Simulation Study of Average Run Lengths in Statistical Process Control
and Effect of Non-Normally Distributed Data on ARLs

Steven Durtka

Department of Mathematical Sciences
Montana State University

May 4, 2018

A writing project submitted in partial fulfillment
of the requirements for the degree

Master of Science in Statistics

APPROVAL

of a writing project submitted by

Steven Durtka

This writing project has been read by the writing project advisor and has been found to be satisfactory regarding content, English usage, format, citations, bibliographic style, and consistency, and is ready for submission to the Statistics Faculty.

_____

Date

_____

Dr. John Borkowski
Writing Project Advisor

_____

Date

_____

Dr. Mark C. Greenwood
Writing Projects Coordinator

Contents

1   Introduction

Quality Control is an essential component of modern manufacturing. Whether

the product is electronics, automobile parts, or food, manufacturers need to meet

the specifications of customers for product performance, and/or standards

established by regulatory agencies. This is a definition of quality. A product must be

produced consistently to meet these specifications and standards; therefore, the

manufacturing process needs to be monitored to ensure this goal is met.

The monitoring of production involves sampling the process or the product and

applying statistical methods to determine if specifications are being met. This is

called Statistical Process Control (SPC). This study examines a primary tool of SPC,

the Shewhart Control Chart, and its performance characteristics which assume

symmetrical, normally distributed sample data. Using simulation to verify one

theoretical measure of performance, the Average Run Length (ARL), the simulation

tool was then used to explore the effect on the ARL if data are not distributed

normally or symmetrically.

Following this introduction is a discussion of SPC and the Shewhart chart (2.1),

and the Average Run Length calculations (2.2) with Supplementary Runs Rules.

Then the simulation method (3.1) and results (3.2) are shown. Finally, the effect on

the ARL when sample data are simulated from non-normally distributed data,

including non-symmetrical data (4), and the conclusion of the study and thoughts

for future research (5). Section 6 and 7 are References and an Appendix for the

simulation code in R, respectively.

2   Quality Control

To monitor a manufacturing process, various quality characteristics are defined. A quality characteristic is a measure of process or product quality. For example, a manufacturing process may need to be completed at a specified temperature, or perhaps must be maintained for a certain length of time. Such process characteristics then must be monitored throughout the manufacture of a product. Some quality characteristics are measured on the product, and might include the dimensions, strength, capacity, density, or weight of the product. For example, there may be critical dimensions of a product, such as the thickness of a wire. Possibly, the quality characteristic is strength, such as the breaking strength of a plastic component.

2.1 Statistical Process Control and Shewhart Control Charts

Statistical Process Control (SPC) applies the concepts of statistics to these measures of process or product quality to monitor the manufacturing process. Manufacturers use SPC to help them determine whether a process is in or is not in control by taking sample measurements periodically and applying statistics. That is, determine whether or not the process is producing products meeting production requirements, and ultimately is meeting customer specifications.

Shewhart Control Charts are a popular and successful graphical tool used for monitoring quality characteristics. The Shewhart Control Chart consists of several components: the centerline, the control limits, and the warning limits. The

centerline is the target value of the quality characteristic being monitored. The

warning limits are two sampling distribution standard deviations above or below the

centerline. The control limits are three sampling distribution standard deviations

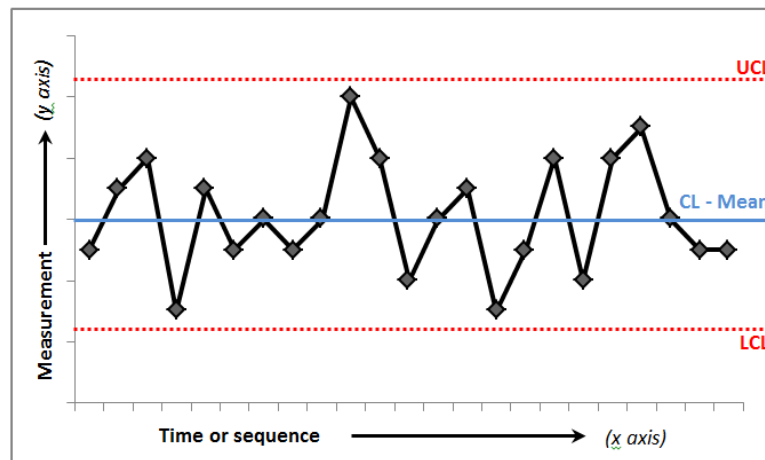above or below the center line. (See Figure 1.)

