calculate for any underlying distribution. The in-control ARL is given by $ARL_0 = \frac{1}{\alpha}$, where α is the probability that an in-control sample will plot outside the control limits. The in-control run length follows a geometric distribution, which has the mean $\frac{1}{p}$, where in this case the probability $p = \alpha$. To calculate the ARL_0 for a normal underlying distribution, α must first be found by

$$\alpha = 1 + \Phi\left(\frac{LCL - \mu_0}{\sigma_0/\sqrt{n}}\right) - \Phi\left(\frac{UCL - \mu_0}{\sigma_0/\sqrt{n}}\right),$$

where Φ denotes the standard normal cumulative density function (Yang et al., 2012). However looking at the control limit equations (2.1) and (2.2), it can be seen that this can be simplified to

$$\begin{aligned} \alpha &= 1 + \Phi(-L) - \Phi(L) \\ &= 2 - 2\Phi(L) \end{aligned} \tag{2.3}$$

When the critical value $L = 3$ is used to calculate the control limits of a $\bar{x}$ chart $\alpha = 0.0027$ and the $ARL_0 = 370$. If an ARL_0 or probability of a false out-of-control signal is specified, then

$$L = \Phi^{-1}\left(1 - \frac{1}{2^* ARL_{0,s}}\right) = \Phi^{-1}\left(1 - \frac{\alpha_s}{2}\right), \tag{2.4}$$

where $ARL_{0,s}$ is the specified ARL_0 and α_s is the specified probability.

The out-of-control ARL is denoted by $ARL_1 = \frac{1}{1-\beta}$, where β is the probability of not detecting a change in the process on the first sample after the change has occurred. The shift in the mean and SD are denoted by δ_μ and δ_σ, in terms of σ_0, and have the constraints $-\infty \leq \delta_\mu \leq \infty$ and $0 < \delta_\sigma \leq \infty$. Therefore β is the probability that a sample with mean $\mu_1 = \mu_0 + \delta_\mu\sigma_0$ and standard deviation $\sigma_1 = \delta_\sigma\sigma_0$ falls within the control limits of the $\bar{x}$ chart (Yang et al., 2012). This can be denoted as

$$\beta = P\{LCL \leq \bar{x} \leq UCL | \mu = \mu_1, \sigma = \sigma_1\}. \tag{2.5}$$

Since $\bar{x} \sim N(\mu_1, \sigma_1^2/n)$, equation (2.5) can be expressed as

$$\beta = \Phi\left(\frac{UCL - (\mu_0 + \delta_\mu\sigma_0)}{\delta_\sigma\sigma_0/\sqrt{n}}\right) - \Phi\left(\frac{LCL - (\mu_0 + \delta_\mu\sigma_0)}{\delta_\sigma\sigma_0/\sqrt{n}}\right). \tag{2.6}$$

Substituting equations (2.1) and (2.2) into equation (2.6) and simplifying gives

$$\beta = \Phi\left(\frac{L - \delta_\mu\sqrt{n}}{\delta_\sigma}\right) - \Phi\left(\frac{-L - \delta_\mu\sqrt{n}}{\delta_\sigma}\right). \tag{2.7}$$

Using $L = 3$ and $n = 5$, the ARL_1 is found when a shift occurs such that $\mu_1 = \mu_0 + \sigma_0$ and $\sigma_1 = \sigma_0$. Therefore $\delta_\mu = 1$ and $\delta_\sigma = 1$. Substituting these values into equation (2.7) gives

$$\beta = \Phi\left(3 - \sqrt{5}\right) - \Phi\left(-3 - \sqrt{5}\right) = 0.778$$

The resulting $ARL_1 = \frac{1}{1-0.778} = 4.495$, and so it would take on average 4.495 samples to detect a shift in the mean of σ_0.

The ARL_1 is effected by δ_μ and δ_σ as well as the critical value, L. When the in-control ARL is specified, L, and thus, ARL_1 are adjusted. To demonstrate the effect of the variables on the out-of-control ARL, ARL_1 was calculated for $\delta_\sigma = 1$ and varying δ_μ values when the specified ARL_0 is chosen to be 370. This is then repeated for $ARL_0 = 100$ and $ARL_0 = 1000$.

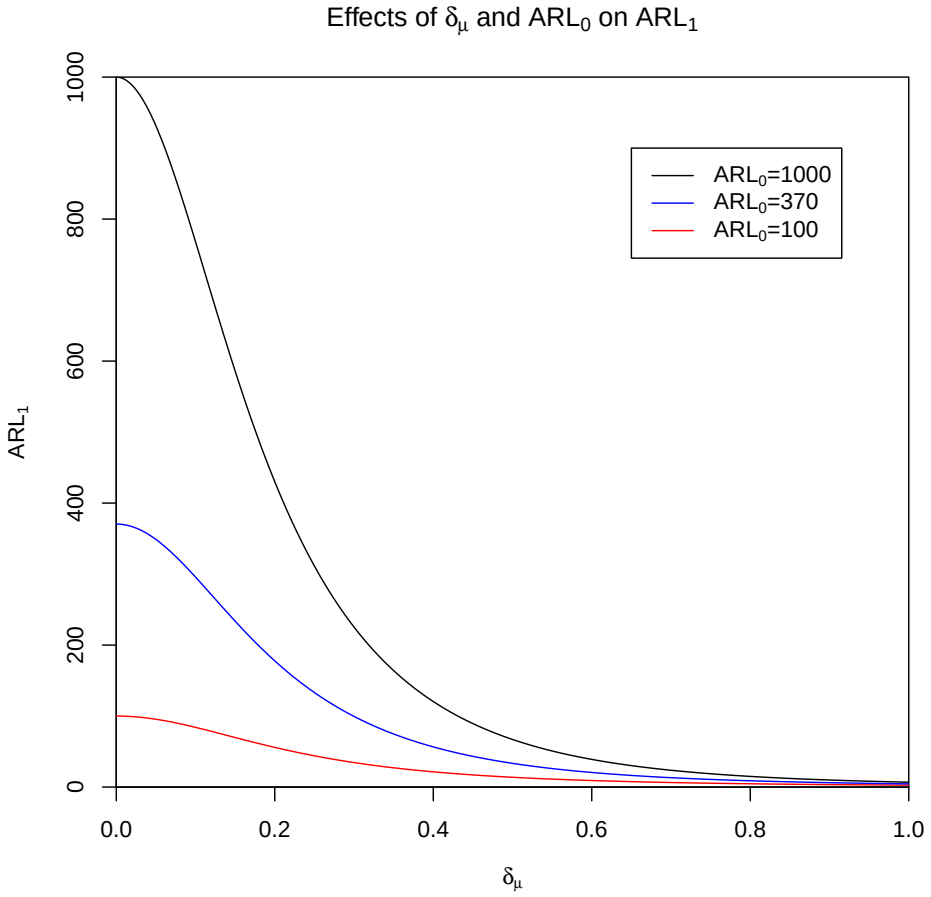

Figure 2.2: ARL_1 for varying ARL_0 and δ_μ

It can be seen in Figure 2.2 that the greater the shift in the mean, δ_μ, the smaller the ARL_1 is and so the quicker a shift will be detected. It can also be seen that the greater the in-control ARL is, the greater the out-of-control ARL will be, and so the longer it will take to detect a shift in the mean of δ_μ. It is common for the shift in the mean to be set in the range of $0.5 \leq \delta_\mu \leq 1$.

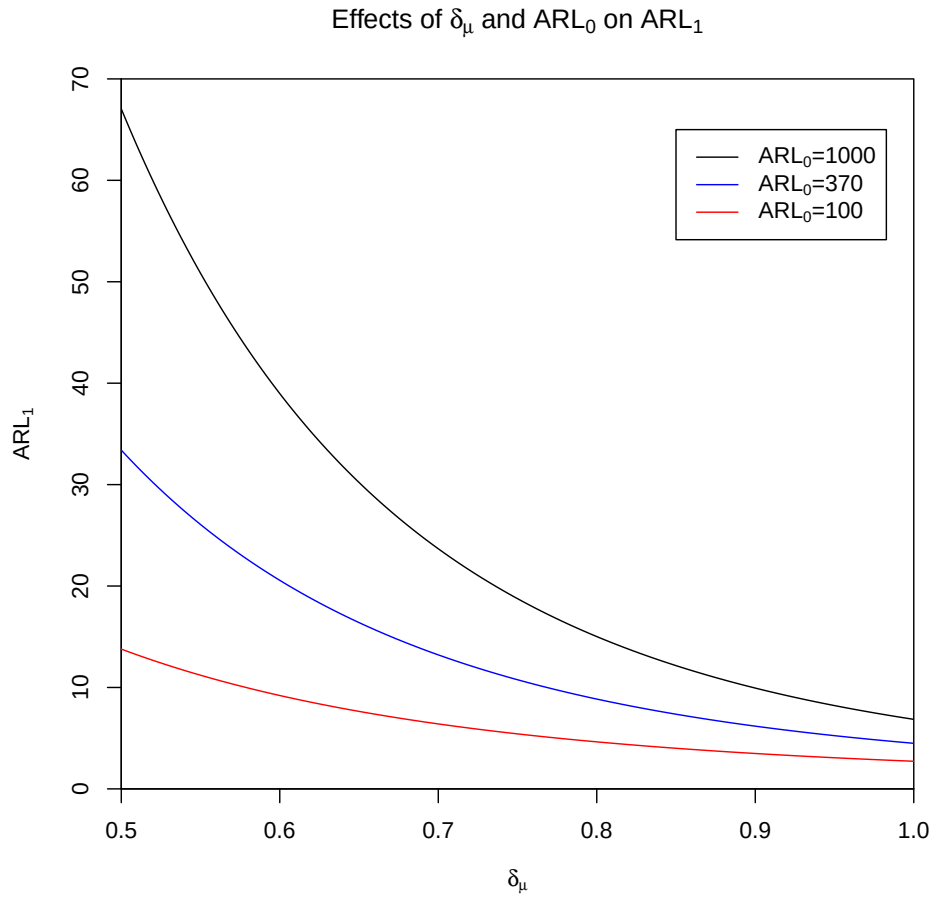

Figure 2.3: ARL_1 for varying ARL_0 and $0.5 \leq \delta_\mu \leq 1$

Figure 2.3 shows more clearly the effect that ARL_0 and δ_μ have on the ARL_1. When $\delta_\mu = 0.5$ the ARL_1 values are approximately 65, 35 and 15 when the in-control ARL is 1000, 370 and 100 respectively. This shows that although a larger ARL_0 is preferable, the corresponding ARL_1 should be taken into consideration when the user defines the in-control ARL.

The hitting probability is the probability that within T samples an out-of-control signal occurs (Gandy and Kvaløy, 2013). The probability that an out-of-control signal occurs for any sample is p, and so the probability that an out-of-control signal does not occur for any sample is $1 - p$. Therefore the hitting probability is

$$P = 1 - (1 - p)^T$$

When a process is in-control, the probability of a sample plotting outside the control limits is α. It follows that the in-control hitting probability is

$$P_0 = 1 - (1 - \alpha)^T \tag{2.8}$$

From equation (2.3) it can be seen that the in-control hitting probability depends on the critical value, L. Taking the example from ARL_0, where $L = 3$ and $\alpha = 0.0027$, then $P_0 = 1 - (1 - 0.0027)^T$. If $T = 100$ then$P_0 = 1 - (1 - 0.0027)^{100} = 0.237$. Thus there is a probability of 0.237 that an out-of-control signal will occur in 100 samples when the process is in-control. If an in-control hitting probability is specified then equation (2.8) can be rearranged to give

$$\alpha = 1 - (1 - P_0)^{\frac{1}{T}}. \tag{2.9}$$

Once α has been calculated, equation (2.4) can then be used to calculate the critical value and thus the control limits.

The probability for any sample that an out-of-control signal will occur when the process is out-of-control is $1 - \beta$. Therefore the out-of-control hitting probability is

$$P_1 = 1 - (1 - (1 - \beta))^T = 1 - \beta^T. \tag{2.10}$$

Equation (2.7) shows that the out-of-control hitting probability is dependent on the critical value L, the sample size n and some specified shift in the mean and/or SD.

2.2 CUSUM Charts

CUSUM charts are better at detecting smaller persistent shifts in the process than Shewhart charts. CUSUM charts incorporate the information collected in all previous samples by plotting the cumulative sum of the deviation of each sample from the expected or target value.

The most common CUSUM chart is the tabular (or algorithmic) CUSUM. The tabular CUSUM is designed to detect a specified change in the quality characteristic of interest. For detecting changes in the mean, the statistic C^+ accumulates the deviations from μ_0 that are above the mean and the statistic C^- accumulates deviations from μ_0 that are below the mean (Grigg et al., 2003). The values C^+ and C^- are known as the one-side upper and lower CUSUMs, and are calculated as

$$C_i^+ = max[0, x_i - (\mu_0 + K) + C_{i-1}^+]$$

and

$$C_i^- = max[0, (\mu_0 - K) - x_i + C_{i-1}^-],$$

where

$$C_0^+ = C_0^- = 0.$$

$K = \frac{|\mu_0 - \mu_1|}{2}$ is known as the reference value, where $\mu_1 = \mu_0 + \delta\frac{\sigma_0}{\sqrt{n}}$ is the out-of-control mean which is specified by the user and $n \geq 1$ is the number of observations in each sample. The process is said to be out-of-control when either C^+ or C^- exceeds a threshold, H. The values of K and H can also be expressed in terms of σ_0 such that $K = k\frac{\sigma_0}{\sqrt{n}}$ and $H = h\frac{\sigma_0}{\sqrt{n}}$, where n is the number of observations in each sample and h is the critical value. If only the deviation of the process either above or below the mean is to be detected, a one-sided CUSUM chart can be constructed by plotting just the relevant one-sided CUSUM values.

The CUSUM can also be standardised such that

$$C_i^+ = max[0, y_i - k + C_{i-1}^+] \tag{2.11}$$

$$C_i^- = max[0, -k - y_i + C_{i-1}^-], \tag{2.12}$$

where $y_i = \frac{x_i - \mu_0}{\sigma_0/\sqrt{n}}$ and $k = \frac{\delta}{2}$. The control limit must also be standardised and thus becomes $\frac{H}{\sigma/\sqrt{n}} = h$. From now on any CUSUM chart will be computed as a standardised CUSUM chart.

Performance Measures

The values of k and h effect the ARL and hitting probability of the chart and can be slightly altered to improve the chart performance (Montgomery, 2013). Unfortunately the ARL and hitting probability of the CUSUM chart are more complicated to calculate than the those of the Shewhart chart. The most straightforward approach to finding the ARL of the CUSUM

chart is the Siegmund Approximation which is outlined by Montgomery (2013, pg.423). This can provide an estimate for the ARL of each one-sided upper and lower CUSUM, ARL^+ and ARL^-, respectively. The ARL of a two-sided CUSUM chart is then found using *equation (9.7)* in Montgomery (2013, pg. 423) which states

$$\frac{1}{ARL} = \frac{1}{ARL^+} + \frac{1}{ARL^-}.$$

Other methods to calculate the ARL of the CUSUM chart are the Markov Chain approach and the Integral Equation method. Brook and Evens (1972) and Fu et al. (2002) outline the Markov Chain approach. For further information on the Integral Equation methods see Rao et al. (2001) and Champ et al. (2001).

Basic Example

The 25 samples from the standard normal distribution found in Table A.1 in Appendix A were used to produced an example of a standardised basic CUSUM chart. The control limits were set to $h = 5$ and the specified out-of-control mean was chosen such that $k = 0.5$. These parameters give an $ARL_0 = 465$, which has been calculated by using the spc R package (Knoth, 2015), where the ARL is estimated using the integral equation method.

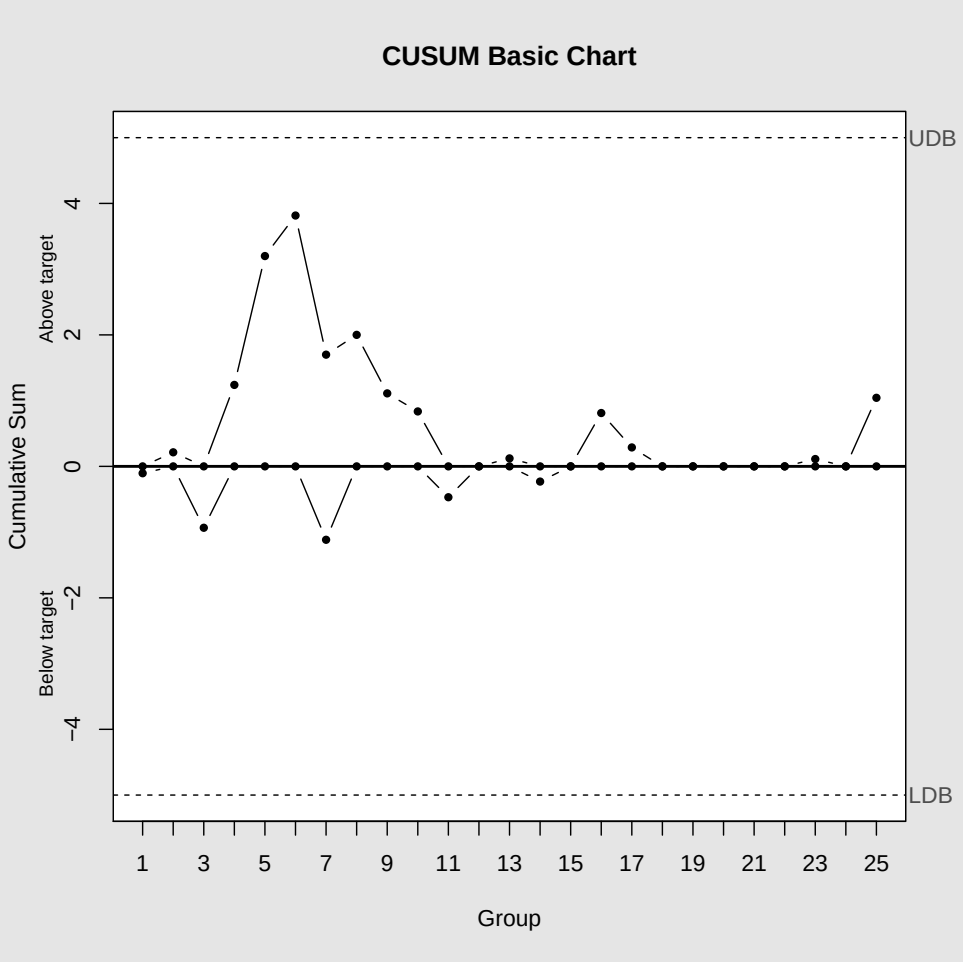

Figure 2.4: Example of a standardised CUSUM chart for standard normal distribution

Figure 2.4 is a two-sided CUSUM chart with C_i^+ plotted above 0 and C_i^- below 0. The chart suggests that the process is in-control for all 25 samples.

2.3 EWMA Charts

Exponentially Weighted Moving Average (EWMA) charts give each sample a weight, depending on the age of the sample; the older the sample, the less weight it is given. Like the CUSUM, the EWMA is preferred to the Shewhart chart when the detection of smaller process shifts is required. The exponentially weighted average is given as

$$z_i = \lambda \bar{x}_i + (1 - \lambda) z_{i-1}, \tag{2.13}$$

where $\bar{x}_i$ is the mean of the sample taken at time i and λ is some constant between 0 and 1.

To construct the EWMA chart, the exponentially weighted average, z_i, is plotted on the control chart. The centre line of the chart is μ_0. The control limits are calculated by

$$UCL = \mu_0 + L\sigma_{z_i}$$

$$LCL = \mu_0 - L\sigma_{z_i},$$

where

$$\sigma_{z_i} = \sigma_0 \sqrt{\frac{\lambda}{(2-\lambda)n}[1 - (1-\lambda)^{2i}]}, \tag{2.14}$$

and L is the critical value. In Equation 2.14, the term $[1 - (1 - \lambda)^{2i}]$ approaches unity as i increases and therefore will sometimes lead to the standard deviation of z_i being defined as $\sigma_z = \sigma_0 \sqrt{\frac{\lambda}{(2-\lambda)n}}$ (Montgomery, 2013; Serel and Moskowtiz, 2008). This is a simpler calculation but may lead to an early out-of-control signal not being detected. In cases where only the deviations above or below the mean are to be monitored, a one-sided EWMA chart can be used instead. For further information on one-sided EWMA charts see Shu et al. (2007).

Performance Measures

The ARL and the hitting probability of the EWMA chart are dependent on the variables L and λ. As with the CUSUM chart it is difficult to calculate the ARL and hitting probability. See Chen and Chen (2007) for more information on the methods which can be used. The ARL and hitting probability can also be estimated using simulation, where varying the values of L and λ shows the effect on the ARL and hitting probability.

Basic Example

The 25 samples used to construct this example can be found in Table A.1 in Appendix A. The control limits were calculated using $n = 5$, $L = 3$ and $\lambda = 0.2$. These parameters give an $ARL_0 = 560$, which has been calculated by using the spc R package (Knoth, 2015), where the ARL is estimated using the integral equation method.

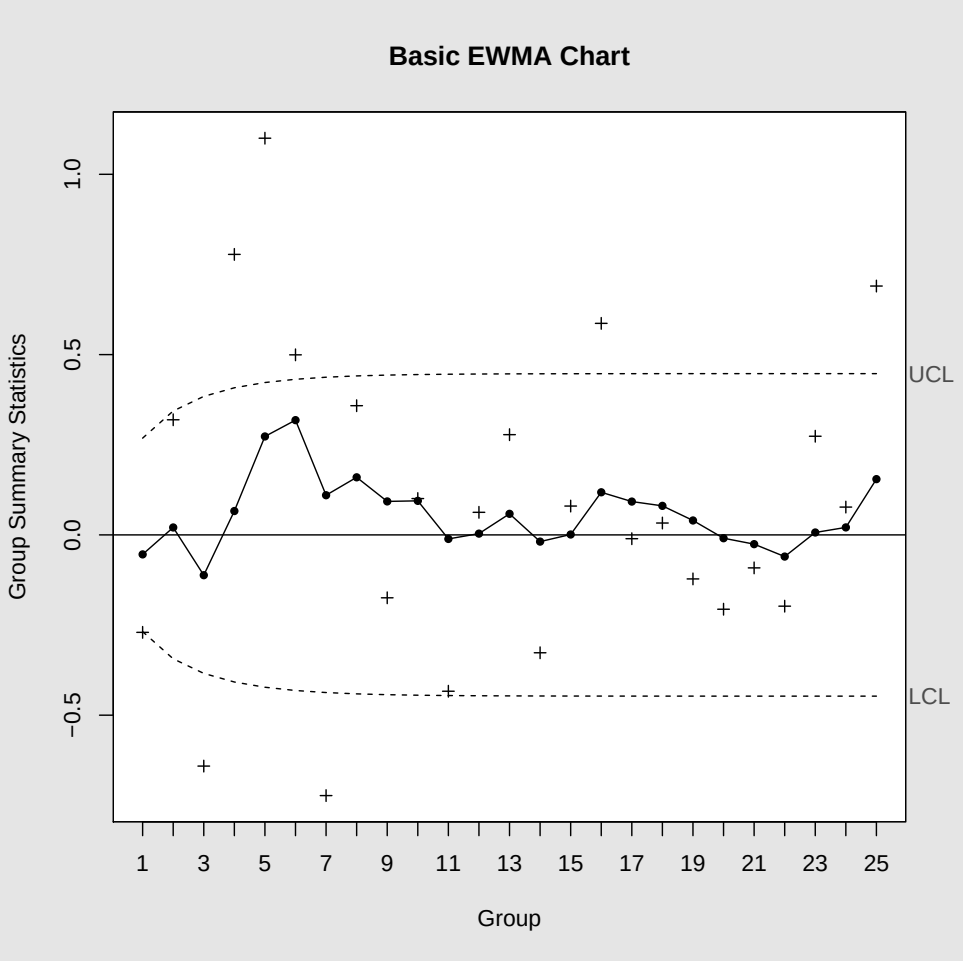

Figure 2.5: Example of a EWMA chart for standard normal distribution

The control chart in Figure 2.5 suggests that the process is in-control. The control limits have been constructed using equation 2.14. This can be seen from the expansion of the limits for the first 7 samples until they reach a steady state. The crosses on the chart represent the $\bar{x}_i$ value for each sample.