Figure 1 Shewhart Control Chart with Upper and Lower Control Limits (UCL and LCL), and Center Line (CL)
equal to target mean for the process. UCL and LCL are three standard errors from the mean. (Source:
https://qi.elft.nhs.uk/resource/control-charts/)

In practice, periodic measurements are made on the process or product, and

the sample mean is calculated and plotted on the chart. If the sample mean is more

extreme than the upper or lower control limit, the process is deemed to be out of

control. Statistically speaking, then, each sample mean plotted on the Shewhart

Control Chart is a hypothesis test, with the null hypothesis being that the process is

in control. If the sample mean falls outside the control limits, the null hypothesis is

rejected, and the conclusion is that the process is out of control. Factory or

machine operators then will attempt to determine the cause of the out of control

signal and make any necessary adjustments to the process needed to regain process

control.

If a process is on target, the probability of a sample mean being outside of the

control limits is 0.002699796 under the normality assumption. There is, therefore, a

finite probability of getting an out of control signal even if the process is in control.

Because the cost of halting a process to find problems and adjust the process can be

significant, it is desirable to avoid a false out of control signal, which is equivalent to

a Type I error in statistics. On the other hand, if the process is not in control, and

there is a shift in the mean making the process off target, then you want to quickly

detect that a shift has occurred. Failure to detect a shift is equivalent to a Type II

error, when the null hypothesis is false but is not rejected.

2.2 Average Run Length and Supplementary Runs Rules

The run length is the number of samples taken before an out of control signal is

detected. The Average Run Length (ARL) is the expected value of the run lengths.

When a process is in control, we want the Average Run Length (ARL) to be large,

which would lead to fewer false signals. However, when a process is not in control,

and there is a shift in the mean quality characteristic to be off target, then we want

the ARL to relatively small.

For the Shewhart Control Chart with the three-sigma control limits the run

lengths follow a geometric distribution (Koutras, et al., 2007) with probability of

0.002699796 if the process mean is on target, so the average run length is just the

expected value of such a distribution, which is 370.4. Thus, on average, when the process is in control and the process mean is on target, the average number of samples before an out of control signal is detected is 370.4.

Probabilities of a sample mean landing outside of the three-sigma control limits if the true process mean has shifted can be calculated using probability theory and the normal distribution, and the average run lengths of the shifted mean process can be calculated in a similar manner using the geometric distribution. For convenience, the shifts are in terms of multiples of the standard error, $d = \sqrt{n}|\mu - \mu_0|/\sigma$, where n is the sample size, σ is the process standard deviation, μ is the current process mean from which a sample is taken, and μ_0 is the target process mean (Champ and Woodall, 1987). Absolute values are used because a shift in the process mean above or below the target results in the same probabilities if the measures being sampled are normally distributed. A shift of d = 1 means that the process mean has shifted one standard error from the target mean. The probability that a sample mean falls outside of a control limit if the true process mean has shifted from the target mean by one standard error is

$$P\left(\bar{x} \geq \mu_0 + 3\frac{\sigma}{\sqrt{n}}\right) + P\left(\bar{x} \leq \mu_0 - 3\frac{\sigma}{\sqrt{n}}\right) =$$

$$P\left(\frac{\bar{x} - (\mu_0 - \frac{\sigma}{\sqrt{n}})}{\frac{\sigma}{\sqrt{n}}} \geq 2\right) + P\left(\frac{\bar{x} - (\mu_0 - \frac{\sigma}{\sqrt{n}})}{\frac{\sigma}{\sqrt{n}}} \leq -4\right)$$

$$= P(Z \geq 2) + P(Z \leq -4) = 0.02275013 + 0.00003167$$

$$= 0.02278180.$$

Therefore, the average run length for a one standard error shift in the process mean is expected value of the geometric distribution with the above probability, or 43.9.

A list of ARLs for the Shewhart Control Chart for a set of possible shifts from the target process mean is shown in Table 1. For a large shift, such as three standard errors, the ARL is about two, meaning that on average it would only take two samples to detect the process was out of control. However, for small shifts, the ARL is still relatively large. Even for a one standard error shift, it would take on average nearly 44 samples to obtain an out of control signal, which could result in a considerable amount of defective product being produced before the problem is detected.

To improve the ability of the Shewhart Control Chart to detect small shifts in a process mean, additional methods were developed based on potential patterns in the sample means (Koutras, et al., 2007). These additional detection methods are commonly referred to as Supplementary Runs Rules and are usually applied to the Shewhart chart in addition to the upper and lower control limits. For example, rule two (C2) indicates an out of control process if two of three consecutive sample means are between two and three standard errors above the process mean or below the process mean even though none of the sample means are beyond the three-sigma control limits.

The average run lengths of combining supplementary runs rules with the standard control chart with three sigma limits (C1) have been calculated using a computer program by applying probability theory and Markov chains (Champ and Woodall, 1987). When C1 and C2 are combined (C12), the average run lengths for small shifts in process mean decreases greatly. For example, for a one standard error shift from the mean, the ARL decreases from 43.9 to 20.0. (See Table 1.)

Shift (standard errors from target mean)	ARL 3 sigma (C1)	C1 plus C2 Supplementary Run Rule (C12)
0.0	370.4	225.4
0.2	308.4	177.6
0.4	200.1	104.5
0.6	119.7	57.9
0.8	71.6	33.1
1.0	43.9	20.0
1.2	27.8	12.8
1.4	18.3	8.7
1.6	12.4	6.2
1.8	8.7	4.7
2.0	6.3	3.6
2.2	4.7	3.0
2.4	3.7	2.5
2.6	2.9	2.1
2.8	2.4	1.9
3.0	2.0	1.7

Table 1 Calculated ARL for C1 and C12 by shift from process target mean. Shifts are number of standard errors, and under the normality assumption data are symmetrically distributed, so positive and negative shifts result in same ARL.

3 Methods and Results

The calculations for average run lengths produced by Champ and Woodall (1987), shown in in Table 1, assume that samples are drawn from normally distributed data. An important question is what happens to Average Run Length when sample data originate from non-normally distributed data? The initial purpose of this study is to simulate samples drawn from a normal distribution and verify the calculated ARL's for the Shewhart Control Chart, alone and when combined with rule 2. After verifying the validity of the simulation, use the simulation to explore the effect on ARL's when samples of the same size are drawn from non-normal distributions with the same mean and standard deviation. Three distributions were selected for this study: Uniform, Double Exponential, and Log-Normal. The uniform distribution is symmetrically distributed but has no tail. The double exponential distribution is also symmetrical, however with long tails on both the higher and lower end. Finally, the Log-Normal distribution is skewed to the right and is used to examine the effect of drawing samples from a non-symmetrical distribution.

3.1 Simulation

The availability of computing power to carry out large quantities of calculations has made simulation studies an important research method in statistics. When Champ and Woodall (1987) calculated exact ARLs for Runs Rules, they wrote a program in the BASIC computer language to make the calculations based on

probability theory. Simulations such as bootstrap distributions, however, have become increasingly common to use when theory-based methods may not be convenient or intractable. Today, software specifically created for statistics helps statisticians. For this study, the free, open source software known as R was used to run the simulations and calculate the ARLs. The R package has functions for producing samples from specific probability distributions, as well as program control and decision tools such as loops and if statements which make programming a simulation straight forward.

The R code for the simulation is included in the appendix. For each shift in mean from 0 to 3.0 in 0.2 increments, 1000 simulations were completed. Each simulation involved repeatedly taking samples of size n from a distribution, calculating the sample mean, then applying either the three-sigma test or the runs rule. If an out of control condition was detected, the run length – the number of samples taken until the out of control was detected – was stored. The average of the 1000 simulated run lengths was calculated.