2.4 Risk Adjusted Charts

When the data is heterogeneous then the explainable variation present in the samples must be accounted for. In such cases risk-adjusted control charts are used. A regression model is applied to the sampled data, which takes into account the explainable variation. Therefore when a risk-adjusted chart signals that the process is out-of-control, this means that some assignable variation is present in the process which has not been included in the regression model (Gandy and Kvaløy, 2013). The residuals of the regression model are used, rather than the observed data, to construct the control chart.

Linear Regression Example

To demonstrate a risk-adjusted chart the Cascade process data shown in *Table 11.5* of Montgomery (2013, p.529) has been used. The Cascade process data has 40 samples which each have two outputs y_1 and y_2. However for this example only the y_1 output will be used, and will be

referred to as y. For each y_i there are nine input process variables, x_1 to x_9. The estimated mean of the data is $\hat{\mu}_0 = \frac{\Sigma y_i}{m} = 952.78$, where $m = 40$. The estimated SD of the data is $\hat{\sigma}_0 = 0.743$ and was calculated using the Moving Range method which is explained in Chapter 3. A standardised CUSUM chart of y was first constructed in R using the formulas

$$C_i^+ = max\left[0, \frac{y_i - \hat{\mu}_0}{\hat{\sigma}_0} + k + C_{i-1}^+\right]$$

$$C_i^- = max\left[0, -k - \frac{y_i - \mu_0}{\hat{\sigma}_0} + C_{i-1}^-\right],$$

where $C_0^+ = C_0^- = 0$, $k = 0.5$ and the control limits are $h = 5$.

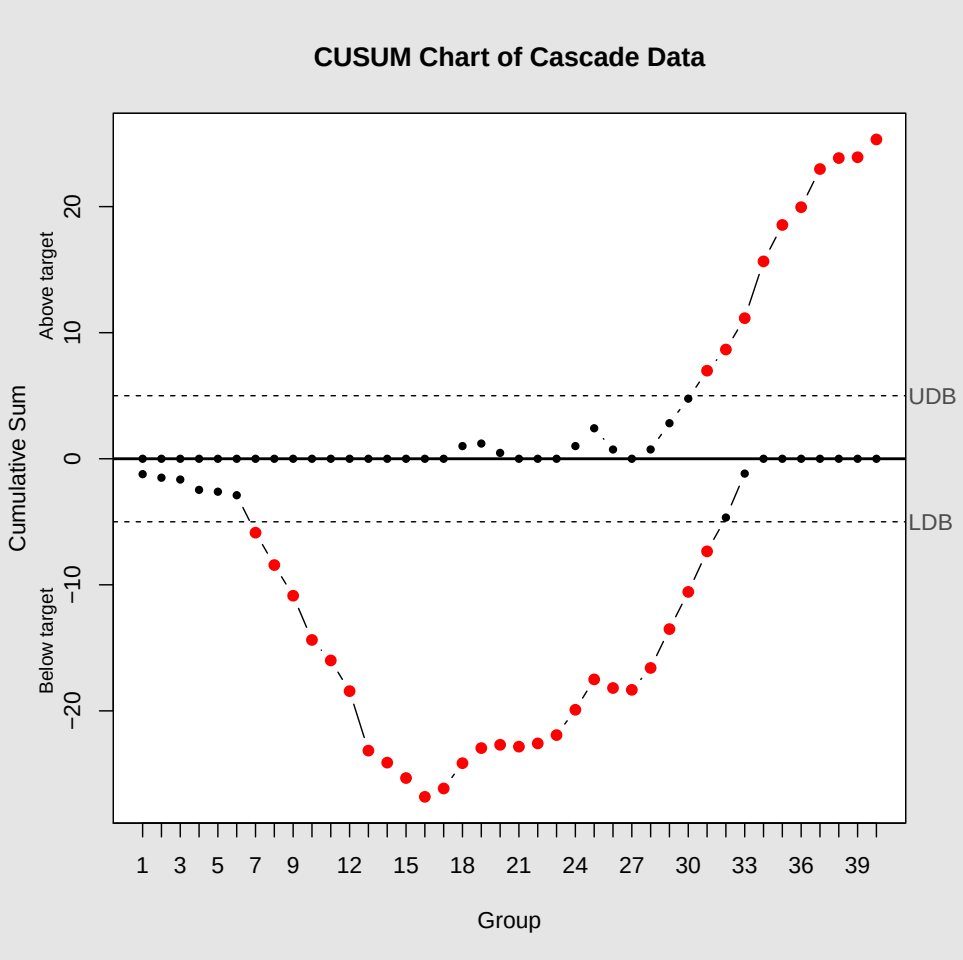

Figure 2.6: Standardised CUSUM chart of cascade process data

From Figure 2.6 it can be seen that most of the points plot out-of-control. However as the data has the nine x_i input variables a regression model can be fitted to the data. For this example a linear regression model was fitted, and the following equation was used to calculate the estimated data points:

$$\hat{y} = 825.89 + 0.474x_1 + 1.441x_2 - 0.117x_3 - 0.082x_4 - 2.39x_5 - 1.3x_6 + 2.18x_7 + 2.98x_8 + 113.22x_9$$

The residuals of the linear regression were then used as the data points for a standardised CUSUM chart with the equations

$$C_i^+ = max\left[0, \frac{r_i - 0}{\hat{\sigma}_r} - k + C_{i-1}^+\right] C_i^- = max\left[0, -\frac{r_i - 0}{\hat{\sigma}_r} - k + C_{i-1}^-\right]$$

where r_i is the residual of each data point, the estimated mean of the residuals is $\hat{\mu}_r = 0$, $k = 0.5$ and the estimated SD of the residuals is $\hat{\sigma}_r = 0.663$. The control limits are $h = 5$ which corresponds to an $ARL_0 = 465$.

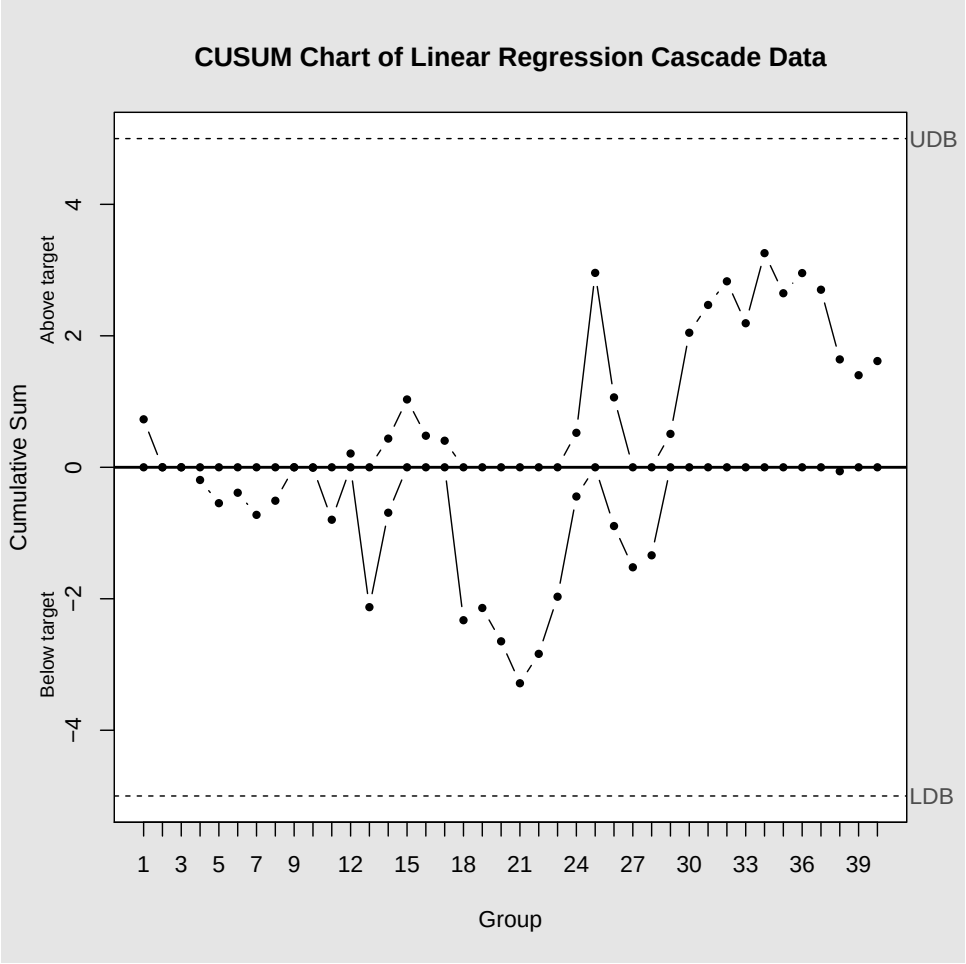

Figure 2.7: Standardised CUSUM chart of linear regression cascade process data

Comparing Figures 2.6 and 2.7, it can be seen that when the input process variables are taken into account the CUSUM chart appears to be in-control during all 40 data points. This is a vast improvement on the standardised CUSUM chart where no input process variables are taken into account.

Logistic Regression

For risk-adjusted charts which monitor binary data, a logistic regression can be used to account for the explainable variation present in the samples. These charts are favoured in the health care industry where the outcome of a procedure is considered to be either successful or unsuccessful. Explainable variation in the data due to patient differences can be accounted for in the regression

model and thus decrease the number of false signals compared to when a basic control chart is used. An example of this is using the Parsonnet score as the input variable when the 30-day mortality for patients after cardiac surgery is being monitored. The Parsonnet score is the risk score of a patient and depends on factors such as age, health status and type of operation the patient has undergone (Grigg et al., 2003). In logistic regression models the log likelihood ratio, rather than the observed minus expected data, is used to construct the control chart.

For binary variables Y_i, $i = 1, ..., n$, which are the observed responses in the process, the likelihood ratio needs to be defined to produce a risk-adjusted chart. To find the likelihood ratio, the logistic or logit regression model for the data must first be found. The logit model is the estimation of the log of the odds. This gives

$$\text{logit}\,(p_i) = \ln\left(\frac{p_i}{1-p_i}\right) = X_i\beta, \tag{2.15}$$

where p_i is the probability that the i$^{\text{th}}$ sample is successful, such that $Y_i = 1$. X_i is the vector of input variables and β is the coefficient for the input variables in the regression model (Steiner et al., 2000). Equation 2.15 can be rearranged to give the probability that $Y_i = 1$ as

$$p_i = \frac{\exp\left(X_i\beta\right)}{1+\exp\left(X_i\beta\right)}.$$

If the variable Y_i follows the Bernoulli distribution, the probability of the outcome is defined as $P\left(y\right) = p^y\left(1-p\right)^y$, where y is the binary variable. The likelihood ratio, LR, is then calculated by dividing the likelihood for the out-of-control situation, L_1, by the likelihood for the in-control situation, L_0. L_0 can be calculated by

$$\begin{aligned} L_0 &= p^y\left(1-p\right)^{1-y} \\ &= \left(\frac{\exp\left(X\beta\right)}{1+\exp\left(X\beta\right)}\right)^y\left(\frac{1}{1+\exp\left(X\beta\right)}\right)^{(1-y)} \\ &= \frac{\left(\exp\left(X\beta\right)\right)^y}{1+\exp\left(X\beta\right)} \end{aligned}$$

To detect a change to the logit model for a out-of-control situation, the logit function can be given as

$$\text{logit}\,(p_i) = \Delta + X_i\beta. \tag{2.16}$$

where Δ is the change in the logit such that the process is out-of-control and is defined by the user. Using the logit function as given in Equation 2.16, the likelihood for the out-of-control situation can be calculated as

$$L_1 = \frac{\exp\left(\Delta + X\beta\right)^y}{1+\exp\left(\Delta + X\beta\right)}$$

The likelihood ratio for a single observation is then expressed as

$$LR_i = \frac{\exp\left(\Delta\right)^{Y_i}\left(1+\exp\left(X_i\beta\right)\right)}{1+\exp\left(\Delta + X_i\beta\right)}$$

To construct the risk-adjusted CUSUM chart the C statistic is $C_i = max\left(0, C_{i-1} + R_i\right)$, where $R_i = \ln(LR)$ and $C_0 = 0$ (Gandy and Kvaløy, 2015). When the out-of-control situation is more likely than in the in-control situation, $R_i > 0$, which gives a larger C_i value than when the situation is more likely to be in-control. To monitor an increase in the odds when the process in out-of-control Δ should be set to greater than 0. To monitor a decrease in the odds when the process is out-of-control, Δ should be set to less than 0.

3. Estimation of Chart Parameters

When using control charts in practice, the underlying distribution is generally unknown, and therefore must be estimated. The development of control charts is then divided into two phases. In Phase I the underlying distribution is estimated from historical data. Using the m historical samples, of size n, estimates of the unknown parameters are calculated. Phase I control charts are then constructed by using the methods outlined in Chapter 2, with the known mean and SD, μ and σ, being replaced with the calculated estimations, $\hat{\mu}$ and $\hat{\sigma}$.

If the Phase I control chart shows that all of the m historical samples plot in-control then the estimated parameters are used to monitor samples in Phase II. If an out-of-control signal does occur, the process is investigated for assignable causes. The sample that plotted out-of-control is then removed from the Phase I data and the parameter estimates are revised. This process is repeated until all the samples used to calculate the estimations indicate an in-control signal. These estimates are then used in Phase II to construct a control chart to monitor the performance of the process (Montgomery, 2013).

3.1 Estimation Methods

When the underlying distribution is assumed to follow the normal distribution, there are several methods to estimate the unknown parameters of μ and σ. To estimate the mean,

$$\hat{\mu}_0 = \frac{\sum_{i=1}^{m} \bar{x}_i}{m}, \tag{3.1}$$

where $\bar{x}_i = \frac{\sum_{j=1}^{n} x_{ij}}{n}$, the mean of each sample, or if $n = 1$, then $\bar{x}_i = x_i$, the individual observation. To estimate the standard deviation there are a variety of methods, of which three will be discussed. For an overview on other methods see Saleh et al. (2015).

Method I

Method I involves either the range of each sample, when $n > 1$, or the moving range of the m samples when $n = 1$. When $n > 1$, the average of the sample ranges,

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

is used to estimate the SD such that

$$\hat{\sigma}_0 = \frac{\bar{R}}{d_2(n)}, \tag{3.2}$$

where R_i is the range of sample i and $d_2(n)$ is some constant which depends on the size of each sample (Saleh et al., 2015). A table containing the values of d_2, depending on various values of n can be found in *Appendix VI* of Montgomery (2013).

When $n = 1$, the moving range is used instead of the range and is defined as

$$\overline{MR} = \frac{\sum_{i=1}^{m-1} MR_i}{m-1},$$

where $MR_i = |x_{i+1} - x_i|$. The SD estimate is then calculated as

$$\hat{\sigma}_0 = \frac{\overline{MR}}{d_2(2)} = \frac{\overline{MR}}{1.128}. \tag{3.3}$$

For the moving range estimate, $d_2(n)$ is used where $n = 2$, as the moving range of two samples is used (Montgomery, 2013).

Method II

Method II uses the sample variances to estimate the SD and is referred to as the S_{pooled} method in Saleh et al. (2015). When $n > 1$, the samples variance of each historical sample is calculated as

$$s_i^2 = \frac{\sum_{j=1}^{n} (x_{ij} - \bar{x}_i)^2}{n-1}$$

From this the SD estimate is found using

$$\hat{\sigma}_0 = \sqrt{\frac{\sum_{i=1}^{m} s_i^2}{m}} \tag{3.4}$$

When $n = 1$ the data is treated as one sample with m observations, instead of m samples which each contain 1 observation. From Equation 3.4, it can be seen that if m is substituted with n, and $n = 1$, the SD estimate is then just square root of the sample variance. Therefore

$$\hat{\sigma}_0 = \sqrt{\frac{\sum_{i=1}^{m} (x_i - \bar{x})^2}{m-1}} \tag{3.5}$$

where x_i is the i^{th} observation and $\bar{x} = \frac{\sum_{i=1}^{m} x_i}{m}$.

Method III

Another method of calculating the SD is to adjust the biased sample SD by dividing by a constant, which is dependent on n, to give an unbiased SD estimate. When $n > 1$, the unbiased SD estimate is given as

$$\hat{\sigma}_0 = \frac{\bar{s}}{c_4(n)}, \tag{3.6}$$

where $\bar{s} = \frac{\sum_{i=1}^{m} s_i}{m}$ and $c_4(n)$ is a constant which depends on n, the values of which can be found in *Appendix VI* of Montgomery (2013). When $n = 1$ the data is treated as one sample with m

observations, rather than m samples that each have one observation. This then gives $\bar{s} = s$ and so

$$\hat{\sigma}_0 = \frac{s}{c_4(m)}, \tag{3.7}$$

where $s = \sqrt{\frac{\sum_{i=1}^{m}(x_i - \bar{x})^2}{m-1}}$ and $\bar{x} = \frac{\sum_{i=1}^{m} x_i}{m}$

3.2 Examples of Estimated Control Charts

The Phase I data for the process which manufactures engine piston rings can be found in *Table 6.3* of Montgomery (2013, pg.260). The data consists of 25 samples, each with 5 observations, where the quality characteristic is the diameter of the ring. To estimate the SD, method III is used, and so Equations 3.1 and 3.6 are used to estimated the mean and SD respectively.

Table 3.1: Summary of Mean, SD Estimation and Standardised Values of Piston Data

Sample Number	$\bar{x}_i$	s_i	y_i
1	74.01020	0.0148	2.1457
2	74.00060	0.0075	-0.1370
3	74.00800	0.0147	1.6226
4	74.00300	0.0091	0.4337
5	74.00340	0.0122	0.5288
6	73.99560	0.0087	-1.3259
7	74.00000	0.0055	-0.2796
8	73.99680	0.0123	-1.0405
9	74.00420	0.0055	0.7190
10	73.99800	0.0063	-0.7552
11	73.99420	0.0029	-1.6587
12	74.00140	0.0042	0.0533
13	73.99840	0.0105	-0.6601
14	73.99020	0.0153	-2.6099
15	74.00600	0.0073	1.1470
16	73.99660	0.0078	-1.0881
17	74.00080	0.0106	-0.0894
18	74.00740	0.0070	1.4799
19	73.99820	0.0085	-0.7076
20	74.00920	0.0080	1.9079
21	73.99980	0.0122	-0.3272
22	74.00160	0.0074	0.1008
23	74.00240	0.0119	0.2910
24	74.00520	0.0087	0.9568
25	73.99820	0.0162	-0.7076
Σ	1850.029	0.235	
$\bar{x}$	74.001	$\bar{s}$ 0.0094	

Table 3.1 shows the estimated mean and average of the sample SD estimation for the Piston

diameter data. From this the estimated SD can be calculated using

$$\hat{\sigma}_0 = \frac{0.0094}{c_4(5)} = \frac{0.0094}{0.9896} = 0.0095.$$

To construct the Phase I Shewhart control chart the in-control ARL was specified as $ARL_0 = 350$. This gives the critical value of $L = 2.98$. Using this with the estimated mean, $\hat{\mu}_0 = 74.001$, as the centre line, the upper and lower control limits are

$$UCL = \hat{\mu}_0 + L\frac{\hat{\sigma}_0}{\sqrt{n}} = 74.001 + 2.98(0.0097) = 74.014$$

and

$$LCL = \hat{\mu}_0 - L\frac{\hat{\sigma}_0}{\sqrt{n}} = 74.001 - 2.98(0.0097) = 73.99.$$

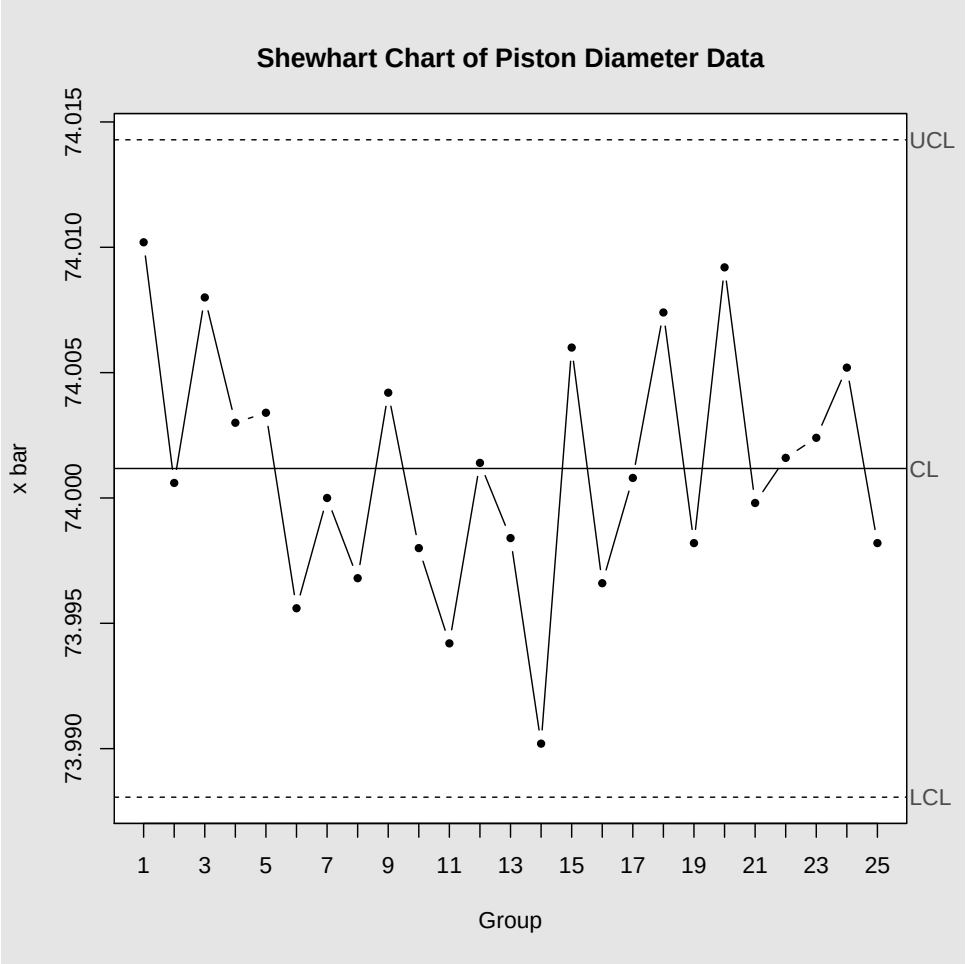

Figure 3.1: Example of a Shewhart control chart with estimated parameters

From Figure 3.1 it can be seen that using the estimated mean and SD, all 25 samples appear to be in-control and so the centre line and control limits calculated for this chart would be used to monitor the piston ring diameter in Phase II.

To construct a CUSUM chart for the piston data, standardised values of the data, the y_i values in Table 3.1, are found using the estimated mean and SD such that

$$y_i = \frac{\bar{x}_i - \hat{\mu}_0}{\hat{\sigma}_0\sqrt{n}} = \frac{\bar{x}_i - 74.001}{0.0095/\sqrt{5}}.$$

For a specified $ARL_0 = 350$, the critical value is $h = 4.72$ and the upper and lower CUSUM values are calculated using equations 2.11 and 2.12. To calculate the critical value, h, the spc R package was used (Knoth, 2015), which estimates the ARL using the integral equation method.

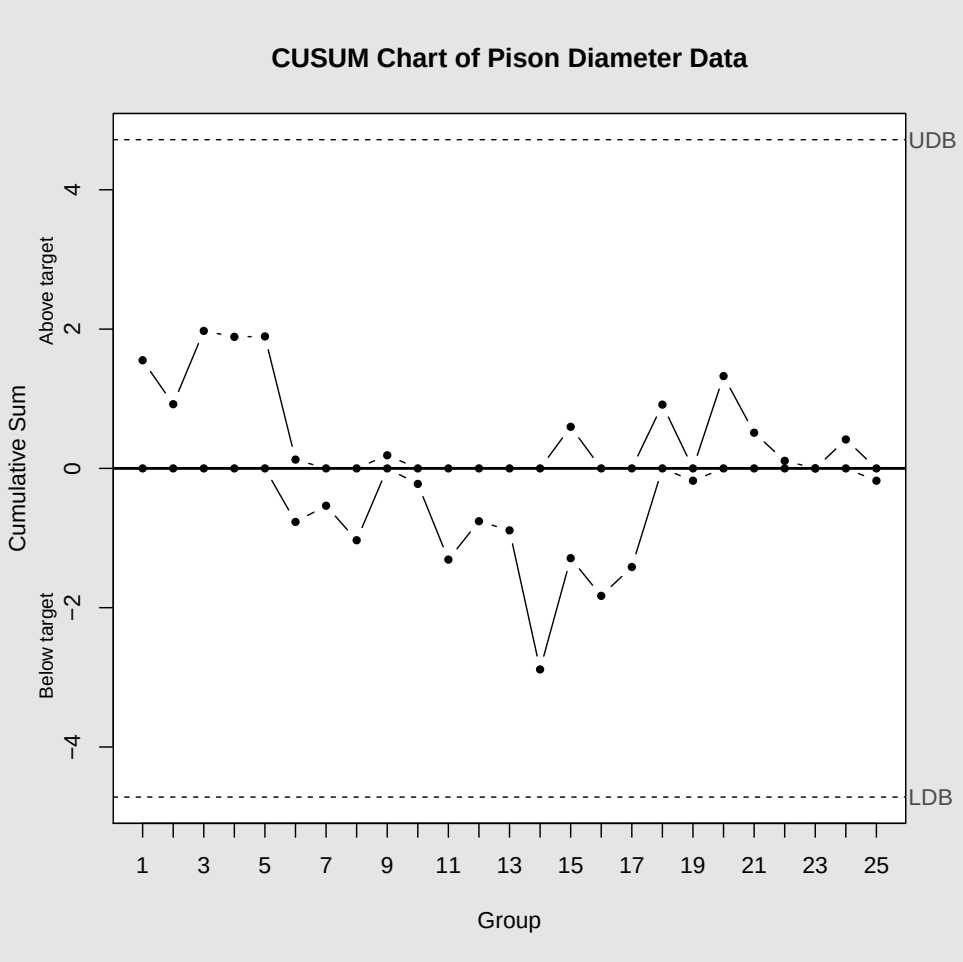

Figure 3.2: Example of a CUSUM control chart with estimated parameters

Figure 3.2 illustrates that all the Phase I data points appear to be in-control and so the estimated mean and SD of $\hat{\mu}_0 = 74.001$ and $\hat{\sigma}_0 = 0.0098$ can be used to construct the Phase II control chart to monitor the diameter of piston rings.

To construct an EWMA chart for the piston ring data, the mean of each sample, $\bar{x}_i$, is needed as well as the estimated mean and SD. The exponentially weighted average is calculated using Equation 2.13 is plotted on a EWMA chart with a centre line which is equal to $\hat{\mu}_0$ and control limits

$$UCL = \hat{\mu}_0 + L\sigma_{z_i}$$
$$LCL = \hat{\mu}_0 - L\sigma_{z_i},$$

where

$$\sigma_{z_i} = \hat{\sigma}_0 \sqrt{\frac{\lambda}{(2-\lambda)n}[1-(1-\lambda)^{2i}]}.$$

For an in-control ARL of $ARL_0 = 350$, a critical value of $L = 2.84$ is used to calculate the control limits. This critical value was found using the spc R package (Knoth, 2015), which uses the integral equation method to estimate the ARL.

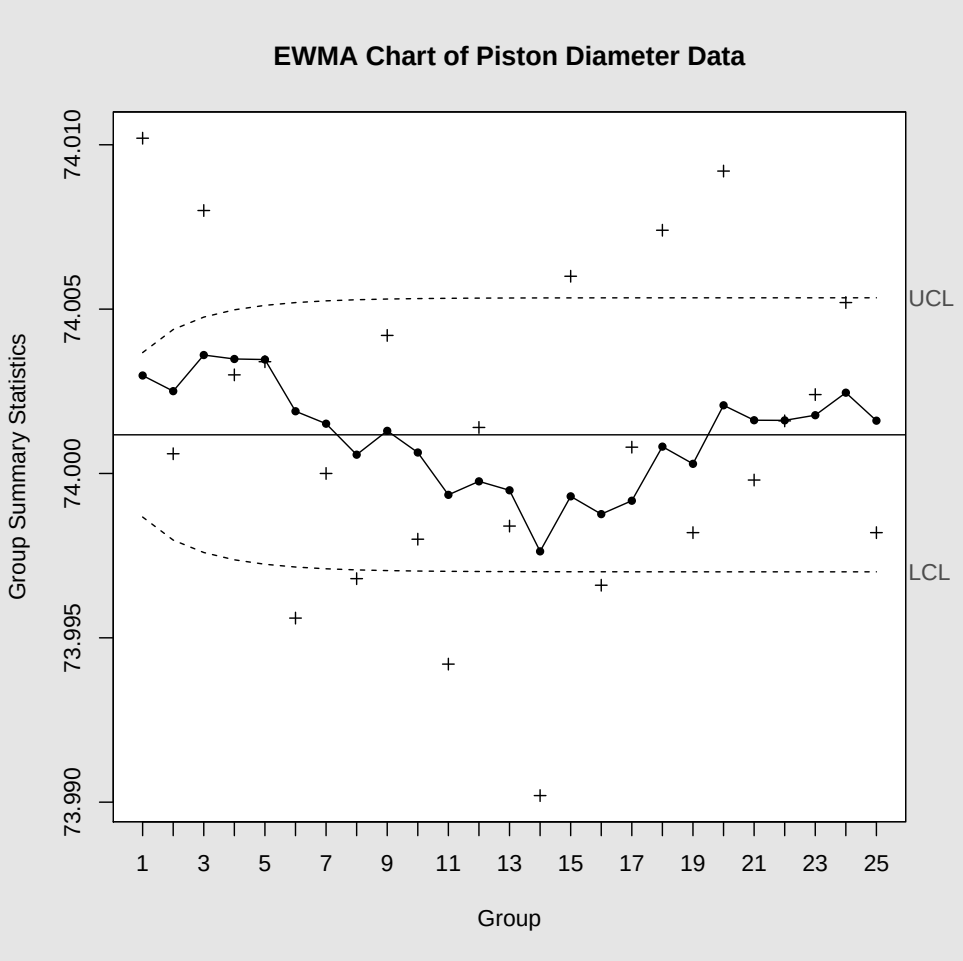

Figure 3.3: Example of a EWMA control chart with estimated parameters

Again the EWMA chart in Figure 3.3 shows that all 25 Phase I samples are in-control and so the centre line and control limits calculated for Figure 3.3 would be used to monitor the data in Phase II. The crosses at each group i represent the mean for each sample, $\bar{x}_i$.

4. Impact of Estimation Errors

When the underlying distribution is unknown and thus chart parameters are estimated, the estimations can effect the efficiency of the control chart. The impact that the estimation error has on the chart efficiency has been discussed by many, including Jensen et al. (2006) and Jones and Steiner (2011). To demonstrate the impact that estimating the underlying distribution can have, the true in-control ARL is calculated using estimated parameters for samples drawn from a known distribution. The following generic algorithm is used to illustrate the impact of the estimation error for normally distributed data.

Algorithm 1 Estimation Error Impact

1. Generate Phase I sample data from normal distribution with known μ_0 and σ_0.
2. Estimate $\hat{\mu}_0$ and $\hat{\sigma}_0$ from Phase I data.
3. Calculate critical value for a specified in-control ARL using $\hat{\mu}_0$, $\hat{\sigma}_0$ and specified reference value if applicable.
4. Calculate true in-control ARL ($ARL_{0,true}$), with the critical value from the estimated chart parameters, $\hat{\mu}_0$ and $\hat{\sigma}_0$, and an in-control data stream with known μ_0 and σ_0 and specified reference value if applicable.
5. Repeat B times.

The critical value for Shewhart and EWMA charts has been previously denoted as L, while the critical value for the CUSUM charts has been denoted as h. For each use of Algorithm 1, the in-control ARL was specified as 100, and the algorithm was repeated 1000 times. The Phase I data was generated from the known standard normal distribution, such that $\mu_0 = 0$ and $\sigma_0 = 1$. When the true in-control ARL of a CUSUM chart is investigated, the reference value $k = 0.5$ has been used, unless otherwise specified.

4.1 Samples with Multiple Observations

It is stated in Montgomery (2013, pg.239) that "it is desirable to have 20-25 samples of size n" for the Phase I data, where "n is typically between 3 and 5". The effect of the number of samples, m, on the estimated parameters has been investigated and discussed in detail, see Saleh et al. (2015) for more information. To illustrate how the number of Phase I samples affects the ARL_0, samples containing 5 observations were generated with the number of samples varying from 20 to 2000. Algorithm 1 was then used to illustrate the effect of Phase I sample size on the in-control ARL of both Shewhart and CUSUM charts.

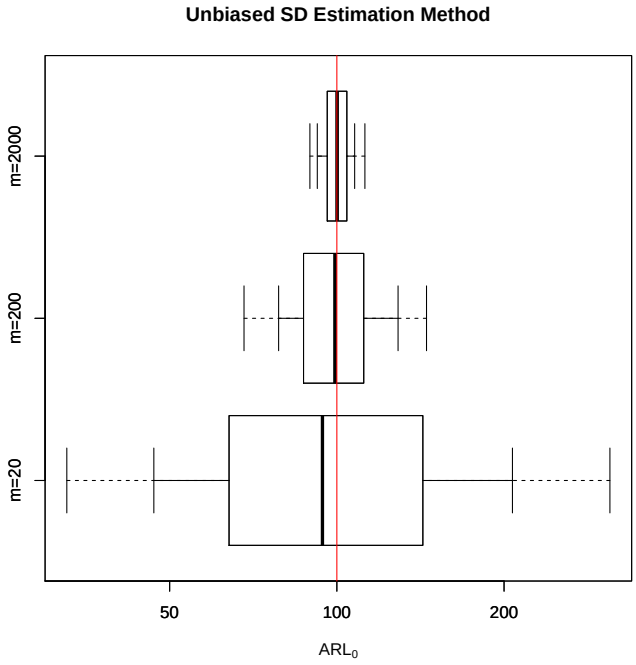

Figure 4.1: $ARL_{0,true}$ distribution of Shewhart control charts with estimated parameters for Phase I data of sample size 20, 200 and 2000, where each sample has 5 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the horizontal axis.

Figure 4.1 shows the true in-control ARL of Phase I samples for Shewhart charts when the number of samples vary. The SD of each sample was calculated using *Method III* as shown in Chapter 3. As the number of samples increases, the mean of the $ARL_{0,true}$ becomes closer to the specified ARL_0, however all three boxplots have a mean which is close to the specified in-control ARL. The main difference between the boxplots is that, as the number of Phase I samples increases, the variation of the true in-control ARL decreases. For $m = 20$ the variation in the true in-control ARL is very large. When the true in-control ARL is considerably smaller than the specified in-control ARL, the probability of a false out-of-control signal increases. When the true in-control of the ARL is considerably larger than the specified in-control ARL, the probability of a false in-control signal increases. These false signals can have a profound effect on the reliability of the control chart. The impact of parameter estimation using each of the three SD estimation methods were very similar and can be seen in Appendix B.

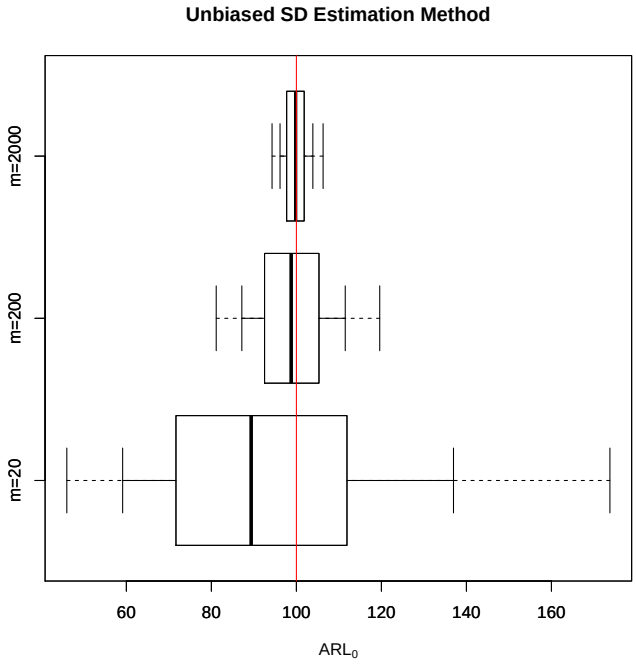

Figure 4.2: $ARL_{0,true}$ distribution of CUSUM control chart with estimated parameters for Phase I data of sample size 20, 200 and 2000, where each sample has 5 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution

Figure 4.2 shows the true in-control ARL of Phase I for CUSUM charts when the number of samples varies. The SD of each sample used to produce Figure 4.2 was calculated using *Method III* as shown in Chapter 3. When the Phase I data contains only 20 samples, the average in-control run length is around 70. As the number of Phase I samples increases, the average of the true in-control run length tends towards the specified in-control ARL. As in Figure 4.1, as the number of samples increases, the variability of the true in-control ARL decreases, in turn decreasing the probability of false signals in the control chart. Boxplots of all three methods of SD estimation were very similar and have again been included in Appendix B.

4.2 Samples with Single Observations

In some situations, each Phase I sample will contain only one observation. This occurs for example in medical settings where each patient is regarded as an observation, or in production settings where there is a low number of products produced. Algorithm 1 was used to illustrate the effect of parameter estimation on both the Shewhart and CUSUM charts for varying Phase I sample sizes.

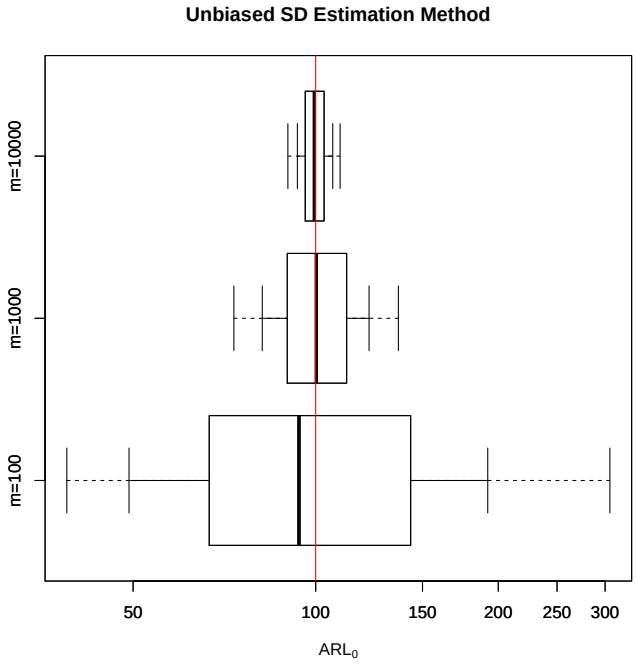

Figure 4.3: $ARL_{0,true}$ distribution of Shewhart control chart with estimated parameters for Phase I data of sample size 100, 1000 and 10000, where each sample has 1 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the x-axis.

Figure 4.3 shows the true in-control ARL for Phase I Shewhart charts with single observations. The SD was estimated using *Method III* as shown in Chapter 3. The results from using all three estimation methods gave similar results and can be seen in Appendix B. As with multiple observation data, the more samples available to estimate the chart parameters from, the better the estimation and the less variation between the true in-control ARL of the charts.

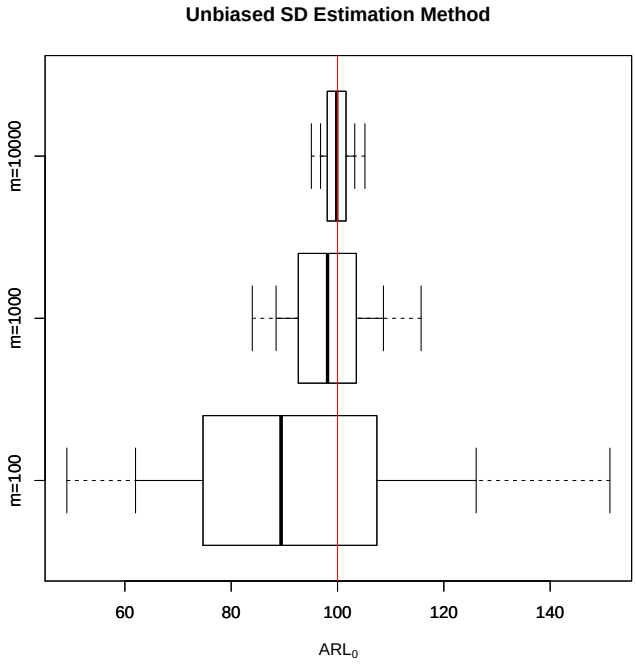

Figure 4.4: $ARL_{0,true}$ distribution of CUSUM control chart with estimated parameters for Phase I data of sample size 100, 1000 and 10000, where each sample has 1 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution.

Figure 4.4 shows the true in-control ARL for Phase I CUSUM charts with sample size of $n = 1$ where the SD was calculated using the unbiased variance method, *Method III* in Chapter 3. Appendix B contains boxplots which illustrate the effect of each SD estimation method for CUSUM charts with each sample containing a single observation. Figure 4.4 once again shows similar results to the previous figures: the greater the number of Phase I samples, the better the estimation of the underlying observation.

It can be noted that for Figures 4.1 to 4.4 that the suggested number of Phase I of 25 can lead to a true in-control ARL which is greatly different from the specified in-control ARL. This will affect the probability of false in-control and out-of-control signals occurring and the chart may not be as effective as the user believes it to be. A comparison of Figure 4.1 with 4.3 and Figure 4.2 with 4.4 shows that when the total of number of observations is the same, but sample size differs, the distribution of the actual ARL_0 does not differ greatly.

Effect of Reference Value on Estimated CUSUM Charts

When looking at the impact of estimating chart parameters, the reference value has so far been specified as $k = 0.5$. However the reference value also has an effect on the ARL of the CUSUM chart. To investigate the effect, Algorithm 1 was applied whilst varying the reference value between $k = 0.25$ and $k = 1.5$.

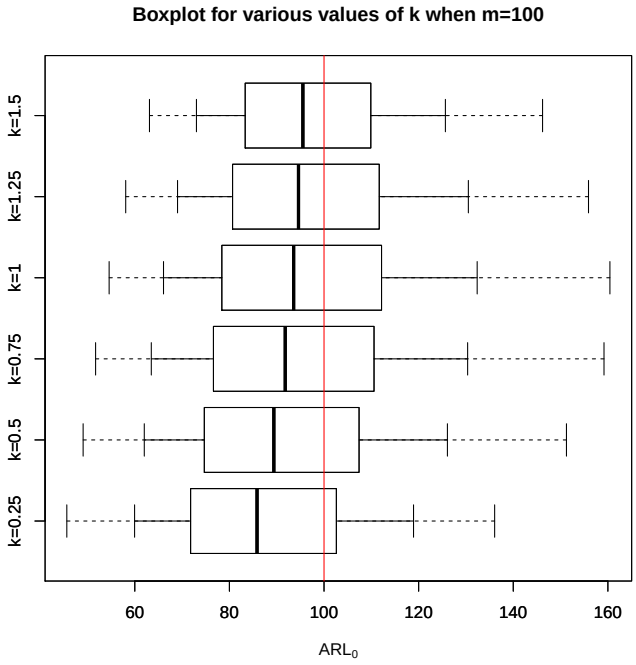

Figure 4.5: $ARL_{0,true}$ distribution of CUSUM control charts with estimated parameters for various reference values. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution.

From Figure 4.5 it can be see that as the reference value increases, the median of the true ARL_0 becomes closer to the specified ARL_0. For smaller reference values, the change in the mean wishing to be detected is small and so the errors present in the estimation of the parameters can increase the probability of a false out-of-control signal.

5. Accounting for Estimation Error

Although many authors have demonstrated the impact of estimation error on control charts, few have investigated how to account for the estimation errors (Albers and Kallenberg, 2005). Examples of authors who discuss accounting for the estimation error are Albers and Kallenberg (2005) and Gandy and Kvaløy (2013).

5.1 Corrections for Shewhart Charts

Albers and Kallenberg have produced various papers on how to account for parameter estimation error in Shewhart control charts. They have developed two main approaches for this. The first aims to reduce ARL bias. The second is to limit the exceedance probability of the chosen performance measure. For each approach an adjustment factor is included in the control limit equations. These corrections work well for Shewhart charts and can be implemented easily without needing to understand exactly how they work. However these corrections are limited by the fact that they cannot be applied to other types of control charts.

5.1.1 Reduction of ARL Bias

Looking at Figure 4.3, it can be seen that the median of the true ARL_0 distribution is roughly equal to the specified ARL_0. However when calculated, the mean of the distribution is slightly larger than the specified ARL_0 at 113. The reduction of ARL bias method, aims to correct the distribution such that the mean, or expected value of the true ARL_0, is approximately equal to the specified ARL_0. Albers and Kallenberg (2005) state that to reduce the ARL bias when the parameters are estimated, the control limits should be corrected to

$$\begin{aligned} UCL &= \hat{\mu}_0 + L\frac{\hat{\sigma}_0}{\sqrt{n}}\left(1 + \frac{B}{m}\right) \\ LCL &= \hat{\mu}_0 - L\frac{\hat{\sigma}_0}{\sqrt{n}}\left(1 + \frac{B}{m}\right) \end{aligned} \tag{5.1}$$

where $B = \frac{1}{2}\left(1 - L^2\tau^2\right)$ and τ depends on the method used to estimate the SD and the number of observations in each sample. It is worth noting in Equation 5.1 that as the number of Phase I samples, m, increases, the correction factor added to the limits decreases. As discussed in Chapter 4, the greater the number of observations used to estimate the parameters, the less variation is present in the estimations.