To facilitate conducting simulations for different distributions, a function was written that allowed specifying the mean, variance, sample size, and distribution. The function also allowed for specifying if the shifts would be positive or negative. For symmetrical distributions, shifts of positive or negative have the same ARL. But for skewed distributions like the lognormal, it was expected that the ARL for

negative shifts would differ from the ARL for positive shifts in process mean. (See R

Code in Appendix.)

3.2 Results for Normally Distributed Data

The simulation results for the three-sigma control chart rule and the combined

C12 runs rule is shown in Table 2, alongside the exact values from Champ and

Woodall (1987). The ARLs found using the simulation are very close to the exact

calculated values, indicating the simulation is getting results similar to the expected

ARL's calculated from theory-based methods.

	Simulated		Exact Results	
shift	C1	C12	C1	C12
0.0	357.7	217.5	370.4	225.4
0.2	308.1	171.5	308.4	177.6
0.4	193.3	108.9	200.1	104.5
0.6	114.1	55.8	119.7	57.9
0.8	71.2	33.0	71.6	33.1
1.0	47.2	18.4	43.9	20.0
1.2	26.9	12.3	27.8	12.8
1.4	17.7	8.3	18.3	8.7
1.6	12.2	5.8	12.4	6.2
1.8	8.9	4.2	8.7	4.7
2.0	6.3	3.5	6.3	3.6
2.2	4.5	2.9	4.7	3.0
2.4	3.5	2.3	3.7	2.5
2.6	3.0	2.0	2.9	2.1
2.8	2.5	1.9	2.4	1.9
3.0	2.0	1.6	2.0	1.7

Table 2 ARL's from simulation for three-sigma control chart and C12 runs rule, along with exact ARL's calculated by Champ and Woodall (1987).

4 Simulation Results for Non-Normal Data

The simulation was then applied to the Uniform, Double-Exponential, and Log-Normal distributions to explore the effect of drawing samples from non-normal distributions. The results for the symmetric distributions are shown in Table 3. The ARLs for shifts in process mean greater than 0.6 standard deviations are very similar when samples are drawn from the three different distributions. For samples drawn from data that is distributed uniformly when the shift is less than 0.6, ARL's tend to be greater. For a process that is on target, the simulated ARL is 71.2% greater for uniformly distributed data than for normally distributed data.

Shifts	Exact Normal	Uniform	Double Exponential
0.0	225.4	385.9	122.9
0.2	177.6	256.7	110.2
0.4	104.5	117.7	80.1
0.6	57.9	55.5	52.4
0.8	33.1	32.1	33.7
1.0	20.0	18.5	20.6
1.2	12.8	12.1	13.4
1.4	8.7	7.8	9.0
1.6	6.2	5.6	6.3
1.8	4.7	4.3	4.5
2.0	3.6	3.5	3.5
2.2	3.0	2.8	2.9
2.4	2.5	2.3	2.4
2.6	2.1	2.1	2.0
2.8	1.9	1.9	1.8
3.0	1.7	1.6	1.6

Table 3ARLs from simulation for Uniform and Double Exponential Distributed data compared to the exact results for normal data from Champ and Woodall (1987). Mean and Standard Deviation used in simulations are the same for all distributions.

However, when the data is sampled from a Double Exponential distribution, shifts in mean less than 0.6 tend to have smaller ARLs. When the process is on target, the simulation indicates the ARL will be about 45.5% smaller when data comes from a Double Exponential distribution compared to normally distributed data.