For each SD method as described in Chapter 3, Albers and Kallenberg (2005) have specified how to calculate τ^2. Each SD estimate can be expressed as $\hat{\sigma} = \frac{\sigma^*}{A}$, where σ^* is a biased SD estimate and A is a constant which reduces bias from the SD estimation. From this τ^2 can be calculated as

$$\tau^2 = \lim_{m\to\infty}\left\{m \cdot \mathrm{Var}\left(\frac{\sigma^*}{E(\sigma^*)}\right)\right\}$$

A brief outline of how to find the control limits when *Method III* has been used to estimate the SD follows. To see the calculation of control limits for the other SD estimation methods see Appendix C. For a more detailed description of how to find the control limits see Albers and Kallenberg (2005).

Method III

For this SD estimation method where the sample size is n=1, $\tau^2 = \frac{1}{2}$. From this B is calculated as

$$\begin{aligned} B &= \frac{1}{2}\left(1 - \frac{L^2}{2}\right) \\ &= \frac{1}{2} - \frac{L^2}{4} \\ &= \frac{2 - L^2}{4} \end{aligned}$$

This gives the control limits with correction factor for $n = 1$ as

$$\begin{aligned} UCL &= \hat{\mu}_0 + L\hat{\sigma}_0 \left\{1 + \frac{2 - L^2}{4m}\right\} \\ LCL &= \hat{\mu}_0 - L\hat{\sigma}_0 \left\{1 + \frac{2 - L^2}{4m}\right\} \end{aligned} \tag{5.2}$$

In these control limit equations if $L^2 > 2$, or $L > \sqrt{2}$, then the bias reduction limits will be reduced compared to the unadjusted limits.

For $n > 1$, $\tau^2 = n\left(c_4^{-2}(n) - 1\right)$, which gives

$$B = \frac{1}{2}\left(1 - n(c_4^{-2}(n) - 1)L^2\right)$$

From this the control limits can then be expressed as

$$\begin{aligned} UCL &= \hat{\mu}_0 + L\hat{\sigma}_0 \left\{1 + \frac{1 - n\left(c_4^{-2}(n) - 1\right)L^2}{2m}\right\} \\ LCL &= \hat{\mu}_0 - L\hat{\sigma}_0 \left\{1 + \frac{1 - n\left(c_4^{-2}(n) - 1\right)L^2}{2m}\right\} \end{aligned}$$

5.1.2 Reduction of Exceedance Probability

When the user wishes to reduce the probability that the true ARL_0 will exceed a certain percentage of the specified in-control ARL, $ARL_{0,s}$, the following control limits can be used

$$\begin{aligned} UCL &= \hat{\mu}_0 + L\hat{\sigma}_0 \{1 + E\} \\ LCL &= \hat{\mu}_0 - L\hat{\sigma}_0 \{1 + E\} \end{aligned}$$

where

$$E = \frac{\Phi(\alpha)\tau}{m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1 - \epsilon)}$$

and α and ϵ are specified by the user such that the true ARL_0 is not below $(1 - \epsilon)ARL_{0,s}$ in more than $100\alpha\%$ of the applications. For example if the specified $ARL_0 = 500$ and the user

wishes for the true ARL_0 to not be less than 450 in at most 20% of the applications, then $\alpha = \frac{20}{100} = 0.2$ and $\epsilon = 1 - \frac{450}{500} = 0.1$. As with the bias reduction method, τ depends on which SD estimation is used. The derivation of the control limits for the SD estimation *Method III* from Chapter 3 is shown below. It should be noted that with $\epsilon = 0$ the adjusted control limits will always be increased compared to the unadjusted control limits. For other SD estimation methods, see Appendix C.

Method III

When the SD has been estimated using method III and $n = 1$, $\tau = \frac{1}{4}$. The adjusted control limits are therefore given by

$$UCL = \hat{\mu}_0 + L\hat{\sigma}_0 \left\{ 1 + \frac{\Phi(\alpha)}{4m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$
$$LCL = \hat{\mu}_0 - L\hat{\sigma}_0 \left\{ 1 + \frac{\Phi(\alpha)}{4m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$

For samples where $n > 1$, then $\tau = \left[n \left(c_4^{-2}(n) - 1 \right) \right]^{\frac{1}{2}}$ and the control limits are

$$UCL = \hat{\mu}_0 + L\frac{\hat{\sigma}_0}{\sqrt{n}} \left\{ 1 + \frac{\Phi(\alpha) \left[n \left(c_4^{-2}(n) - 1 \right) \right]^{\frac{1}{2}}}{m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$
$$LCL = \hat{\mu}_0 - L\frac{\hat{\sigma}_0}{\sqrt{n}} \left\{ 1 + \frac{\Phi(\alpha) \left[n \left(c_4^{-2}(n) - 1 \right) \right]^{\frac{1}{2}}}{m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$

5.2 Bootstrap Method for Guaranteed Conditional Performance

Gandy and Kvaløy (2013) have developed a method which is not limited to Shewhart charts, but instead can be used on many others "without having to derive specific approximation formulas in each setting". The method produces control limits which are constructed such that for approximately $100(1-\alpha)\%$ of the applications, the true ARL_0 will be at least the specified ARL_0.

For a one-sided control chart without any adjustments, the control limit is calculated as $c(\hat{P}; \hat{\xi})$ for a specified in-control ARL, where c is the control limit, which is dependent on the estimated distribution, $\hat{P}$, and parameters, $\hat{\xi}$. However as shown in Chapter 4 this does not always produce a true in-control ARL which is close to the specified ARL_0. If the true distribution of the process, P, is known then the control limit $c(P; \hat{\xi})$ can be found, which gives the correct in-control ARL. The distribution $c(\hat{P}; \hat{\xi}) - c(P; \hat{\xi})$ can be calculated and from this, a constant p_α can be found such that

$$P\left(c(\hat{P}; \hat{\xi}) - c(P; \hat{\xi}) > p_\alpha \right) = 1 - \alpha$$

This can be rearranged to give

$$P\left(c(P; \hat{\xi}) < c(\hat{P}; \hat{\xi}) - p_\alpha \right) = 1 - \alpha$$

If such a p_α exists then $\left(-\infty, c(\hat{P};\hat{\xi}) - p_\alpha\right)$ would be considered as an approximate lower one-sided $(1-\alpha)100\%$ confidence interval for $c(P;\hat{\xi})$ for a one-sided chart. However as the true distribution is unknown, p_α is also unknown, but can be approximated by creating bootstrap samples from the estimated distribution.

Each bootstrap sample has the same dimension as the Phase I data and follows the estimated distribution $\hat{P}$. For each bootstrap sample the distribution, $\hat{P}^*$, and parameters, $\hat{\xi}^*$, can be estimated. The distribution $c(\hat{P}^*;\hat{\xi}^*) - c(\hat{P};\hat{\xi}^*)$ can thus be simulated, and a constant p^*_α can be found such that

$$P\left(c(\hat{P}^*;\hat{\xi}^*) - c(\hat{P};\hat{\xi}^*) > p^*_\alpha | \hat{P}\right) = 1-\alpha$$

The constant p^*_α is approximately equal to p_α. Thus $\left(-\infty, c(\hat{P};\hat{\xi}) - p^*_\alpha\right)$ is an approximate one-sided $1-\alpha$ confidence interval that $c(P;\hat{\xi})$, for a one-sided chart, is below $c(\hat{P};\hat{\xi}) - p^*_\alpha$. Thus in practise the adjusted control limit is set as $c\left(\hat{P},\hat{\xi}\right) - p^*_\alpha$ for a one-sided control chart. In approximately $(1-\alpha)100\%$ of the applications, the true ARL_0 will then be at least as large as the specified ARL_0. The adjustment factor is thus said to produce guaranteed conditional performance.

To find the adjusted control limits for a two-sided control chart, α should be replaced with $\frac{\alpha}{2}$ in the above equations to give the UCL and $1-\frac{\alpha}{2}$ for the LCL. For the UCL this gives that $\left(-\infty, c_U(\hat{P}^*;\hat{\xi}^*) - p^*_{\frac{\alpha}{2}}\right)$ is an approximate confidence interval that $c_U(P;\hat{\xi})$, for a two-sided chart, is below $c_U(\hat{P};\hat{\xi}) - p^*_{\frac{\alpha}{2}}$. For the LCL $\left(-\infty, c_L(\hat{P}^*;\hat{\xi}^*) - p^*_{1-\frac{\alpha}{2}}\right)$ is an approximate confidence interval that $c_L(P;\hat{\xi})$, for a two-sided chart, is above $c_L(\hat{P};\hat{\xi}) - p^*_{1-\frac{\alpha}{2}}$. Thus in practise the control limits for a two-sided control chart are

$$UCL = c_U\left(\hat{P},\hat{\xi}\right) - p^*_{\frac{\alpha}{2}}$$

$$LCL = c_L\left(\hat{P},\hat{\xi}\right) - p^*_{1-\frac{\alpha}{2}}$$

5.3 Comparison of Methods which Account for Estimation Error

For a quick comparison of each method to reduce estimation error, 100 samples, each of size 1, from a standard normal distribution were used to produce a Shewhart chart with a specified in-control ARL of 370. The R package qcc (Scrucca, 2004) was used to produce the chart. *Method III* was used to estimate the SD. The ARL bias reducing, exceedance probability and guaranteed condition performance control limits for the same samples were then calculated and added to the control chart as shown in Figure 5.1.

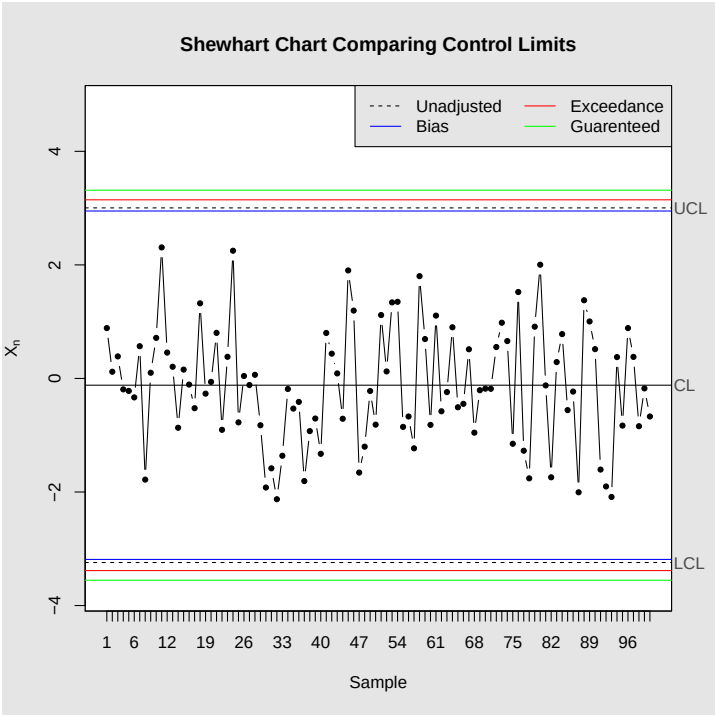

Figure 5.1: A Shewhart Control Chart with specified in-control ARL of 370, demonstrating the change of the limits whilst employing each method of reducing estimation error.

For each method of calculating the control limits the true ARL_0 was calculated using Algorithm 1. For the standard Shewhart control limits the $ARL_0 = 519$. For the reduction of ARL bias the $ARL_0 = 433$. For the exceedance probability with $\alpha = 0.2$ and $\epsilon = 0.1$ the $ARL_0 = 845$, and for the guaranteed conditional performance method the $ARL_0 = 1546$. Although these in-control ARLs are larger than the specified 370, this means that there will be fewer false out-of-control signals. To compute the critical value for the guaranteed performance method the spcadjust R package was used (Gandy and Kvaløy, 2013).

Algorithm 1 can be used to compare the true ARL_0 distribution for each method of accounting for estimation error. Samples of size 100, each with one observation were generated from the standard normal distribution with a specified $ARL_0 = 100$ for each iteration of the algorithm. The algorithm was then iterated 1000 times to show how each method affects the ARL_0 distribution.

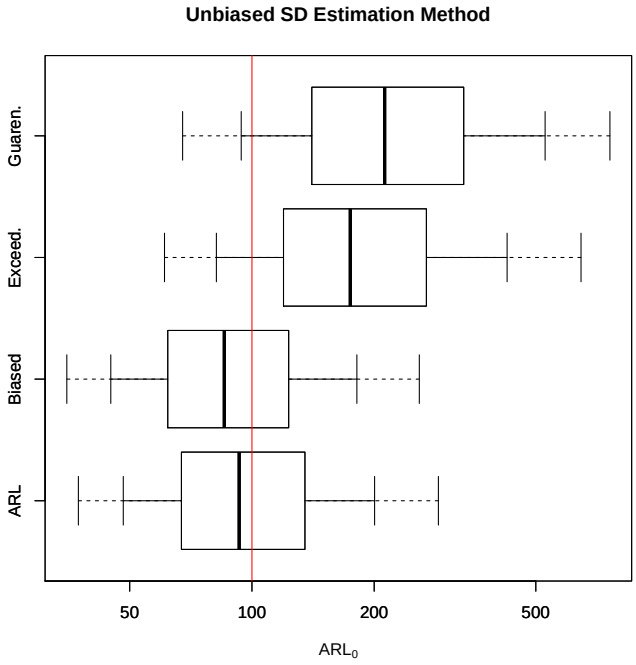

Figure 5.2: The true in-control ARL distribution of Shewhart Charts for various control limits. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the horizontal axis.

Figure 5.2 shows the distribution of the true in-control ARL of Shewhart charts for each method of calculating the control limits. For the exceedance case $\alpha = 0.1$ and $\epsilon = 0.1$ and for the Bootstrap method $\alpha = 0.1$. It can be seen that the variance in the true in-control ARL distributions does not reduce with each method. In the ARL biased reduction method, the true ARL_0 distribution has shifted to the left as expected. The mean of the unadjusted ARL_0 distribution is 111 and the mean of the ARL_0 distribution for the biased reduction adjustment method is 101. This shows that although the method produces a distribution with a mean which is closer to the specified ARL_0, it has a higher chance of producing a true ARL_0 which is less than the specified ARL_0.

For the exceedance and guaranteed conditional performance, the distributions of the true in-control ARL both have a large number of run lengths greater than the specified in-control ARL. These methods are favoured as the number of false out-of-control signals is greatly reduced. This does lead to a longer true out-of-control run length, but the out-of-control ARL is not considerably affected by the increase to the in-control ARL.

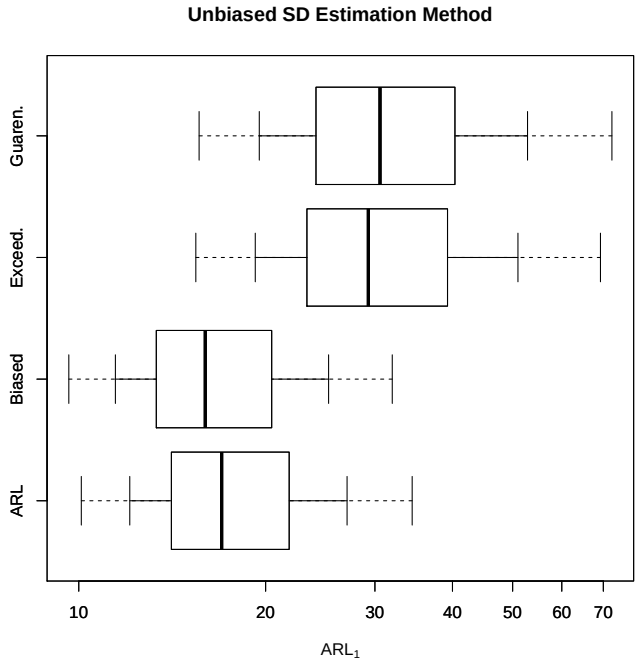

Figure 5.3: The true out-of-control ARL distribution of Shewhart Charts for various control limits when the out-of-control mean to be detected is $\mu_1 = \mu_0 + \sigma_0$. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{1,true}$ distribution. A log-scale has been applied to the horizontal axis.

Figure 5.3 shows the distribution of the out-of-control ARL for each method of calculating the control limits. Using the method described in Chapter 2 for a control chart with known normal distribution, which is used to detect a shift in mean of $\mu_1 = \mu_0 + \sigma_0$, the out-of-control ARL is 17.33. For the unadjusted control limits and the bias adjusted control limits the median of the ARL_1 distribution is close to that of the expected out-of-control ARL when the parameters are known. For the exceedance probability and bootstrap guaranteed adjusted control limits, the median of the ARL_1 distribution is close to double the ARL_1 when the underlying normal distribution is known. It would take around 30 samples from when the shift in the underlying mean occurs, to detect the change when using either of these two methods. This is not a huge difference from the unadjusted limits and has the added bonus of less false out-of-control signals and so will not greatly effect the control chart's ability to detect significant changes to the underlying distribution.

For the exceedance probability reduction method, if $\epsilon = 0$ and $\alpha = 0.1$, then the method aims to give a true in-control run length which does not fall below the specified ARL_0 for at least 90% of the applications. This is also the specification for the guaranteed performance method and so can be directly compared when the exceedance probability is used with the specified parameters.

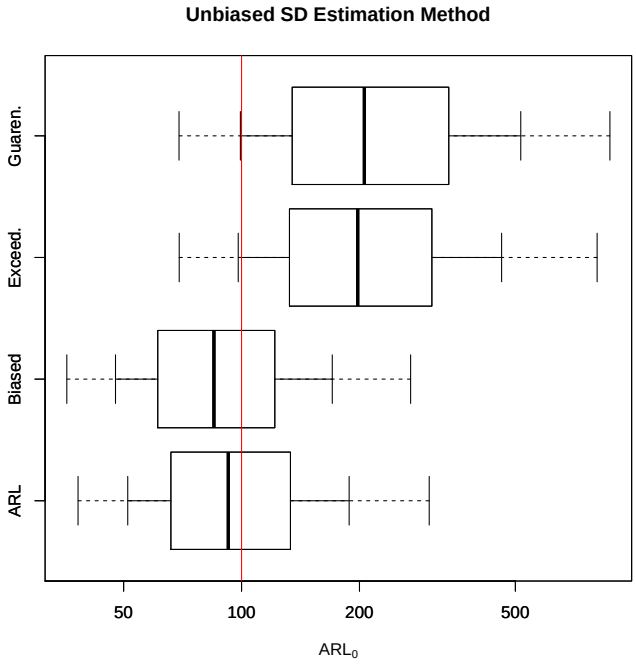

Figure 5.4: The actual in-control ARL distribution of Shewhart Charts for various control limits. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the horizontal axis.

From Figure 5.4 it can be seen that when $\epsilon = 0$ and $\alpha = 0.1$ for the exceedance probability method this gives approximately the same result as the bootstrap guaranteed performance method with $\alpha = 0.1$. However the exceedance probability method has the disadvantage that it can only be applied to Shewhart charts whereas the bootstrap guaranteed performance method can be applied to other types of charts, as well as risk-adjusted charts.

The effect of the bootstrap guaranteed performance method on the in-control ARL distribution can be shown for risk-adjusted CUSUM charts. When the response variable, Y_i is binary, a logistic regression model must be fitted as described in Section 2.4. Each observation is dependent on one variable X, which is between 0 and 70 and thus is set to produce an example of how the Parsonnet score can effect the recovery of a patient who has undergone surgery.

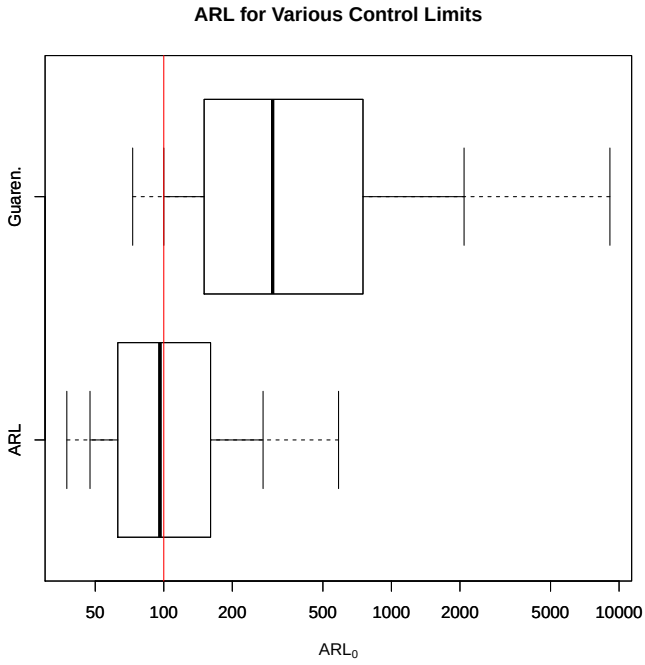

Figure 5.5: The actual in-control ARL distribution of risk-adjusted CUSUM charts for various control limits. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the horizontal axis.

Figure 5.5 shows the distribution of the true in-control run lengths for risk-adjusted CUSUM charts, where a logistic regression model has been fitted to the response variable. It can be seen that when the chart limits have not been adjusted, the median of the ARL_0 distribution is close to 100, and so on average, roughly half the applications of the unadjusted control limits will have an in-control ARL less than the specified value. For the bootstrap guaranteed performance method, as expected, on average approximately 10% of the applications will have an in-control ARL which is less than the specified 100. Figure 5.5 shows that for charts which have bootstrap adjusted control limits, on average only 2.5% of the applications will have an in-control ARL of less than 75. Thus for the 10% of applications where the ARL_0 is less than 100, the true ARL_0 for the majority of these applications will still be close to the specified ARL_0.

6. Cardiac Surgery Analysis

The use of the risk-adjusted CUSUM chart is extremely useful in the monitoring of medical data. To show an example of where an increase in the probability of an event occurring is to be detected, the surgical data as seen in Steiner et al. (2000) can be used. The data contains the surgical performance for nearly 7000 patients who all underwent a cardiac operation in the same surgical centre situated in the UK, from 1st January 1992 to 31st December 1998. The surgical performance is monitored using the 30 day mortality of each patient. For this situation, $Y_i = 1$ if the event occurs, i.e. if the patient does not survive for 30 days after the surgery.

The first two years of the data is treated as the Phase I period. The adjusted control limits are calculated using the Phase I data and the spcadjust R package (Gandy and Kvaløy, 2013). For the spcadjust package to calculate the adjusted control limits a value of Δ must be specified. The value of Δ is found such that it corresponds to the change in the probability of death the user wishes to detect, as shown in Equation 2.16. The probability of death in the Phase I samples is first calculated. During Phase I 2218 patients underwent surgery with 413 patients who did not survive at least 30 days post surgery. This gives a mortality probability of $p = 0.0644$. If we wish to detect that the probability of mortality doubles, we must use that $logit(p_i) = X_i\beta$, as shown in the logistic regression section of Chapter 2.4, as each patient has a binary response variable Y_i. From this the sort of "average" in-control risk score is

$$\overline{X\beta} = ln\left(\frac{p}{1-p}\right)$$

Substituting in $p = 0.064$ gives the average in-control value of

$$\overline{X\beta} = ln\left(\frac{0.0644}{1-0.0644}\right) = -2.67$$

To detect a change such that the probability of mortality doubles then the average out-of-control risk score can be calculated as

$$\begin{aligned}\overline{X\beta} + \Delta &= ln\left(\frac{2p}{1-2p}\right)\\ &= ln\left(\frac{0.129}{1-0.129}\right)\\ &= -1.91\end{aligned}$$

Taking the average in-control risk score from the average out-of-control risk score will then give the value of Δ for the risk-adjusted CUSUM chart used to detect a doubling of the 30 day mortality probability.

$$\Delta = -1.91 - (-2.67) = 0.76 \tag{6.1}$$

A target of the in-control ARL must also be specified for the spcadjust R package to calculate the adjusted control limit. For this example the target in-control ARL was specified as 10000.

Steiner et al. (2000) states that using backwards elimination, only the Parsonnet score of each patient was included as an explanatory variable in the logistic regression model. This gave the regression model of

$$logit(p_i) = -3.68 + 0.077X_i$$

where X_i is the Parsonnet score for patient i. Using this logistic regression model and the spcadjust R package, the Phase I and Phase II risk-adjusted CUSUM charts are produced for the cardiac surgery data.

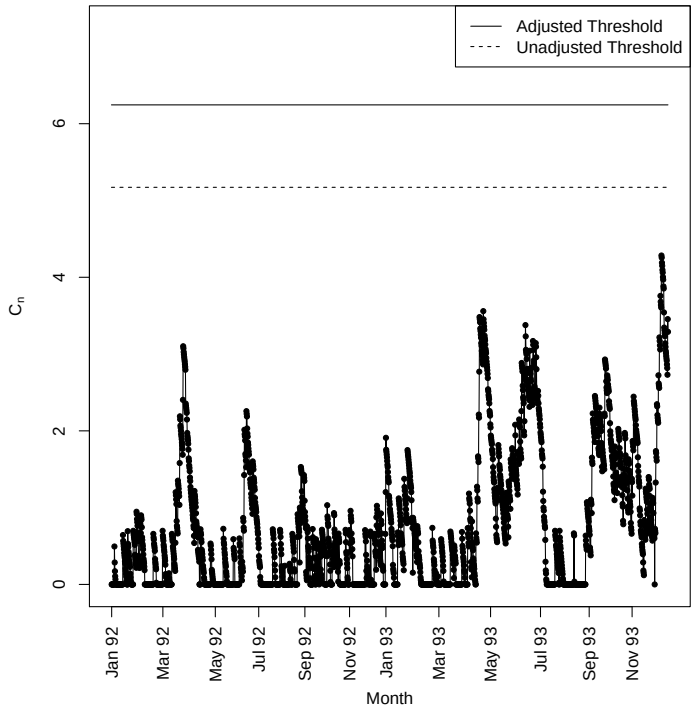

Figure 6.1: Risk-adjusted CUSUM chart showing Phase I data for 30 day mortality of cardiac surgery patients. Adjusted control limit is 6.2, unadjusted control limit is 5.2.

Figure 6.1 shows the risk-adjusted CUSUM chart for the Phase I data. From the chart it can be seen that for both the adjusted and unadjusted control limits, the chart is in-control. These limits can now be used to produce the chart containing the Phase II data.

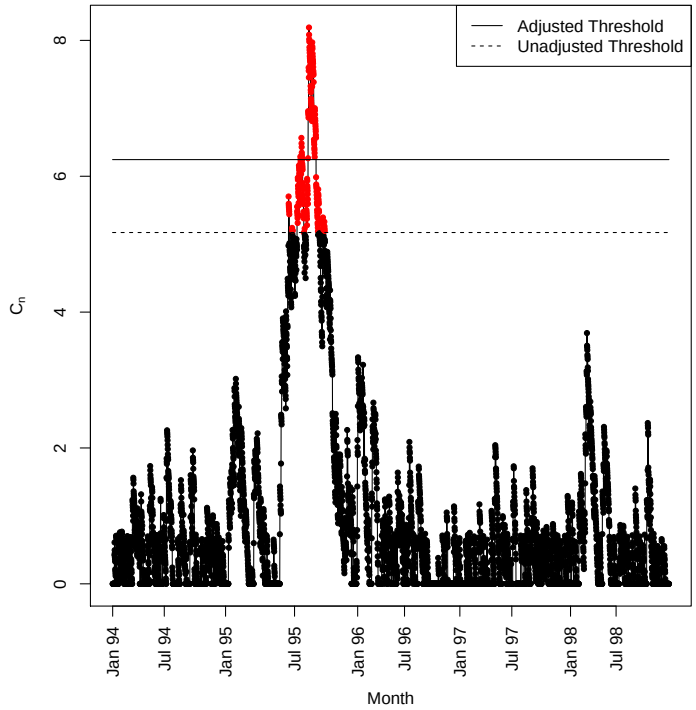

Figure 6.2: Risk-adjusted CUSUM chart showing Phase II data for 30 day mortality of cardiac surgery patients. Adjusted control limit is 6.2, unadjusted control limit is 5.2.

Figure 6.2 show the risk-adjusted CUSUM chart for the Phase II cardiac surgery data. It can be seen that just before July 1995 there is a steep increase in the number of 30 day mortalities occurring for this surgical centre. Between July 1995 and January 1996 the chart signals out-of-control and so it should be investigated as to why there is such an increase in 30 day mortality.

7. Helping Babies Breath Analysis

The risk-adjusted CUSUM chart can be used to monitor a change in the probability of a medical data set. This process is applied to the Helping Babies Breath (HBB) data which was presented in Mduma et al. (2015). The data contains information on infants born after the introduction of the HBB program in Haydom Lutheran Hospital in Northern Tanzania. In February 2010 the National HBB study started which included a one-day course where the care providers were assessed before and after the course on their basic neonatal resuscitation management. Mduma et al. (2015) states that the pass rate for newborn resuscitation in a test set up increased after completion on the day of the course compared to the assessment carried out prior to the course. However this did not translate into a decrease in infant mortality for observations after the course.

In February 2011 Frequent and Brief On-Site (FBOS) HBB simulation training was implemented, which included providing equipment to the hospital such that it was easily available for frequent practice. Monthly training sessions were also carried out within the hospital. The data analysed in Mduma et al. (2015) contains information on all infants born in Haydom Lutheran Hospital from February 2010 until January 2012. The neonatal outcome in the first year, when the one-day HBB course was completed, with the neonatal outcome of the second year, when the FBOS training was started is compared. An updated version of this data is used to illustrate the use of the risk-adjusted CUSUM chart when monitoring for a change in probability, where the last observations were the infants born on 31$^{\text{st}}$ January 2016.

The data collected from the introduction of the one-day HBB course on 1$^{\text{st}}$ February 2010 up until the 31$^{\text{st}}$ January 2011, the day before the introduction of the FBOS training, will be treated as the Phase I data. The Phase II data is from 1$^{\text{st}}$ February 2011 until 31$^{\text{st}}$ January 2016. This will aid in comparing the effectiveness of the FBOS training to the one-day HBB course and to see if there is a noticeable decrease in the number of infant mortalities since the introduction of the FBOS training. The Phase I data contains 4844 observations and the Phase II data 22176 observations. The data contains many variables including the outcome after 24 hours for each infant. This is the binary response variable Y_i, where $Y_i = 1$ if the infant did not survive 24 hours or was classified as a fresh stillbirth.

7.1 Initial Analysis

The first stage in analysing the HBB data is to produce a basic CUSUM chart for Phase I with both unadjusted and adjusted control limits. The adjusted control limits will be calculated using the bootstrap method as this is the method which can be applied to CUSUM charts. To be able to construct the CUSUM chart, the change in the probability to be detected must be specified. Although the data is not to be risk-adjusted in this initial analysis, the logistic regression set up from Section 2.4 is used, including only a constant term such that $\text{logit}\,(p_i) = \beta_1$. This set-up is needed as the response variable is binary.

Determining Delta

To monitor if a "clinically relevant" change in the probability of the data occurs, a value of Δ must be found for Equation 2.16 in Section 2.4. The HBB data has a binary response, with $Y_i = 1$ if the event occurred, i.e. the infant does not survive the first 24 hours. During the Phase I data there were 133 deaths and 4844 samples, giving a probability of death as $p = 0.0275$. This gives an in-control average risk score of

$$\hat{\beta} = ln\left(\frac{0.0275}{1 - 0.0275}\right) = -3.56$$

It was decided that a 20% decrease in the probability of death would be clinically relevant, such that a probability of $p = 0.022$ was to be detected. The resulting out-of-control constant term is thus given by

$$\hat{\beta} + \Delta = ln\left(\frac{0.022}{1 - 0.022}\right) = -3.79$$

From this Δ can be calculated as

$$\Delta = -3.79 - (-3.56) = -0.23$$

where $\Delta < 0$ gives a CUSUM chart which gives an out-of-control signal when the probability of death decreases.

Phase I

Using the value of $\Delta = -0.23$, the control limits for the basic CUSUM chart were calculated using the spcadjust R package (Gandy and Kvaløy, 2013) to be 4.5 for the unadjusted control limit, unfortunately the current set up the spcadjust package cannot determine the adjusted control limit for a regression model which is only dependent on a constant term. The unadjusted control limit is then used with a specified $ARL_0 = 40000$ to produce a CUSUM chart for the Phase I data.

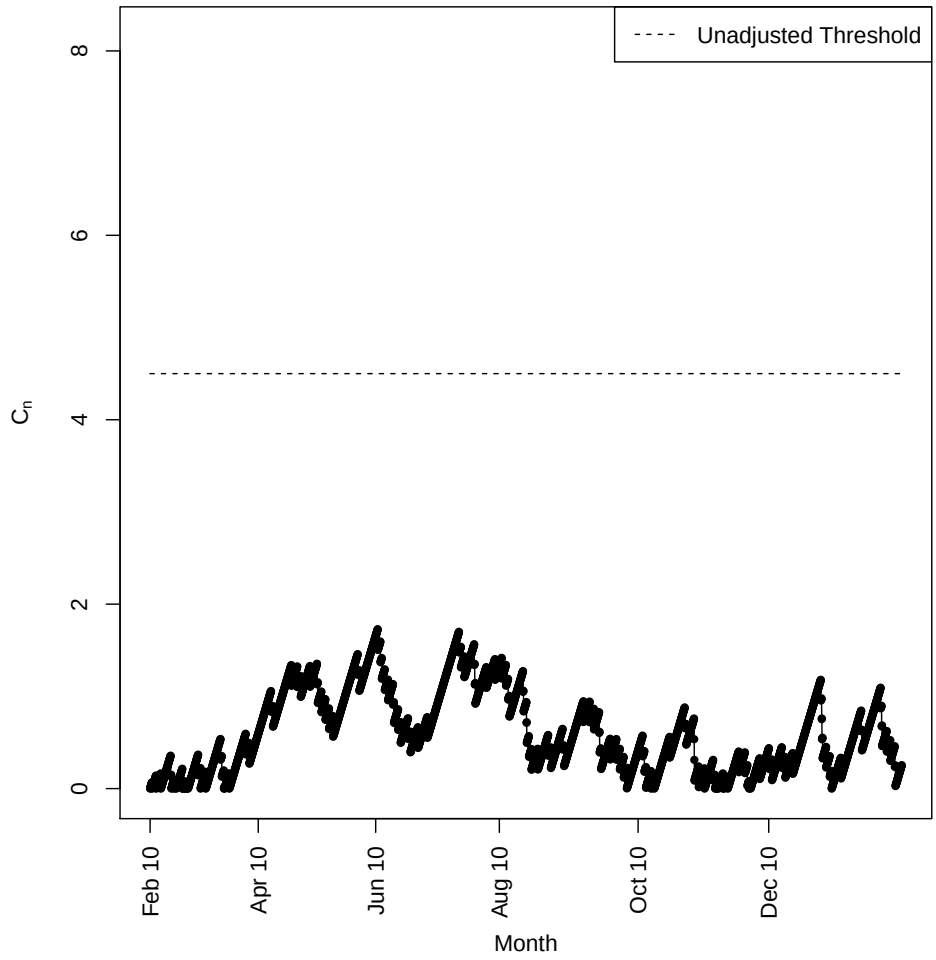

Figure 7.1: Phase I basic CUSUM chart for detecting a decrease in the probability of infant mortality of the HBB data. The unadjusted control limit is 4.5.

From Figure 7.1 it can be seen that all Phase I observations are in control for the control limit, and thus this control limit can be used to monitor the Phase II observations.

Phase II

The control limits found using the Phase I samples can be used to construct the Phase II CUSUM chart, with the parameters of Δ and ARL_0 as specified for the Phase I chart.

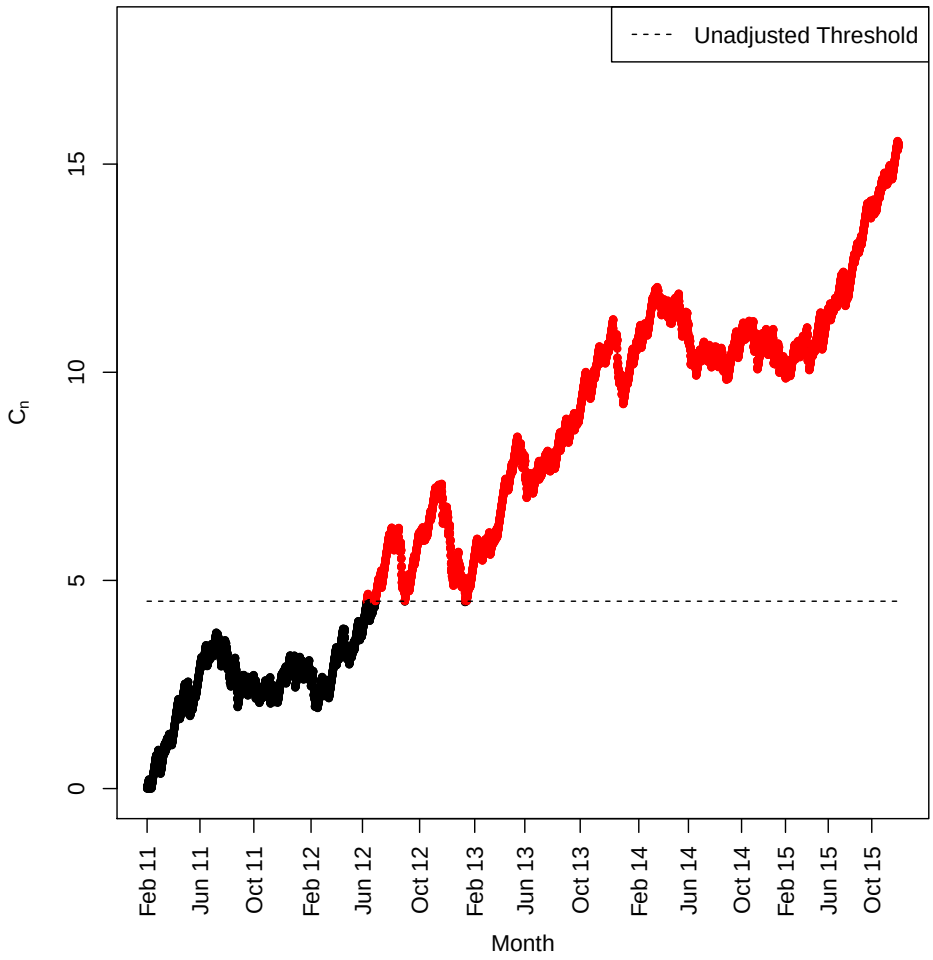

Figure 7.2: Phase II basic CUSUM chart for detecting a decrease in the probability of infant mortality of the HBB data. The unadjusted control limit is 4.5.

Figure 7.2 shows the basic CUSUM chart of the Phase II HBB data. It can be seen that the chart first gives an out-of-control signal around July 2012. It should be noted that the chart is constructed using only the unadjusted control limit and that the chart will signal out-of-control earlier than if an adjusted control limit was used. However the chart remains out-of-control and from March 2013 onwards the chart shows a steady increase in the value of C_n. This suggests that the FBOS HBB training program has a positive effect on the probability of an infant surviving for the first 24 hours, but it is not until September 2012 for this effect to be formally established as clinically relevant.

Variable Life Adjusted Display

The CUSUM chart is a good visual indicator with formal control limits for the effect of the HBB training program. However when showing the CUSUM chart to those involved in the HBB program, it may be hard to interpret the y-axis. To help with this interpretation a Variable Life Adjusted Display (VLAD) can also be constructed. A VLAD is a plot which shows the cumulative sum of the expected number of outcomes minus the observed number of outcomes against time Lovegrove et al. (1997).

For the HBB data the expected number of outcomes, i.e. death of the infant, for each sample is the probability of death during Phase I, which is 0.0275. The observed number of outcomes

is the binary response variable Y_i of the Phase II samples. The equation used to calculate the VLAD at each observation j is

$$\begin{aligned} V_j &= \sum_{i=1}^{j} (Y_{i,E} - Y_{i,O}) \\ &= \sum_{i=1}^{j} (0.0275 - Y_{i,O}) \end{aligned} \tag{7.1}$$

where V_i is the cumulative number of excess survivors, $Y_{i,E}$ is the expected number of outcomes for observation i and $Y_{i,O}$ is the observed number of outcomes for observation i. Thus if V_i is greater than 0, this means that the expected number of outcomes is greater than the observed and thus there has been a decrease in deaths during the Phase II observations compared to the baseline Phase I samples. If V_i is less than 0, then the number of observed outcomes is greater than the expected number of outcomes and indicates an increase in deaths during Phase II compared to Phase I.

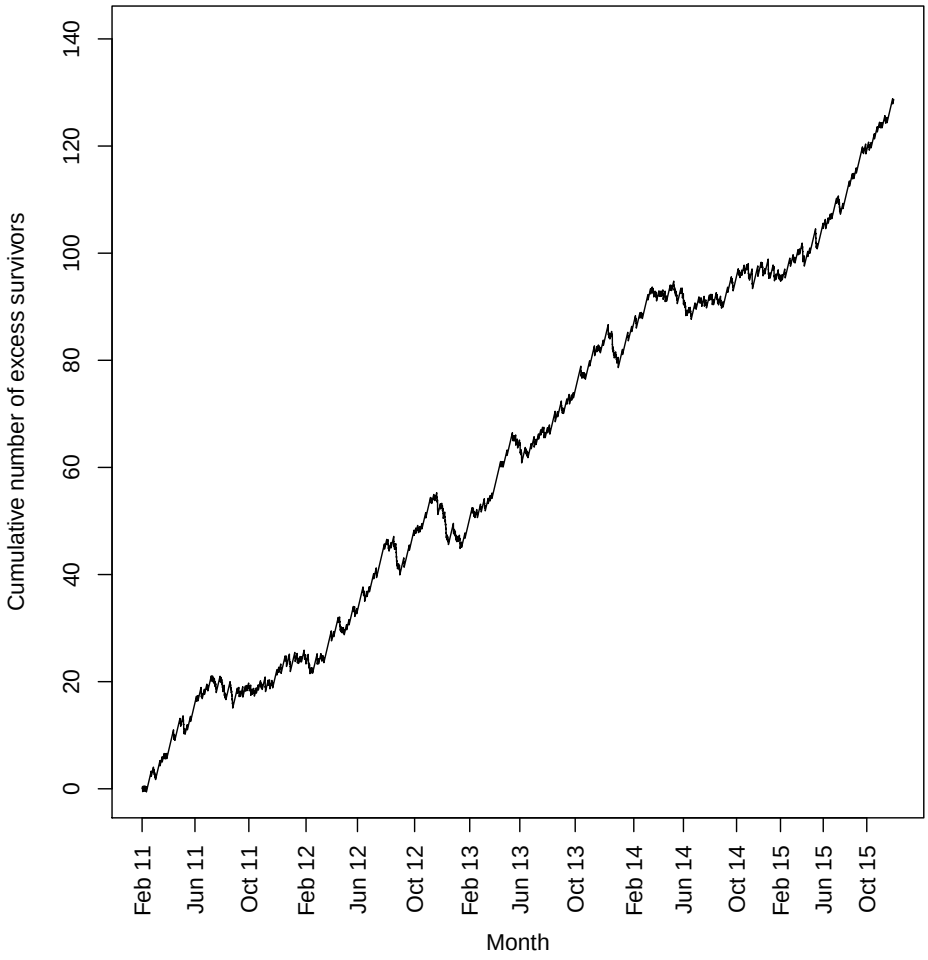

Figure 7.3: VLAD plot for Phase II HBB Data with expected value of 0.0275.

From Figure 7.3 it can be seen that from February 2011 to January 2016 there has been a great increase in the number of infants surviving, compared to expected value of Phase I data. The shape of the VLAD is similar to the shape of the basic CUSUM shown in Figure

7.2, however the y-axis of the VLAD corresponds to the cumulative number of infants surviving compared to the baseline level. In this case it can be seen from this plot that there were around 130 lives saved over the Phase II period compared to the baseline level.

7.2 Risk-Adjusted CUSUM Chart

To produce a risk-adjusted CUSUM chart of the HBB data a suitable regression model must first be fitted to the data. It is reasonable to fit a regression model to the data as there are many factors which could effect the outcome for each infant. These factors need to be taken into consideration, as the difficulty of each case will effect the outcome. As with the cardiac data in Chapter 6, the response variable is binary and so a logistic regression must be fitted to the data. The effectiveness of the training is to be analysed and so any "management" variables are not included. The regression model should only contain variables which are related to the health of the child and the mother. To find the regression model for the HBB data, the stats R package was used (R Core Team, 2015).

Fitting an Appropriate Regression Model

Regression models for each individual variable are first produced. For each of the individual regression models, if the variable was significant at the 20% level it was included in the initial multiple regression model. The initial multiple model contains the 15 variables as shown in Table D.1 in Appendix D. To find an appropriate multiple model, the stepwise method was used to remove any variables which were not significant from the initial multiple model. This gave the multiple regression model as

$$Y_i = 4.45 - 0.0014X_{i1} - 2.43X_{i2} + 3.58X_{i3} + 7.34X_{i4} + 0.18X_{i5} + 4.44X_{i6} + 1.29X_{i7} \quad (7.2)$$

where the covariates are explained in Table 7.1.

Table 7.1: Covariates Present in Equation 7.2

Covariate	Description
X_{i1}	Birth weight, grams
X_{i2}	If the mother suffered from pre-eclampsia
X_{i3}	Abnormal foetal heart rate
X_{i4}	Foetal heart rate not measured/not detected
X_{i5}	Caesarean section delivery
X_{i6}	Assisted breech delivery
X_{i7}	Vacuum extraction

However when this model is fitted to the Phase I data it has 945 missing observations of 4844. Therefore the model expressed in Equation 7.2 is fitted on only 80% of the data. When looking at the data most of these missing observations are present in the variable delivery mode, which has 920 of 4844 observations missing. It was decided to exclude any variable with large missing observations in either Phase I or Phase II. Of the 15 variables shown in Table D.1, admission and delivery mode had around 20% of observations missing in Phase I and prolonged labour has over 50% of the observations missing in Phase II. These variables were removed from the initial multiple model, before finding the final model using the step-wise method. The resulting regression model is

$$Y_i = 2.94 - 0.0011X_{i1} - 2.05X_{i2} + 2.57X_{i3} + 7.20X_{i4} + 0.43X_{i5} - 1.06X_{i6} + 2.03X_{i7} + 0.37X_{i8} \quad (7.3)$$

where the covariates of the model are explained in Table 7.2

Table 7.2: Covariates Present in Equation 7.3

Covariate	Description
X_{i1}	Birth weight, grams
X_{i2}	If the mother suffered from pre-eclampsia
X_{i3}	Abnormal foetal heart rate
X_{i4}	Foetal heart rate not measured/notdetected
X_{i5}	Presentation of baby was breeched
X_{i6}	Presentation of the baby was shoulder dystocia
X_{i7}	Presentation of baby was tranversed
X_{i8}	Other presentation of the baby

Phase I

Using $\Delta = -0.23$ and the spcadjust R package (Gandy and Kvaløy, 2013), the adjusted and unadjusted control limits for the risk-adjusted CUSUM chart are 4.0 and 3.4 respectively. The Phase I control chart is constructed using these control limits and the regression equation as shown in Equation 7.3, with a specified $ARL_0 = 40000$.