Positive Shifts		Negative Shifts	
Shift	ARL	Shift	ARL
0.0	183.3	0.0	182.8
0.2	120.1	-0.2	196.4
0.4	75.0	-0.4	147.6
0.6	45.2	-0.6	80.7
0.8	29.3	-0.8	41.0
1.0	18.1	-1.0	21.6
1.2	12.1	-1.2	12.3
1.4	8.6	-1.4	8.5
1.6	6.2	-1.6	5.8
1.8	4.4	-1.8	4.3
2.0	3.6	-2.0	3.4
2.2	3.0	-2.2	2.7
2.4	2.4	-2.4	2.3
2.6	2.1	-2.6	2.0
2.8	1.8	-2.8	1.8
3.0	1.7	-3.0	1.6

Table 4 ARL for data sampled from a Log-Normal distribution for positive and negative shifts in process mean. Mean and standard deviation were the same as for the others.

When data are sampled from a skewed distribution, such as the Log-Normal distribution, the simulation shows that there is a slight difference in ARLs for small shifts below the mean compared to shifts above the mean. The simulation results are shown in Table 4. When the shift is negative 0.2, for example, the ARL tends to be greater than the ARL for Log-Normal data when the process is in control, but the ARL for a positive

0.2 shift is less than the in control ARL. For shifts between 0.2 and 1.0, negative shifts

tend to have a greater ARL than for positive shifts in process mean. For shifts in process

mean greater than 1.2, either positive or negative shifts yield similar ARL's as for data

sampled from a normal distribution. When the process is on target, the ARL is less than

under the assumption of normality, about 183 compared to 225 under normality.

5 Conclusion and Suggestions for Further Research

The simulation produced ARLs in agreement with those published by Champ and Woodall (1987) for the Shewhart Control Chart with three sigma control limits combined with runs rule number two. These results are evidence that the simulation method can be used to explore average run lengths for non-normally distributed data. The simulation applied to non-normally distributed data showed that when the data come from symmetrical distributions such as uniform or double exponential, the ARLs for shifts in process mean greater than 0.6 standard errors are about the same as for samples taken from normally distributed data. However, when the process mean is on target, or shifted up to 0.6 standard errors from the target, the ARLs are smaller in the double exponential case than for normally distributed data, but larger in the uniform case. These differences are probably due to the longer tail in the double exponential distribution, and shorter tail in the uniform distribution, when compared to the normal distribution.

The implication is that non-normally but symmetrically distributed data will primarily affect the ARL for processes on target or only having a small shift from the target mean. Larger shifts do not appear effect the ARL compared to the normal assumption ARL, according to the simulation results.

When data are sampled from the lognormal distribution, there are differing effects on ARL's depending on whether the shift is positive or negative and has magnitude between 0.2 and 1.0 standard errors. For shifts greater than 1.0 or -1.0 simulation shows that the ARLs are approximately the same as the other distributions.

Quality managers may want to verify whether the data they are sampling comes from normally distributed data, because of the effect of non-normality on ARLs.   In addition, future research could use simulation studies to study the effects of non-normality on the other runs rules, helping managers to choose which rules will be more effective for quality management. In addition, the effects of other non-normal distributions with interesting properties could be studied, as well as the effect of small sample sizes combined with non-normal distributions.

References

Champ, C. W. (1986) Exact Results for Shewhart Control Charts with Supplementary Runs Rules. Unpublished PhD dissertation, University of Southwestern Louisiana Department of Statistics.

Champ, C. and Woodall, W. (1987) Exact Results for Shewhart Control Charts with Supplementary Runs Rules. Technometrics Vol. 29, No. 4 (Nov. 1987), pp. 393-401.

Champ, C. and Woodall, W. (1990) A Program to Evaluate the Run Length Distribution of a Shewhart Control Chart with Supplementary Runs Rules. Journal of Quality Technology 29: 393 - 399

Montgomery, D.C. (2009), Introduction to Statistical Quality Control, 6th ed., John Wiley & Sons, 2009.

Koutras, M.V., Bersimis, S, and Maravelakis, P.E. (2007) Statistical Process Control using Shewhart Control Charts with Supplementary Runs Rules, Methodology and Computing in Applied Probability 9: 207 – 224.

Nelson, L.S. (1984). The Shewhart Control Chart – Tests for Special Causes. Journal of Quality Technology 15: 237-239.

Nelson, L.S. (1985). Interpreting Shewhart X Control Charts. Journal of Quality Technology 17: 114-116.