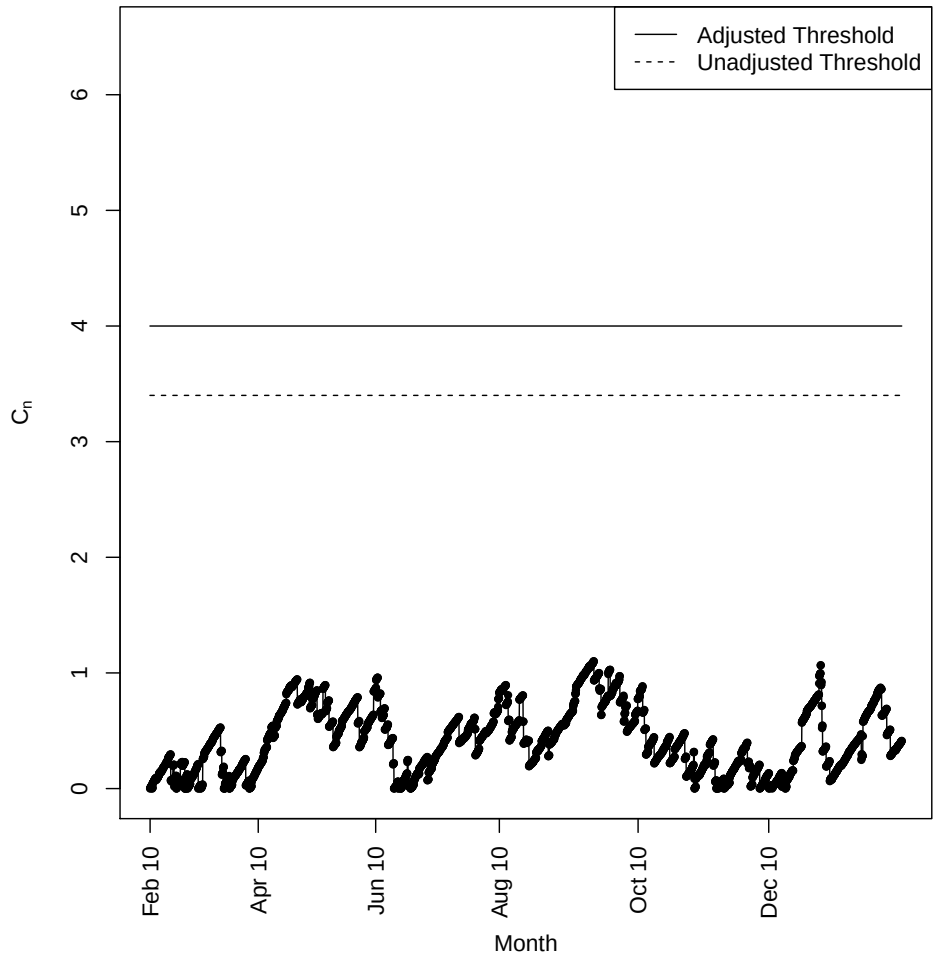

Figure 7.4: Phase I risk-adjusted CUSUM chart detecting a decrease in the probability of infant mortality of the HBB data. The unadjusted and adjusted control limits are 3.4 and 4.0 respectively.

From Figure 7.4 it can be seen that when using both the unadjusted and the adjusted control limits, the chart is in-control. This shows that the Phase I data is assumed to be in-control and the adjusted control limits can be used for the Phase II data.

Phase II

Using the parameter values for Δ, the specified ARL_0 and the control limits found for the Phase I data, a risk-adjusted CUSUM chart for the Phase II data was constructed, using the logistic regression model as shown in Equation 7.3.

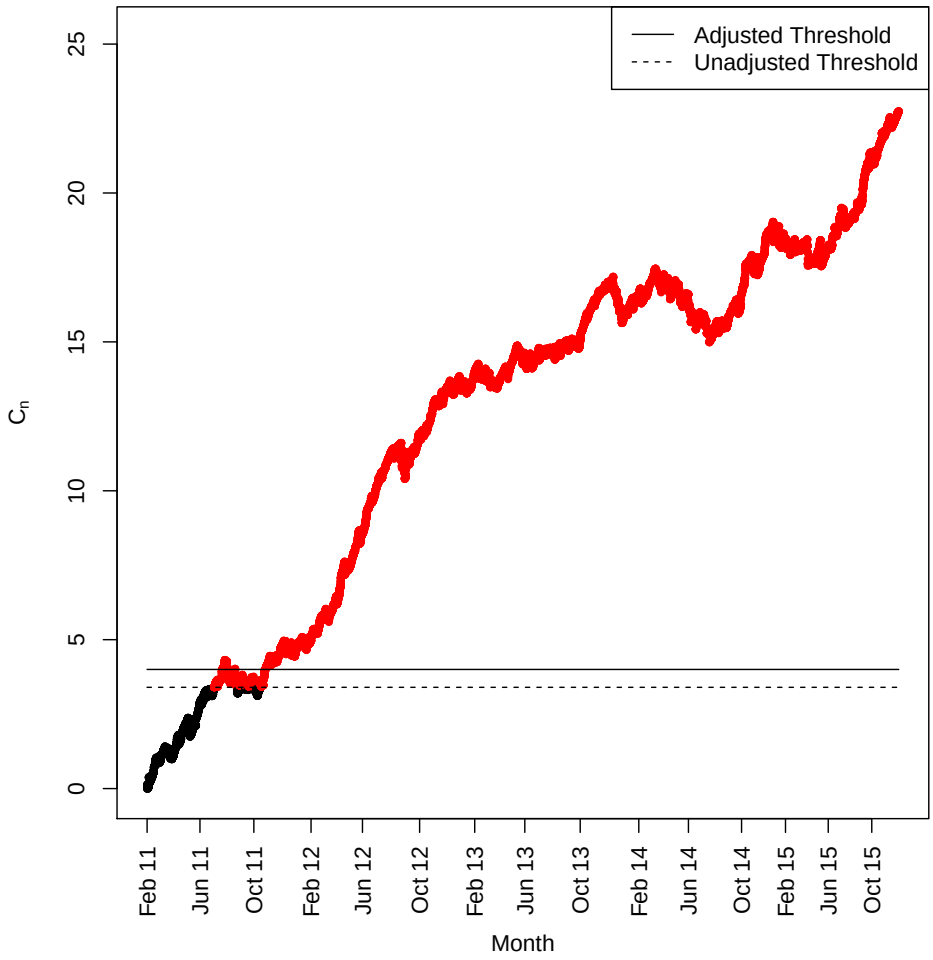

Figure 7.5: Phase II risk-adjusted CUSUM chart detecting a decrease in the probability of infant mortality of the HBB data. The unadjusted and adjusted control limits are 3.4 and 4.0 respectively.

From Figure 7.5 it can be seen that as early as September 2011, the chart is out-of-control for a small period of time, before again signalling out-of-control from around December 2011 onwards. This signal means that the probability of an infant not surviving the first 24 hours of its life has decreased by a clinically relevant amount during Phase II. It can be concluded that the FBOS HBB training is having a positive effect on the survival of infants born at Haydom Lutheran Hospital.

Figure 7.5 can be compared to the basic chart. It can be seen when comparing Figures 7.5 and 7.2 that the risk-adjusted CUSUM chart signals earlier than the basic chart. The most noticeable difference in the two CUSUM charts is between July 2012 and February 2013. During this time the C_n value of risk-adjusted CUSUM chart increases steadily whereas the C_n value of the basic CUSUM chart varies between 5 and 7. This is due to variation present in the basic CUSUM chart which has been accounted for in the risk-adjusted CUSUM chart. In practice this means that the Phase II data contains more difficult cases during this time period, and when this difficulty is taken into account, the number of infants surviving is greater than to be expected. The risk-adjusted CUSUM chart shows that the HBB training has been having a more positive effect on the number of infants surviving than if just the basic CUSUM chart was used to analyse the data.

VLAD for Risk-Adjusted Data

A VLAD can also be constructed in which the expected value for each observation in Phase II is calculated using the regression model in Equation 7.3. The VLAD at each observation j is then

$$V_j = \sum_{i=1}^{j} (Y_{i,E} - Y_{i,O})$$

where $Y_{i,E}$ is the estimated value of Yi found using the regression model in Equation 7.3. This takes into account the difficulty of each case, where a more difficult case will have a higher probability of the infant not surviving the first 24 hours.

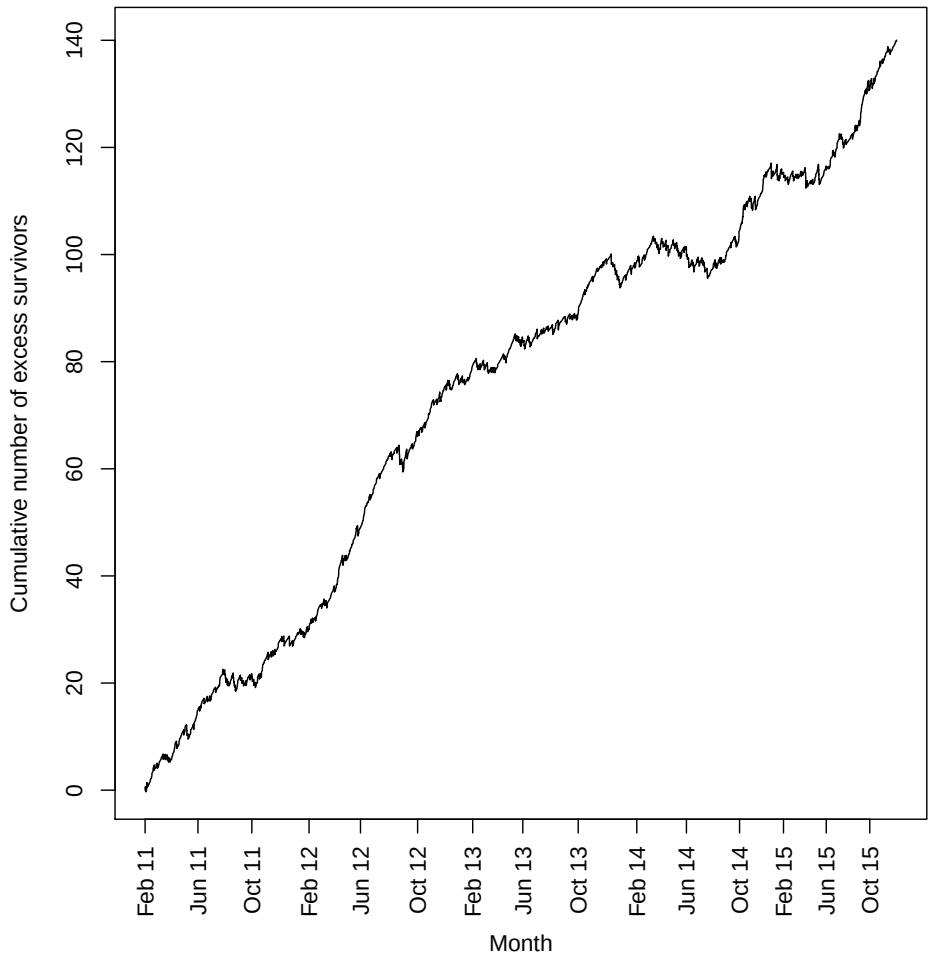

Figure 7.6: VLAD plot for Phase II HBB Data with expected values calculated from the regression model.

Figure 7.6 shows the VLAD for when the expected value is calculated for each Phase II sample using the regression model, Equation 7.3. Figure 7.6 is similar to Figure 7.3, but it does not have the fluctuation from July 2012 to February 2013 which is present in Figure 7.3. During this time period more difficult cases were present and when this is taken into account, the cumulative number of infants surviving continued to increase steadily. It also has a slightly higher overall number of additional infants which survive 24 hours than expected of approximately 140 compared to the 130 which is shown in Figure 7.3. This indicates that Phase II had a greater number of difficult cases than expected, but still the cumulative number

of infants surviving the first 24 hours increased throughout the time period. Both figures agree with the risk-adjusted CUSUM chart that since the introduction of the HBB training program there has been a significant increase in the number of infants surviving 24 hours in Haydom Lutheran Hospital.

7.3 Further Analysis of HBB Data

From Figure 7.5 it can be seen that after February 2014 there is a steady decrease in the C_n value until around August 2014 when the C_n value once again increases. To check that there is not a significant decrease in the number of infants surviving a CUSUM chart which monitors for an increase in the probability that an infant will not survive the first 24 hours is constructed. To monitor a 20% increase in the probability of death gives a probability of $p = 0.033$. The resulting out-of-control average risk score is given by

$$\overline{X\beta} + \Delta = ln\left(\frac{0.033}{1 - 0.033}\right) = -3.38$$

From this Δ can be calculated as

$$\Delta = -3.38 - (-3.56) = 0.18$$

where $\Delta > 0$ gives an out-of-control signal when the probability of death increases. The adjust and unadjusted control limits were then calculated as 3.8 and 3.0 respectively using the Phase I data.

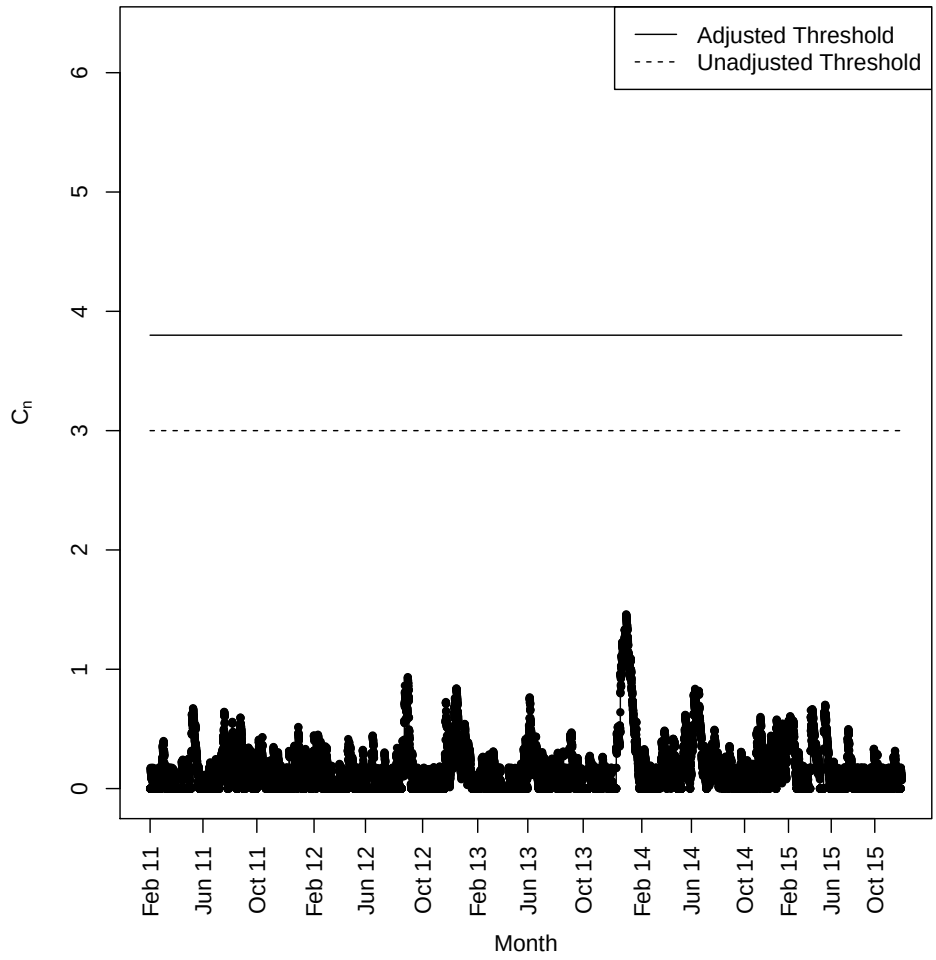

Figure 7.7: Phase II risk-adjusted CUSUM chart for HBB data which detects an increase in the probability of infant mortality.

CUSUM charts for both Phase I and II were created, with the Phase I chart shown in Appendix D used as a check that the Phase I is in control for the adjusted control limits. Figure 7.7 shows that all Phase II observations are in control when monitoring for a 20% increase in the probability of infant mortality. This suggests that since the introduction of the HBB training program there has not been a significant increase in the probability of infant mortality at Haydom Lutheran Hospital.

8. Conclusion

In this thesis the aim was to see how to construct control charts when the chart parameters are estimated and how this estimation affects the performance of the chart. In particular the effect of the estimation errors on the chart's ARL_0 was demonstrated. Three methods which account for the estimation error by adjusting the control limits of the chart were investigated and their effect on the ARL_0 of Shewhart and CUSUM charts was simulated.

The bias reduction method aims to correct the mean of the ARL_0 distribution so it is closer to the specified ARL_0 for Shewhart charts. In chapter 4 the mean of the distribution was greater than the specified ARL_0 and so the bias reduction method shifted the ARL_0 distribution to the left in Figure 5.2. Although the mean of the bias reduction ARL_0 distribution was closer to the specified ARL_0, this increased the number of applications with a run length less than the specified ARL_0, and thus gives an increase in the number of applications which are quicker to give a false out-of-control signal than expected.

The exceedance probability reduction method and bootstrap guaranteed performance method both aim to adjust the ARL_0 distribution so that roughly only a given percentage of the run lengths are less than the specified ARL_0. The exceedance probability method does this by adding adjustment factors to the control limits of Shewhart charts. The bootstrap method estimates a confidence interval for the control limit by producing a number of bootstrapped samples taken from the estimated distribution. This method has the advantage that it can not only be applied to Shewhart charts but other types of control charts as well as risk-adjusted charts. Both methods increase the $ARL1$, but the increase in the ARL_1 is not so great that it outweighs the reduction of false out-of-control signals.

An analysis of the HBB data provided by the SAFER partnership involved constructing CUSUM charts to demonstrate the use of the bootstrap guaranteed performance method. Basic and risk-adjusted CUSUM charts, alongside VLAD plots, clearly illustrate the clinically relevant positive effect that the HBB training program has had on the number of infants surviving 24 hours at Haydom Lutheran Hospital, Tanzania. The risk-adjusted CUSUM chart shows that from as earlier as December 2011 there was a clinically relevant decrease in the probability that an infant does not survive the first 24 hours after birth. VLAD plots show that over the period February 2011 until January 2016, there were 130-140 more infants which survived the first 24 hours than compared to the baseline period of February 2010 to January 2011. The VLAD plots support the conclusions of the CUSUM charts in a way that is easy to interpret.

8.1 Further Work

This thesis focused on three methods for how to account for estimation error in Shewhart and CUSUM charts. A more thorough investigation into other methods which aim to account for the estimation error in these charts, and the methods which aim to account for estimation error in EWMA charts could be completed. The affect of these methods on the in-control and out-of-control ARL should be compared to those of the methods discussed.

When looking at the risk-adjusted CUSUM chart of the HBB data in Figure 7.5 there are periods in which the C_n value is decreasing or fairly stable. This can be seen throughout the period of October 2013 until January 2015. An extension of the HBB analysis could be to see if any of the variables which have been recorded attributes to these trends in the data or if it is due to a greater number of difficult cases than expected. This could be valuable information to the SAFER partnership and could provide a guide on other factors which should be focused on to further increase the number of infants that survive the first 24 hours after birth. These time periods could also be compared to the variable delivery mode that was found to be significant when a multiple regression model was first applied to the data, but had large missingness in Phase I and so was excluding from the regression model. If there was an increase in one of the delivery modes recorded during this time period, it could indicate that this delivery mode affects the probability of the infant not surviving the first 24 hours after birth greater than other delivery modes.

When calculating the control limits of the risk-adjusted chart for the HBB data, the spcadjust R package (Gandy and Kvaløy, 2013) did not produce definitive values for the control limits, but differed when certain parameters of the package varied. For the variable "presentation of the infant was shoulder dystocia" in the multiple regression model, there were only 9 occurrences throughout the Phase I data. This low number of occurrences was a factor in why the spcadjust package did not produce uniform control limits when parameters of the package were varied. Further work could be completed to see if removing variables from the regression model which had few occurrences in Phase I leads to more stable control limits being produced from the spcadjust package, and if the shape of the risk-adjusted CUSUM remains fairly similar to when these variables are included. If so, the risk-adjusted chart produced using the regression model which excludes variables with low occurrences can be compared to the current risk-adjusted CUSUM chart for the Phase II period to see if there is a considerable difference in when the charts signal out-of-control.

Bibliography

W Albers and C Kallenberg. New corrections for old control charts. *Quality Engineering*, 17 (3):467–473, 2005.

D Brook and D A Evens. An approach to the probability distribution of CUSUM run length. *Biometrika*, 59(3):539–549, 1972.

C Champ, S Rigdon, and K Scharnagl. A method for deriving integral equations useful in control chart performance analysis. *Nonlinear Analysis*, 47:2089–2101, 2001.

A Chen and YK Chen. Design of EWMA and CUSUM control charts subject to random shift sizes and quality impacts. *IIE Transcations*, 39(12):1127–1141, 2007.

J Fu, F Spring, and H Xie. On the average run lengths of quality control schemes using markov chain approach. *Statistics and Probability Letter*, 56(4):369–380, 2002.

A Gandy and J T Kvaløy. Guaranteed conditional performance of control charts via bootstrap methods. *Scandinavian Journal of Statistics*, 40:647–668, 2013. doi: 10.1002/sjos.12006.

A Gandy and J T Kvaløy. CUSUM chart with estimated in-control state: Using nonparameteric bootstrap. `http://cran.r-project.org/web/packages/spcadjust/vignettes/CUSUM\_LogisticRegression.html`, 2015.

O Grigg and V Farewell. An overview of risk-adjusted charts. *Journal of Royal Statistical Society: Series A*, 167(3):523–539, 2004.

O Grigg, V Farewell, and D Spiegelhalter. Use of risk-adjusted CUSUM and RSPRT charts for monitoring in medical contexts. *Statistical Methods in Medical Research*, 12:147–170, 2003.

W Jensen, L Jones-Farmer, C Champ, and W Woodall. Effects of parameter estimation on control chart properties: A literature review. *Journal of Quality Technology*, 38(4):349–364, 2006.

M Jones and S Steiner. Assessing the effect of estimation error on risk-adjusted CUSUM chart performance. *International Journal for Quality in Health Care*, 24(2):176–181, 2011.

S Knoth. *spc: Statistical Process Control – Collections of Some Useful Functions*, 2015.

J Lovegrove, O Valenica, T Treasure, C Johnson, and S Gallivan. Monitoring the results of cardiac surgery by variable life adjusted display. *The Lancet*, 350(9085):1128–1130, 1997.

E Mduma, H Ersdal, H Kidanto, B Auestad, and J Perlman. Frequent brief on-site simulation training and reduction in 24-h neonatal mortality-an educational intervention study. *Resuscitation*, 93:1–7, 2015.

D Montgomery. *Statistical Quality Control: A Modern Introduction.* John Wiley & Sons, seventh edition edition, 2013.

R Core Team. *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria, 2015. URL `https://www.R-project.org/`.

B Rao, R Disney, and J Pignatiello. Uniqueness and congergence of solutions to average run length integral equations for cumulative sum and other control charts. *IIE Transactions*, 33 (6):463–469, 2001.

N Saleh, M Mahmoud, M Keefe, and W Woodall. The difficulty in designing shewhart $\bar{X}$ and X control charts with estimated parameters. *Journal of Quality Technology*, 47(2):127–138, 2015.

L Scrucca. qcc: An r package for quality control charting and statistical process control. *R News*, 4/1:11–17, 2004. URL `http://CRAN.R-project.org/doc/Rnews/`.

D Serel and H Moskowtiz. Joint economic design of EWMA control charts for mean and variance. *European Journal of Operational Research*, 184:157–168, 2008.

L Shu, J Wei, and S Wu. A one-sided EWMA control chart for monitoring process means. *Communications in Statistics: Simulation & Computation*, 36(4):901–920, 2007.

S Steiner, R Cook, V Farewell, and T Treasure. Monitoring surgical performance using risk-adjusted cumulative sum charts. *Biostatistics*, 1(4):441–452, 2000.

M Yang, Z Wu, K M Lee, and M Khoo. The x control chart for monitoring process shifts in mean and variance. *International Journal of Production Research*, 50(3):893–907, 2012.

A. Data Sets

The following data set was used as the observations for Figures 2.1, 2.4 and 2.5

Table A.1: Standard Normal Variable Data Set

Sample	Observations				
1	-1.5099	0.7936	-1.7719	1.0944	0.0423
2	1.6475	-0.0421	-1.3108	-0.7379	2.0397
3	-1.6089	0.0002	-1.9091	0.1974	0.1144
4	0.0216	0.8425	2.0311	0.8593	0.1348
5	1.3367	1.2272	1.6683	0.3085	0.9601
6	0.0837	1.0599	0.6030	-0.2527	1.0030
7	-1.8969	-0.6325	1.4666	-1.5282	-1.0241
8	0.7284	-0.5805	-0.2846	0.3179	1.6103
9	0.1622	0.8321	-1.6722	0.8924	-1.0868
10	0.4663	1.3533	-0.2769	0.5870	-1.6260
11	0.9116	-0.1866	-2.0147	-1.0367	0.1585
12	-0.3115	1.1638	-0.4258	0.5765	-0.6901
13	-1.2441	1.1201	-0.6849	1.3108	0.8878
14	0.3440	-0.8275	-0.8371	-0.3521	0.0382
15	-0.0001	0.5875	-0.4194	-0.9640	1.1959
16	-0.5655	0.4085	0.9024	0.6160	1.5717
17	0.0812	-0.1707	-0.8191	-0.1286	0.9844
18	0.7408	0.3426	0.4853	-0.4588	-0.9455
19	0.7782	0.0714	-0.8856	0.6825	-1.2570
20	1.0159	0.2784	0.1522	-0.9618	-1.5155
21	1.1347	-0.1750	-1.8665	0.0416	0.4077
22	0.1539	0.8868	0.0887	-0.3671	-1.7509
23	-0.3886	-0.4860	0.6348	0.2653	1.3424
24	-0.1282	0.1917	0.9025	1.4904	-2.0707
25	2.4954	1.7098	-0.1339	-0.9791	0.3583

B. Impact of Estimation Error: Additional Information

B.1 Samples with Multiple Observations

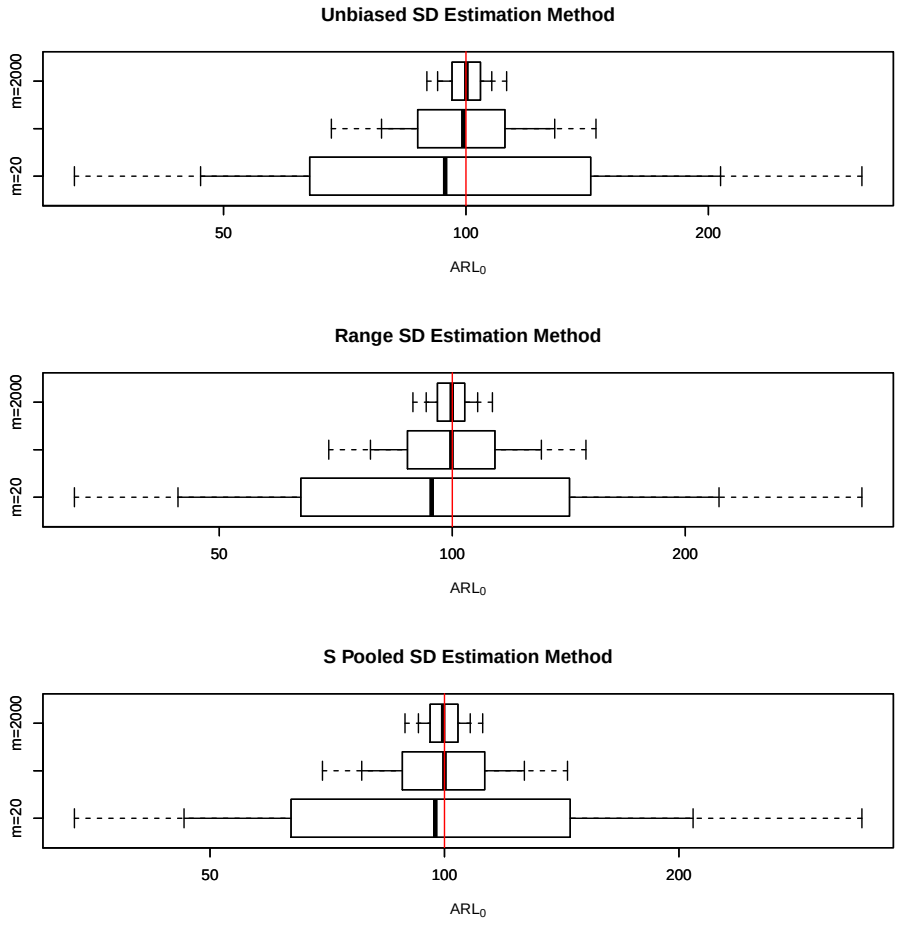

Figure B.1: ARL_0 distribution of Shewhart control Chart with estimated parameters for Phase I data of sample size 20, 200 and 2000, where each sample has 5 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the horizontal axis.

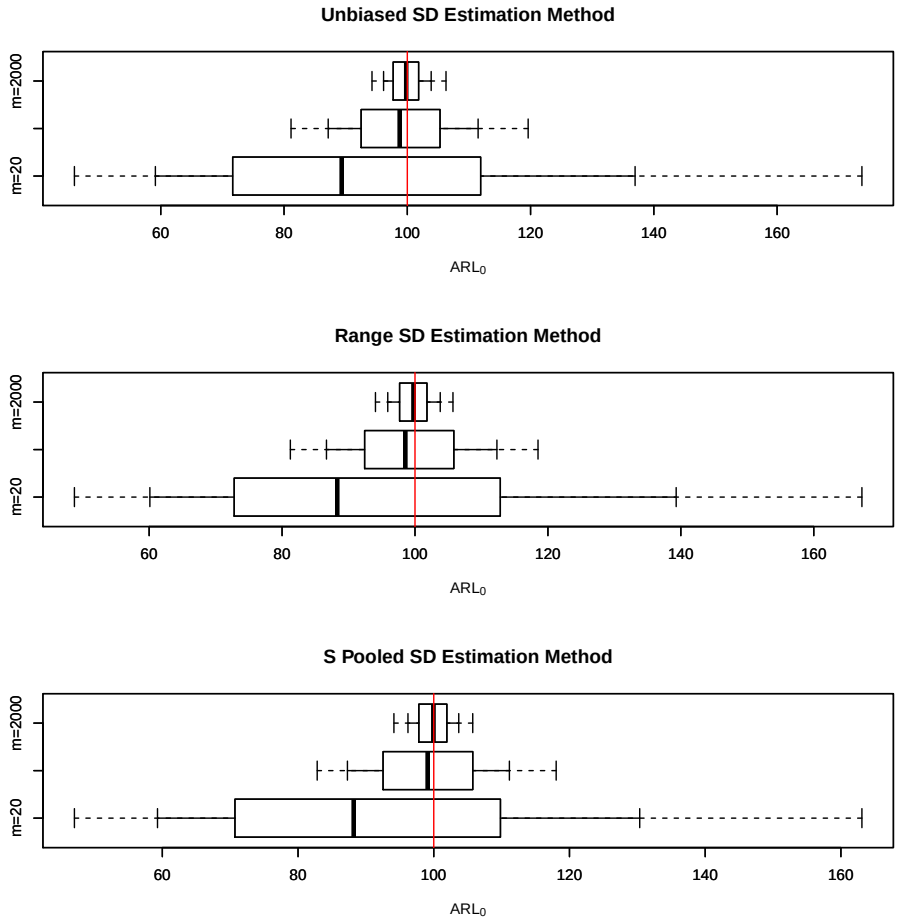

Figure B.2: ARL_0 distribution of CUSUM control charts with estimated parameters for Phase I data of sample size 20, 200 and 2000, where each sample has 5 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution.

B.2 Samples with Single Observations

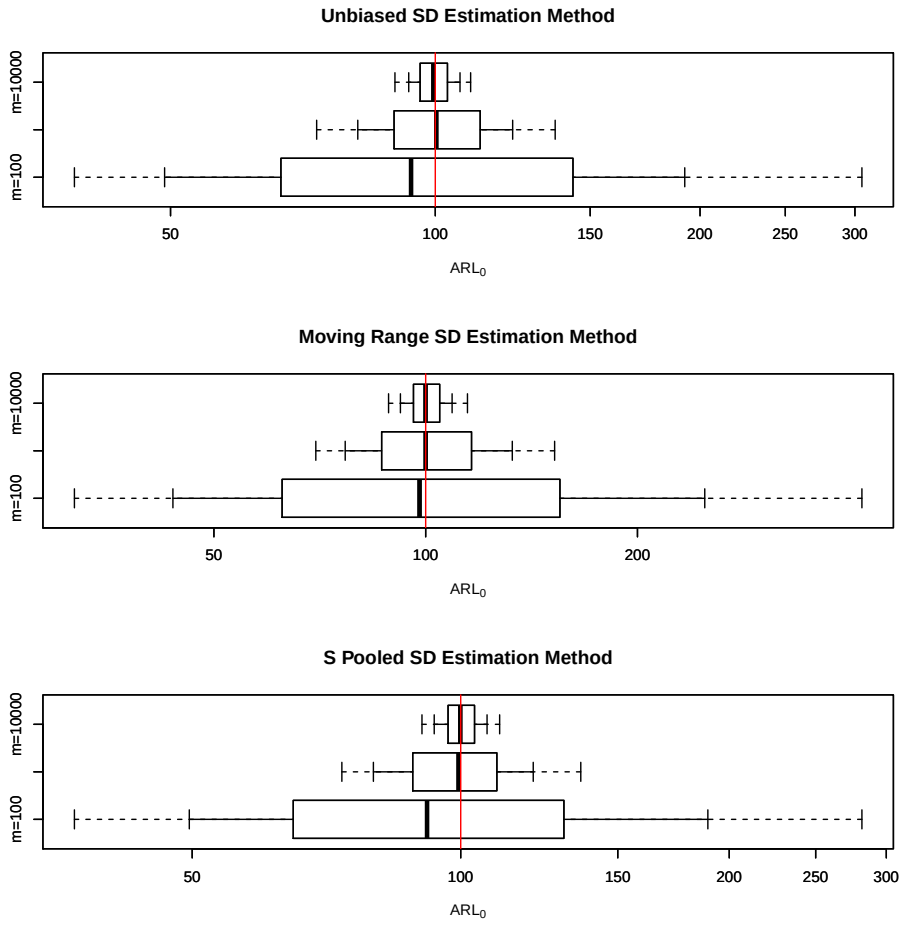

Figure B.3: $ARL_{0,true}$ distribution of Shewhart control charts with estimated parameters for Phase I data of sample size 100, 1000 and 10000, where each sample has 1 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution. A log-scale has been applied to the x-axis.

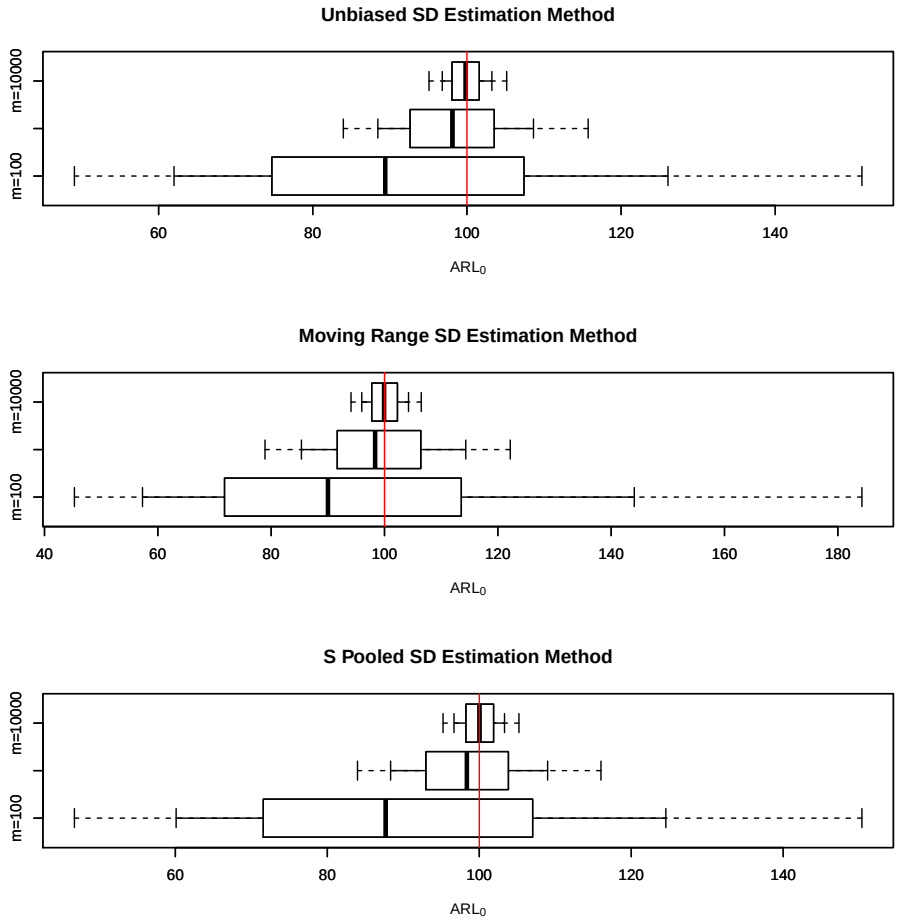

Figure B.4: $ARL_{0,true}$ distribution of CUSUM control charts with estimated parameters for Phase I data of sample size 100, 1000 and 10000, where each sample has 1 observations. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution.

CUSUM Charts with Varied Reference Value

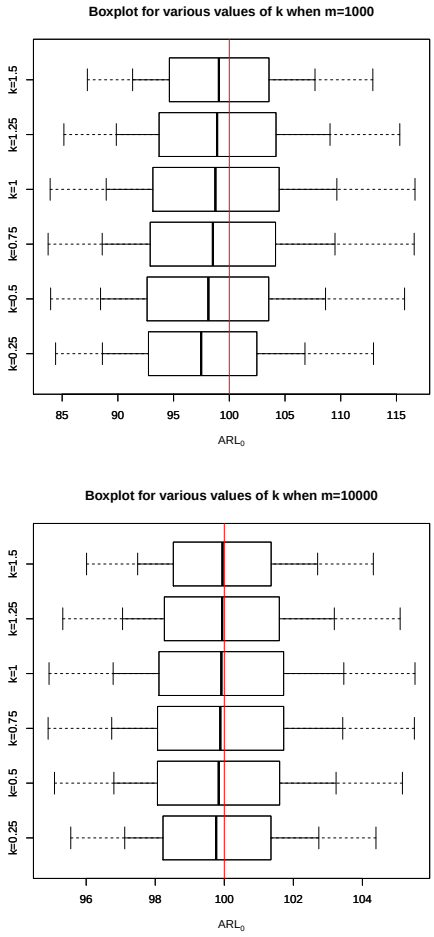

Figure B.5: $ARL_{0,true}$ Distribution of CUSUM Control Charts with Estimated Parameters for Various Reference Values. The boxplots show the 2.5%, 10%, 25%, 50%, 75%, 90% and 97.5% quantiles of the $ARL_{0,true}$ distribution.

C. Accounting for Estimation Error Continued

C.1 Reduction of Standard Deviation Bias

Method I

For the moving range method, where each sample size in $n = 1$, from Equation 3.3 it can be seen that $\sigma* = \overline{MR}$ and $A = d_2(2)$. Using this, it can be found that $\tau^2 \approx 0.826$. This gives $B = \frac{1+0.826L^2}{2}$. This can then be substituted into Equation 5.1 to give

$$UCL = \hat{\mu}_0 + L\hat{\sigma}_0 \left\{ 1 + \frac{1 - 0.826L^2}{2m} \right\}$$
$$LCL = \hat{\mu}_0 - L\hat{\sigma}_0 \left\{ 1 + \frac{1 - 0.826L^2}{2m} \right\}$$

Method II

Corrections can only be applied to the S_{pooled} method when $n > 1$, such that $A = \sqrt{m}$ and $\sigma^* = \sqrt{\sum_{i=1}^{m} s_i^2}$. If $n = 1$ the estimate cannot be split into A and σ^* and so τ^2 cannot be calculated. Albers and Kallenberg (2005) state that when this method is used then $\tau^2 = \frac{n}{2(n-1)}$. This can be used to give

$$B = \frac{1}{4}\left(2 - \frac{nL^2}{n-1}\right)$$

This then gives the control limits as

$$UCL = \hat{\mu}_0 + L\hat{\sigma}_0 \left\{ 1 + \frac{1}{2m} - \frac{nL^2}{4m(n-1)} \right\}$$
$$LCL = \hat{\mu}_0 - L\hat{\sigma}_0 \left\{ 1 + \frac{1}{2m} - \frac{nL^2}{4m(n-1)} \right\}$$

C.2 Reduction of the Exceedance Probability

Method I

When using the moving range average to estimate SD, when $n = 1$, as stated previously $\tau^2 \approx 0.826$. Therefore to find E, $\tau \approx 0.91$ is used. This gives the control limits as

$$UCL = \hat{\mu}_0 + L\hat{\sigma}_0 \left\{ 1 + \frac{0.91\Phi(\alpha)}{m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$
$$LCL = \hat{\mu}_0 - L\hat{\sigma}_0 \left\{ 1 + \frac{0.91\Phi(\alpha)}{m^{\frac{1}{2}}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$

Method II

When the SD has be estimated using the S_{pooled} method and $n > 1$, then as previously stated $\tau^2 = \frac{n}{2(n-1)}$, which leads to $\tau = \left[\frac{n}{2(n-1)}\right]^{\frac{1}{2}}$. Using this to calculate E gives the control limits for the Shewhart chart as

$$UCL = \hat{\mu}_0 + L\frac{\hat{\sigma}_0}{\sqrt{n}} \left\{ 1 + \frac{\Phi(\alpha)}{m^{\frac{1}{2}}} \left[\frac{n}{2(n-1)}\right]^{\frac{1}{2}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$
$$LCL = \hat{\mu}_0 - L\frac{\hat{\sigma}_0}{\sqrt{n}} \left\{ 1 + \frac{\Phi(\alpha)}{m^{\frac{1}{2}}} \left[\frac{n}{2(n-1)}\right]^{\frac{1}{2}} - \frac{\epsilon}{L^2(1-\epsilon)} \right\}$$

D. Helping Babies Breath Analysis Continued

D.1 HBB Regression Model Variables

Table D.1: Covariates Included in Initial Multivariate Regression Model

Covariate
Multiple birth
Pregnancy complication
Source of admission
Maternal infection
Fetal heart rate
Delivery mode
Presentation of the baby
Prolonged labour
Obstructed labour
Pre-eclampsia
Uterine rupture
Cord prolaps
Bleeding before labour
Birth weight, g
Genstational age, weeks

Table D.1 shows the variables which were significant in each individual regression fitted to the HBB data.

D.2 Phase I CUSUM charts

The Phase I CUSUM charts are used to check that the Phase I observations are in-control for the calculated control limits, throughout Phase I.This ensures that the adjusted control limits can then be applied to the Phase II observations.

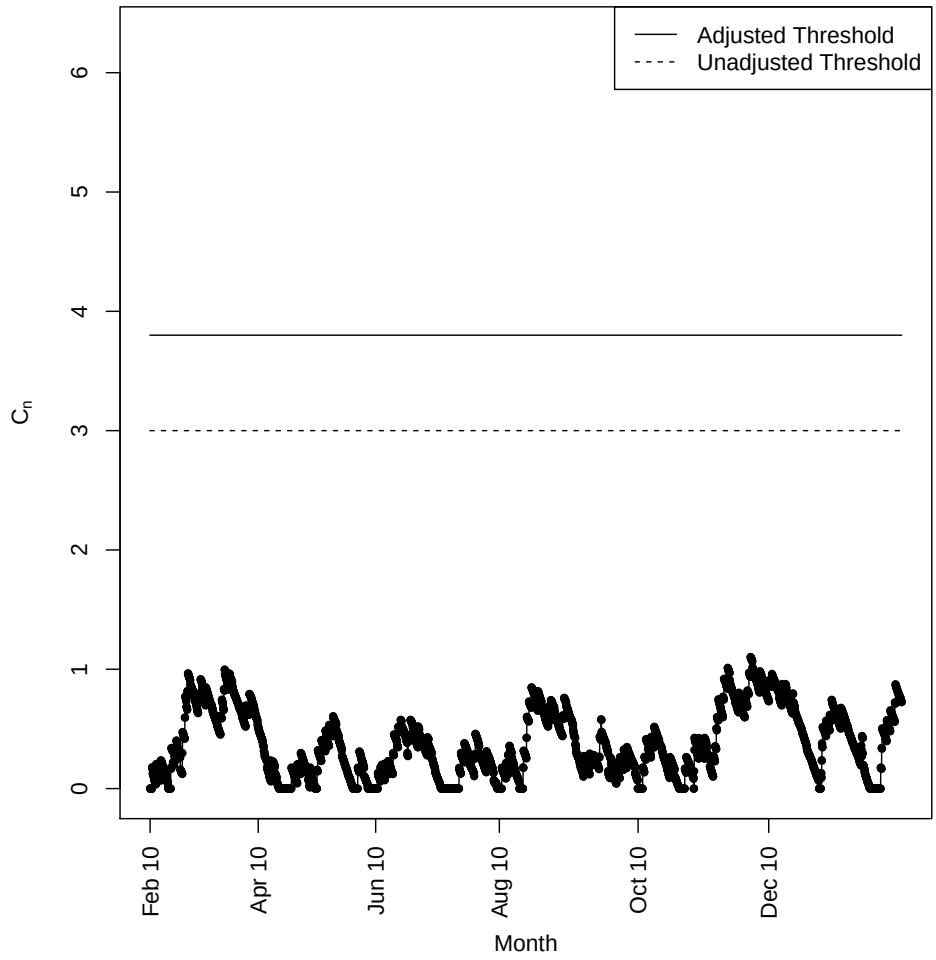

Figure D.1: Phase I basic CUSUM chart for detecting a decrease in the probability of infant mortality of the HBB data.

Figure D.1 shows that for the adjusted control limit all Phase I observations are in control and so this control limit can be used to analysis the Phase II data.