Appendix: R Code

```
```{R ARL.C12}
library(dplyr)
library(smoothmest)

arl.C12<-function(muX, sdX, nX, num.sim=1000, dist.type = "normal", upper=TRUE){

seX<-sdX/sqrt(nX)  #standard deviation of X-bar

#set up control chart

UCL<-muX+3*seX
LCL<-muX-3*seX
UWL<-muX+2*seX
LWL<-muX-2*seX

#number sample standard deviations the process has shifted from the target mean
ifelse(upper, shifts<-seq(from=0, to=3, by=0.2), shifts<-seq(from=0, to=-3, by=-0.2))

#set up variables
run.lengths<-seq(1:num.sim)
C12.arl<-seq(1:length(shifts))

for (d in shifts) {

for (i in 1:num.sim) {

run<-0
signal<-0
count.C2<-0

while (signal==0) {
 if(dist.type=="normal"){samp.X<-rnorm(nX, muX+d*seX, sdX)}
  else if(dist.type=="uniform"){
  unif.a<-muX+d*seX-sqrt(3*sdX)
  unif.b<-muX+d*seX+sqrt(3*sdX)
  samp.X<-runif(nX, unif.a, unif.b)}

 else if(dist.type=="lognorm"){
  lnorm.mean<-log(muX+d*seX)-0.5*log((sdX/(muX+d*seX))^2 + 1)
  lnorm.sd<-sqrt(log((sdX/(muX+d*seX))^2 + 1))
  samp.X<-rlnorm(nX, lnorm.mean, lnorm.sd)}
```

```r
 else {samp.X<-rdoublex(nX, muX+d*seX, sqrt(sdX/2))}

 mean.X<-mean(samp.X)
 run<-run+1
 if(mean.X>UCL|mean.X<LCL) {signal<-1}
 if(run==1){mean.x1<-mean.X}
 else{
 if(between(mean.X, LCL, LWL))
  {
  if(between(mean.x1, LCL, LWL)|between(mean.x2, LCL, LWL)) {signal<-1}
  }
 if(between(mean.X, UWL, UCL))
  {
  if(between(mean.x1, UWL, UCL)|between(mean.x2, UWL, UCL))  {signal<-1}
 }
 mean.x2<-mean.x1
 mean.x1<-mean.X
 }

}

run.length[i]<-run

}

C12.arl[abs(d)*5+1]<-mean(run.length)

}

C12<-cbind(shifts, C12.arl)
return(C12)

}

```

```{R C12norm}

#Mean and standard deviation and sample size
muX<-5  #mean
sdX<-1  #standard deviation
nX<-5   #sample size
```

```
C12arl.norm<-arl.C12(muX, sdX, nX)
C12arl.norm
```]

```{R C12unif}

muX<-5  #mean
sdX<-1 #standard deviation
nX<-5   #sample size

C12arl.unif<-arl.C12(muX, sdX, nX, dist.type = "uniform")
C12arl.unif

```
```{R C12logN}

muX<-5  #mean
sdX<-1 #standard deviation
nX<-5   #sample size

C12arl.logNUpper5<-arl.C12(muX, sdX, nX, dist.type = "lognorm", upper = TRUE)
C12arl.logNUpper5

```
```{R C12Dexp}

muX<-5  #mean
sdX<-1 #standard deviation
nX<-5   #sample size

C12arl.Dexp<-arl.C12(muX, sdX, nX, dist.type = "Dexp")
C12arl.Dexp
```
```{R ExportData}
write.csv(C12arl.Dexp, file = "C12arlDexp.csv", row.names = FALSE)
write.csv(C12arl.norm, file = "C12arlNorm.csv", row.names = FALSE)
write.csv(C12arl.unif, file = "C12arlUnif.csv", row.names = FALSE)
write.csv(C12arl.logNLower5, file = "C12arllogNLower5.csv", row.names = FALSE)
write.csv(C12arl.logNUpper5, file = "C12arllogNUpper5.csv", row.names = FALSE)
```
```