E. R Code

E.1 Impact of Estimation Error

The follow code was used to produce the true ARL_0 distributions used to illustrate the impact of estimation error of Shewhart control charts.

```
library(qcc)
library(spc)
options(scipen=999)

#Specify the number of interations, alpha which for specified in-control
#ARL and corresponding critical values and number and size of samples
B <- 1000
alpha <-0.01
L <- qnorm(1-(alpha/2),0,1, lower.tail=TRUE, log.p=FALSE) #Shewhart cv
crit <- xcusum.crit(mu0=0, L0=RL, k=0.5, sided="two") #CUSUM cv
m<-100
n<- 1

#Define known distribution and critical values
mu <- 0
sigma <- 1

#Produce samples to estimate parameters from
x <-matrix(rnorm((n*m*B), mean=mu, sd=sigma), nrow=(m*B), ncol=n)


for (i in 1:B){
#split x into iteration samples
a<- ((i-1)*m)+1
b<- i*m

if(n==1){
z<-qcc(x[a:b,], type="xbar.one", std.dev="SD", nsigmas=L, plot=FALSE)}
else{
z<-qcc(x[a:b,], type="xbar", std.dev="UWAVE-R", nsigmas=L, plot=FALSE)}

#Estimate parameters from Phase 1 samples
mhat[i] <- z$center
sdhat[i] <- Pooled(matrix(x[a:b,]))
```

```
#Calculate true distribution of critical value and thus true in-control
#ARL for Shewhart charts
mtrue[i] <- (mu-mhat[i])/sdhat[i]
strue[i] <- sqrt((sigma^2)/(sdhat[i]^2))
lower[i] <- pnorm(-L, mtrue[i], strue[i], lower.tail=TRUE)
upper[i] <- pnorm(L, mtrue[i], strue[i], lower.tail=TRUE)
sARL[i] <- 1/(1+lower[i]-upper[i])

#Calculate the true critical values and thus the true in-control
#ARL for CUSUM charts
h[i]<- crit*(sdhat[i]/sigma)
cARL[i]<- xcusum.arl(h[i], k=0.5, mu=(mu0-mhat[i])/(sdhat[i]), sided="two")
}
```

Using the true ARL_0 found in the above code the boxplots in Chapter 4 to illustrate the impact of estimation error were then constructed for various values of m and n.

E.2 Accounting for Estimation Error

This code was used to illustrate the impact of the various methods which take the estimation error of control charts into account for both the ARL_0 and ARL_1 in Chapter 5.

```
library(qcc)
library (spcadjust)
options(scipen=999)

#number of iterations, alpha for specified in-control
#ARL and number and size of samples
B <- 1000
alpha <-0.01
m<-100
n<- 1

#specified eplison and alpha for exceedance probability method
eps <- 0
a<-0.1

#Known distribution parameters and estimated critical value
mu <- 0
sigma <- 1
L <- qnorm(1-(alpha/2),0,1, lower.tail=TRUE, log.p=FALSE)

#Produce samples to estimate parameters from
x <-matrix(rnorm((n*m*B), mean=mu, sd=sigma), nrow=(m*B), ncol=n)


for (i in 1:B){
#split x into iteration samples
```

```
a<- ((i-1)*m)+1
b<- i*m

if(n==1){
z<-qcc(x[a:b,], type="xbar.one", std.dev="SD", nsigmas=L, plot=FALSE)}
else{
z<-qcc(x[a:b,], type="xbar", std.dev="UWAVE-SD", nsigmas=L, plot=FALSE)}

#Estimated parameters
mhat[i] <- z$center
sdhat[i] <- z$std.dev

#Bias adjusted and Exceedance probability adjustements
bLim<-(2-(L^2))/4
eLim<-(pnorm(a,0,1)/sqrt(m))-(eps/(L*L*(1-eps)))

#Bootstrap adjusted control limits
chartShew <- new("SPCShew",model=SPCModelNormal(),twosided=TRUE)
cal<- SPCproperty(data=x[a:b,],nrep=100, property="calARL",
        chart=chartShew, params=list(target=(1/alpha)),quiet=TRUE)
g<-matrix(cal@res, nrow=1, ncol=1)
cLim<-g[1,1]

#True distribution and in-control ARLs
mtrue[i] <- (mu-mhat[i])/sdhat[i]
strue[i] <- sqrt((sigma^2)/(sdhat[i]^2))

lARL[i] <- 1/(1+pnorm(-L, mtrue[i],strue[i], lower.tail=TRUE)-
        pnorm(L, mtrue[i], strue[i], lower.tail=TRUE))

bARL[i]<-1/(1+pnorm(-L*(1+(bLim/m)), mtrue[i], strue[i], lower.tail=TRUE)-
        pnorm(L*(1+(bLim/m)), mtrue[i],strue[i], lower.tail=TRUE))

eARL[i]<-1/(1+pnorm(-L*(1+eLim), mtrue[i],strue[i], lower.tail=TRUE)-
        pnorm(L*(1+eLim), mtrue[i], strue[i], lower.tail=TRUE))

cARL[i]<-1/(1+pnorm(-cLim, mtrue[i], strue[i], lower.tail=TRUE)-
        pnorm(cLim, mtrue[i], strue[i], lower.tail=TRUE))
}

#Calculate out-of-control run lengths for each method
uARL1<- pnorm(L-1, mean=mtrue, sd=strue, lower.tail=TRUE)-
        pnorm(-L-1, mean=mtrue, sd=strue, lower.tail=TRUE)

bARL1<- pnorm((L*(1+bLim))-1, mean=mtrue, sd=strue, lower.tail=TRUE)-
        pnorm(-(L*(1+bLim))-1, mean=mtrue, sd=strue, lower.tail=TRUE)

eARL1<- pnorm((L*(1+eLim))-1, mean=mtrue, sd=strue, lower.tail=TRUE)-
        pnorm(-(L*(1+eLim))-1, mean=mtrue, sd=strue, lower.tail=TRUE)
```

```
cARL1<- pnorm(cLim-1, mean=mtrue, sd=strue, lower.tail=TRUE)-
        pnorm(-cLim-1, mean=mtrue, sd=strue, lower.tail=TRUE)
```

To illustrate the how the bootstrap guaranteed method effected the ARL_0 distribution the following code was use, which was provided by Jan Terje Kvaløy.

```
library(spcadjust)

# The two following functions are calculating the ARL for a CUSUM
# chart using the Markov approximation
getQ <- function(c,gridpoints,pobs){
p <- c(0,(0.5+(0:(gridpoints-2)))*c/(gridpoints-1))
ptarget <- p+c/(gridpoints-1)/2
ptarget[1] <- ptarget[1]- c/(gridpoints-1)/4
sapply(p,function(x) {res <- pobs(ptarget-x); c(res[1],diff(res))})
}

ARL_CUSUM_Markovapprox <- function(c,gridpoints=100,pobs,gridpointsmax=500){
if (c<=0) return(1);
if (pobs(0)>=1) return(Inf);
gridpointsakt <- gridpoints
while(gridpointsakt<gridpoints*8 && gridpointsakt<gridpointsmax){
gridsize <- c/(gridpointsakt-1)/2
if (pobs(gridsize)<0.99) break;
gridpointsakt <- min(gridpointsmax,gridpointsakt*2)
if (gridpointsakt==gridpointsmax) {
warning(paste("Adjustment of grid size reached maximum of",gridpointsmax,
        "grid points."))
}
}
##cannot reach third grid point
if (pobs(c/(gridpointsakt-1)/2.01)>=1) return(Inf);
Q <- getQ(c,gridpointsakt,pobs)
tryCatch(rep(1,gridpointsakt)%*%solve(diag(rep(1,gridpointsakt))-
        Q,c(1,rep(0,gridpointsakt-1))),error=function(e) Inf)
}

#Generate a huge dataset using the true model. These data will be used
#to approximate the true distribution of the updates of the CUSUM when
#the CUSUM is run with data from the true model, but using chart
#parameters and thresholds estimated from a small dataset.
nfuture=1e6
xf <- runif(nfuture,0,70)
beta0=-5
beta1=0.1
pxf=exp(beta0+beta1*xf)/(1+exp(beta0+beta1*xf))
yf=rbinom(n=nfuture,1,prob=pxf)
datafut <- data.frame(y=yf,x=xf)
```

```
nrep=10
ARLtruenaive <- vector(length=nrep)
ARLtrueadjusted <- vector(length=nrep)
nobs=100
beta0=-5
beta1=0.1
Delta=-0.23
chartlogreg<- new("SPCCUSUM", model=SPCModellogregLikRatio
        (Delta=Delta, formula="y~x"))

# iterate from here

for(i in 1:nrep){
x <- runif(nobs,0,70) #put this and the line below outside the loop
#if you want to keep it fixed
px=exp(beta0+beta1*x)/(1+exp(beta0+beta1*x))
y=rbinom(n=nobs,1,prob=px)
dataxy=data.frame(y,x)
xihat<- xiofdata(chartlogreg, dataxy)
cal <- SPCproperty(data=dataxy, nrep=100, chart=chartilogreg,
        property="calARL", params=list(target=100, gridpoints=100))

beta0i=xihat$coefficients[1]
beta1i=xihat$coefficients[2]
xbeta=beta0i+beta1i*xf
updatesfut=Delta*yf + log(1+exp(xbeta)) - log(1+exp(Delta+xbeta))
updatesdist=ecdf(updatesfut) # Approximately the true cdf of the
# updates of the CUSUM when data come
# from the true model and we run with
# estimated parameters and thresholds
ARLtruenaive[i] <- ARL_CUSUM_Markovapprox(c=cal@raw,pobs=updatesdist,
        gridpoints=100)

ARLtrueadjusted[i] <- ARL_CUSUM_Markovapprox(c=cal@res,pobs=updatesdist,
        gridpoints=100)
}
ARLtruenaive
ARLtrueadjusted
```

E.3 Cardiac Surgery Analysis

This code was used to verify the regression model for the cardiac surgery data from Steiner et al. (2000) and produce a risk-adjusted CUSUM chart with adjusted control limits in Chapter 6. This code was provided by Jan Terje Kvaløy with some adjustments made.

```
library(zoo)
SGHdata <- read.table("SGHdata.txt",header=T)

#Use dead within 30 days as response
```

```
dead30 <- as.numeric(SGHdata$Died>=0 & SGHdata$Died<=30)

#Use the two first years of data as phase I sample
dates <- as.Date(SGHdata$OP_DATE, "%d%b%y")
startdate <- "1993-12-31"
daynumber <- as.numeric(dates-as.Date(startdate))
monthnumber=as.numeric((as.yearmon(strptime(dates, format = "%Y-%m-%d"))-
        as.yearmon(strptime(dates[1], format = "%Y-%m-%d")))*12+1)

nmonths=tail(monthnumber,n=1)


estdata <- data.frame(y=dead30[daynumber<0],x=SGHdata$Parsonnet[daynumber<0])
rundata <- data.frame(y=dead30[daynumber>=0],x=SGHdata$Parsonnet[daynumber>=0],
        year=daynumber[daynumber>=0]/365)

#Verify that we get the same estimate as Steiner et al
estmod <- glm(y~x,family=binomial("logit"),data=estdata)
print(summary(estmod))

#Decide on a value of Delta
p<-sum(estdata$y)/length(estdata$y) #Gives Phase I probability of death
EP<-log(p/(1-p)) #Average x^(t)beta value
change<-2 #set change in p wish to detect for mean to be out-of-control
pchange<-p*change
EPdelta<-log(pchange/(1-pchange))
delta<-EPdelta-EP


library(spcadjust)
# Estimate chart parameters
chartlogreg <- new("SPCCUSUM",model=SPCModellogregLikRatio
        (Delta=delta, formula="y~x"))

xihat <- xiofdata(chartlogreg,estdata)
print(xihat)

# Find threshold
cal <- SPCproperty(data=estdata,
nrep=1000,chart=chartlogreg,
property="calARL",params=list(target=10000,gridpoints=200), reportdistr=TRUE)
print(cal)

# First calculating the number of births in each month (the variable nmonth.
nmonth=vector(length=nmonths)
for(i in 1:nmonths){
nosurgery=which(monthnumber==i)
ifelse(length(nosurgery)==0, nmonth[i]<-
        length(which(monthnumber>(i-1)&monthnumber<(i+1))),nmonth[i]<-
```

```
        length(nosurgery))
}

#Then calculating the cumulative number of births by the end of each month:
cnmonth=cumsum(nmonth)
#Cumulative number of births for specified months:
cnmonth2=cumsum(nmonth[1:24])#used for axis for Phase I
cnmonth3=cumsum(nmonth[25:nmonths])#Used for axis for Phase II

# Run chart and plot
S <- runchart(chartlogreg, newdata=rundata,xi=xihat)
par(mfrow=c(1,1),mar=c(4,5,0.5,0.5))
plot(S,ylab=expression(C[n]),xlab="",xaxt="n", type="l",
        ylim=range(S,cal@res+1,cal@raw))

axis(1, las=2, at=c(1, cnmonth3[6]+1, cnmonth3[12]+1, cnmonth3[18]+1,
        cnmonth3[24]+1, cnmonth3[30]+1, cnmonth3[36]+1, cnmonth3[42]+1,
        cnmonth3[48]+1, cnmonth3[54]+1), labels=c("Jan 94","Jul 94","Jan 95",
        "Jul 95","Jan 96","Jul 96","Jan 97","Jul 97","Jan 98","Jul 98"))

mtext(side=1, text="Month", line=4)
points(seq(from=1, to=length(S), by=1),S, type="p", pch=20,
        col = ifelse(S < cal@raw,'black','red'))

lines(c(0,length(S)),rep(cal@res,2),lty=1)
lines(c(0,length(S)),rep(cal@raw,2),lty=2)
legend("topright",c("Adjusted Threshold","Unadjusted Threshold"),lty=1:2)
```

E.4 Helping Babies Breath

To find the adjusted control limits and produce the risk-adjusted CUSUM chart for the HBB data, the following R code was used. This code is an modified version of an original code provided by Jan Terje Kvaløy.

```
library(foreign)  #To read SPSS file
library(zoo)      #For handling dates
hbbdata=read.spss("Complete HBB data.sav",
        use.value.labels = FALSE, to.data.frame = TRUE)
hbbdata=hbbdata[(hbbdata$NEONATAL_OUTCOM<5),]#Removes Neonatal outcome=5
hbbdata=hbbdata[!is.na(hbbdata$NEONATAL_OUTCOM),]#removes Neonatal outcome=NA

#Combine admission 2 and 3
hbbdata$ADMISSION<-replace(hbbdata$ADMISSION, hbbdata$ADMISSION==3, 2)

#Neonatal01 is 1 when neonatal outcome is 3 or 4
neonatal<-replace(hbbdata$NEONATAL_OUTCOM, hbbdata$NEONATAL_OUTCOM==4,3)
hbbdata$NEONATAL01<-as.integer(neonatal==3)
```

```
#Replace Multi so 0 is singular birth and 1 is multiple birth baby
hbbdata$MULTI<- replace(hbbdata$MULTI, hbbdata$MULTI==8,0)
hbbdata$MULTI<- replace(hbbdata$MULTI, hbbdata$MULTI==2,1)
hbbdata$MULTI<- replace(hbbdata$MULTI, hbbdata$MULTI==3,1)

#Factorise significant covariates
hbbdata$INFECTION <-factor(hbbdata$INFECTION,
        levels=c("1", "2", "3", "4", "5"))

hbbdata$FETAL_HEART <-factor(hbbdatabase$FETAL_HEART,
                levels=c("1", "2", "3"))

hbbdata$DELIVERY_MODE <-factor(hbbdatabase$DELIVERY_MODE,
        levels=c("1", "2", "3", "4"))

hbbdata$PRESENTATION <-factor(hbbdatabase$PRESENTATION,
        levels=c("1", "2", "3", "4", "5"))


# Calculate dates, weeknumbers and monthnumbers:
dates=as.Date(as.POSIXct(hbbdata[,1], origin="1582-10-14", tz = "GMT"))
monthnumber=as.numeric((as.yearmon(strptime(dates, format = "%Y-%m-%d"))-
        as.yearmon(strptime(dates[1], format = "%Y-%m-%d")))*12+1)

weeknumber=floor(as.vector(difftime(dates,dates[1],units="weeks")+1))
nmonths=tail(monthnumber,n=1)
nweeks=tail(weeknumber,n=1)

#Compile pre-implementation database containing only significant covariates
hbbdatabase=hbbdata[monthnumber>7 & monthnumber<=19, c(11,13,20,25,46)]
x<-complete.cases(hbbdatabase)

estdata<- data.frame(y=hbbdatabase$NEONATAL01[x], a=hbbdatabase$BIRTH_WEIGHT[x],
        b=hbbdatabase$PREECLAMP[x],c2=as.integer(hbbdatabase$FETAL_HEART[x]==2),
        c3=as.integer(hbbdatabase$FETAL_HEART[x]==3),
        f2=as.integer(hbbdatabase$PRESENTATION[x]==2),
        f3=as.integer(hbbdatabase$PRESENTATION[x]==3),
        f4=as.integer(hbbdatabase$PRESENTATION[x]==4),
        f5=as.integer(hbbdatabase$PRESENTATION[x]==5))


data=hbbdata[monthnumber>=20, c(46,11,13,20,25)]
z<-complete.cases(data)
rundata<- data.frame(y=data$NEONATAL01[z], a=data$BIRTH_WEIGHT[z],
        b=data$PREECLAMP[z],c2=as.integer(data$FETAL_HEART[z]==2),
        c3=as.integer(data$FETAL_HEART[z]==3),
        f2=as.integer(data$PRESENTATION[z]==2),
        f3=as.integer(data$PRESENTATION[z]==3),
        f4=as.integer(data$PRESENTATION[z]==4),
```

```
        f5=as.integer(data$PRESENTATION[z]==5))

#verify model
estmod<-glm(y~a+b+c2+c3+f2+f3+f4+f5, family=binomial("logit"), data=estdata)
summary(estmod)

library(spcadjust)
#Estimate new chart parameters

chartlogreg<- new("SPCCUSUM", model=SPCModellogregLikRatio
        (Delta=-0.23, formula="y~a+b+c2+c3+f2+f3+f4+f5"))

xihat<- xiofdata(chartlogreg, estdata)

#Find threshold
cal<- SPCproperty(data=estdata, nrep=100, chart=chartlogreg,
        reportdistr=TRUE, property="calARL", params=list(target=40000,
        gridpoints=500))

# First calculating the number of births in each month (the variable nmonth.
nmonth=vector(length=nmonths)
monthnumber1<-monthnumber[monthnumber>7 & monthnumber<=19]
monthnumber2<-monthnumber1[x]
monthnumber3<-monthnumber[monthnumber>19]
monthnumber4<-monthnumber3[z]

p1month<-c(monthnumber[monthnumber<=7], monthnumber2, monthnumber4)
for(i in 1:nmonths){
noutcomem=which(p1month==i)
ifelse(length(noutcomem)==0, nmonth[i]<-
        length(which(p1month>(i-1)&p1month<(i+1))),nmonth[i]<-length(noutcomem))

}
#Then calculating the cumulative number of births by the end of each month:
cnmonth=cumsum(nmonth)
#Cumulative number of births starting in February 2010:
cnmonth2=cumsum(nmonth[8:19])
#Cumulative number of births starting in February 2011:
cnmonth3=cumsum(nmonth[20:nmonths])

#Run chart and plot
S<- runchart(chartlogreg, newdata=estdata, xi=xihat)
par(mfrow=c(1,1), mar=c(5,5,0.5,0.5))
plot(S, ylab=expression(C[n]), xlab="", xaxt="n", type="l",
        ylim=range(0, max(cal@raw, cal@res,S)+2))

axis(1, las=2, at=c(1, cnmonth2[2]+1, cnmonth2[4]+1, cnmonth2[6]+1,
        cnmonth2[8]+1, cnmonth2[10]+1), labels=c("Feb 10", "Apr 10",
        "Jun 10", "Aug 10", "Oct 10", "Dec 10"))
```

```
mtext(side=1, text="Month", line=4)
points(seq(from=1, to=length(S), by=1), S, type="p", pch=20,
        col=ifelse(S<cal@raw, 'black', 'red'))

lines(c(0, length(S)), rep(cal@res,2), lty=1)
lines(c(0,length(S)), rep(cal@raw,2), lty=2)
legend("topright", c("Adjusted Threshold", "Unadjusted Threshold"), lty=1:2)
```

By changing the Δ value and the data used and the regression equation to produce the CUSUM chart, this code was used to plot the CUSUM charts in Chapter 7.

To produce the VLAD plots in Chapter 7 the following code was used with estdata, rundata, estmod and cnmonth as defined in the code above.

```
j<-seq(from=1, to=21746, by=1)
dim(j)<-c(21746,1)

#Run chart and plot
yhat<-mean(estdata$y)
S<-c(0*21746)
S[1]<-yhat-rundata$y[1]
for(k in 2:21746){
S[k]<-S[k-1]+(yhat-rundata$y[k])
}

dim(S)<-c(21746,1)
#Produe VLAD for expected Yi=0.0275
par(mfrow=c(1,1), mar=c(5,5,0.5,0.5))
plot(x=j, y=S, ylab="Cumulative number of excess survivors", xlab="",xaxt="n",
        ylim=range(min(S)+0.5, max(S)+0.5), type="l")

axis(1,las=2,at=c(1,cnmonth3[4]+1,cnmonth3[8]+1,cnmonth3[12]+1,
        cnmonth3[16]+1, cnmonth3[20]+1,cnmonth3[24]+1,cnmonth3[28]+1,
        cnmonth3[32]+1,cnmonth3[36]+1,cnmonth3[40]+1, cnmonth3[44]+1,
        cnmonth3[48]+1, cnmonth3[52]+1, cnmonth3[56]+1),
         labels=c("Feb 11","Jun 11","Oct 11","Feb 12", "Jun 12",
         "Oct 12","Feb 13","Jun 13", "Oct 13", "Feb 14", "Jun 14",
         "Oct 14", "Feb 15", "Jun 15", "Oct 15"))

mtext(side=1, text="Month", line=4)

#Calculated Vj for expected Yi from regression model
Sa<-c(0*21746)
Sb<-predict(estmod, newdata=rundata, type="response")
Sa[1]<-Sb[1]-rundata$y[1]
for(l in 2:21746){
Sa[l]<-Sa[l-1]+(Sb[l]-rundata$y[l])
}
dim(Sa)<-c(21746,1)
```

```
#Produce VLAD where expected Yi calculated from regression model
par(mfrow=c(1,1), mar=c(5,5,0.5,0.5))
plot(x=j, y=Sa, ylab="Cumulative number of excess survivors", xlab="",xaxt="n",
        ylim=range(min(Sa)+0.5, max(Sa)+0.5), type="l")

axis(1,las=2,at=c(1,cnmonth3[4]+1,cnmonth3[8]+1,cnmonth3[12]+1,
        cnmonth3[16]+1,cnmonth3[20]+1,cnmonth3[24]+1,cnmonth3[28]+1,
        cnmonth3[32]+1,cnmonth3[36]+1,cnmonth3[40]+1, cnmonth3[44]+1,
        cnmonth3[48]+1, cnmonth3[52]+1, cnmonth3[56]+1),
        labels=c("Feb 11","Jun 11","Oct 11","Feb 12","Jun 12",
        "Oct 12","Feb 13","Jun 13","Oct 13", "Feb 14", "Jun 14",
        "Oct 14","Feb 15", "Jun 15", "Oct 15"))

mtext(side=1, text="Month", line=4)
```