



Phase II monitoring of variability using Cusum and EWMA charts with individual observations

John Lawson

To cite this article: John Lawson (2019) Phase II monitoring of variability using Cusum and EWMA charts with individual observations, *Quality Engineering*, 31:3, 417-429, DOI: [10.1080/08982112.2018.1557205](https://doi.org/10.1080/08982112.2018.1557205)

To link to this article: <https://doi.org/10.1080/08982112.2018.1557205>



Published online: 28 Jun 2019.



Submit your article to this journal [↗](#)



Article views: 9



View Crossmark data [↗](#)



Phase II monitoring of variability using Cusum and EWMA charts with individual observations

John Lawson

Department of Statistics, Brigham Young University, Provo, Utah

ABSTRACT

When monitoring a process mean in Phase II, it is well known that time-weighted control charts (such as the Cusum or EWMA) of individual observations are more sensitive for detecting small mean changes than are the traditional Shewhart control charts for individuals. Further, by collecting one observation every 12 minutes, rather than a subgroup of five every hour, the time-weighted charts of individual values result in a shorter ATS (average time to signal) than would be possible using Shewhart charts of subgrouped data. This article explores a similar strategy of monitoring process variability using time-weighted control charts and individual observations. The average time to signal a change in variability using these charts is studied when there are targets or known values for the in-control process mean and standard deviation. The results show that the ATS of both the Cusum and EWMA are substantially shorter than the ATS for the standard R charts or the more efficient S^2 chart using subgroups of 5. The article also describes how the control limits for the EWMA chart to monitor process variability should be modified if the in-control process mean and standard deviations are unknown and must be estimated from a Phase I study. Computer functions that are available in R packages for creating Cusum-EWMA charts and computing their ARL (average run length) are demonstrated in this study and are included in the appendix.

KEYWORDS

ARL; Cusum; EWMA; Phase II monitoring; variability

Introduction

In his transformative book “Out of the Crisis,” Deming (1986) describes several advantages to having a production process in a state of statistical control. Some of the advantages he mentions include 1) the predictability of process performance and cost, 2) smaller inventories of parts required, 3) maximum productivity and minimum costs, 4) improved relationships with customers, and 5) quicker and more reliable detection of unintentional process changes.

Montgomery (2013) echoed some of the same sentiment when he said, “If a product is to meet or exceed customer expectations, generally it should be produced by a process that is stable and repeatable” (p. 188). Deming (1986) went on to say that, “A control chart can ... be used to attain and maintain statistical control during production” (p. 337), and “once a process has been brought into a state of statistical control, it has definable capability” (p. 339).

To attain and maintain statistical control, control charts are used in two different situations called Phase

I and Phase II, as explained by Montgomery (2013). In Phase I, retrospective data are collected and analyzed to construct trial control limits and determine whether the process was in-control during the time the retrospective data was collected. If the process was not in-control during this period, Shewhart control charts assist personnel in identifying the assignable causes for out-of-control signals. If these causes can be removed, then the associated retrospective data can then be eliminated and revised control limits calculated. New data can then be collected and plotted against revised control limits to verify that the process is now in-control. This is often an iterative process. For more discussion of Phase I control chart usage, see Champ and Chou (2003).

The use of control charts in Phase II is different; in this phase statistical control is maintained by monitoring ongoing data on a control chart based on the in-control process mean and standard deviation estimated in Phase I. As explained by Montgomery (2013), most of the large sources of variability are

eliminated in Phase I. Therefore, the control charts used in Phase II should be more sensitive to small process shifts. It is well known that a time-weighted control chart of subgroup means (like the Cusum or EWMA) is more sensitive for detecting small shifts in the mean than the \bar{X} chart. Likewise, Chang and Gan (1995), Ncube and Li (1999), Crowder and Hamilton (1992) and Castagliola (2005) showed that Cusum or EWMA charts of subgroup variances s^2 , or the $\log(s^2)$, are more sensitive to detecting small shifts in the process standard deviation than the Shewhart S^2 chart. Castagliola, Celano, and Fichera (2006) also give a review of EWMA charts for monitoring process variance. It is also known that the standard R chart is even less sensitive than the S^2 chart for monitoring the process variance.

Once control charts have been used to maintain a process in a state of statistical control, the process capability can be determined. In Ford Motor Company's guide to the use of control charts for improving company and supplier quality and productivity (Corporate Quality Education and Training Center Ford Motor Company 1987), it advocates process capability, as measured by Cpk , be ≥ 1.00 across the board. In addition, that $Cpk \geq 1.33$ for new processes affecting control items or other significant product characteristics.

The Manual for Good Practices (ASQC Chemical and Process Industries Division Chemical Interest Committee 1987) states "Values of Cp exceeding 1.33 indicate the process is adequate to meet the customer specifications. Values of Cp between 1.33 and 1.00 indicate that the process, while adequate to meet specifications, will require close control. Values of Cp below 1.00 indicate that the process is not adequate to meet the specifications and that the process and/or specifications must be changed" (p. 37).

When the process mean is centered between the specification limits and $Cp = Cpk = 1.33$, a shift in the mean by 1.33 standard deviations to the left or right will result in Cpk decreasing to $0.89 < 1.00$. Likewise, a 50 percent increase in the process standard deviation will decrease Cpk from 1.33 to 0.89. Therefore, when monitoring the process mean and standard deviation in Phase II, the method used should be able to quickly detect small shifts in the mean by 1.33 standard deviations or more and small increases in the process standard deviation by 50 percent or more.

Control charts of individual observations rather than subgroup means are often used in Phase II for monitoring and controlling the process mean. This is

Table 1. Expected value of $\sqrt{|y_i|}$ and v_i as function of γ .

% Change in σ	γ	$E[\sqrt{ y_i }]$	$E(v_i)$
-20	0.8	0.735379	-0.24860
0	1.0	0.822179	0.00
50	1.5	1.006960	0.52923
70	1.7	1.071990	0.71548

because there is no need to re-estimate the process standard deviation (used to calculate the control limits for the mean) from subgroups of data in Phase II when it is already known or has been estimated from Phase I data. Therefore, individual values rather than subgroups of data may be used in Phase II. In addition, if the cost to take a subgroup of five measurements is no different than the cost to take five individual measurements over time, then it would be better (e.g.) to take one measurement every 12 minutes rather than a subgroup of five measurements every hour. In that way, any shift in the mean will be detected more quickly. When using individual observations, time-weighted control charts like the Cusum or EWMA are more sensitive for detecting small shifts in the mean than Shewhart individual's charts.

However, if individual observations rather than subgroups of data are collected for Phase II monitoring to detect small changes in the process mean, then a standard Shewhart chart like the R chart or the more efficient S chart, which is based on subgrouped data, cannot be used to monitor the process standard deviation. In this situation, a control chart for detecting changes in the process standard deviation based on individual values would be useful.

When the process characteristic (x) to be monitored is normally distributed and its "in-control" mean (μ) and standard deviation (σ) are known, Hawkins (1981, 1993) advocates monitoring the individual values in Eq. [1] with a Cusum chart to detect changes in the process standard deviation.

$$v_i = \frac{\sqrt{|y_i|} - 0.822179}{0.3491508}, \quad [1]$$

where $y_i = (x_i - \mu)/\sigma$ is the standardized process characteristic.

Hawkins (1981) showed that the statistic $\sqrt{|y_i|}$ defined in Eq. [1] has an expected value equal to 0.822179 and standard deviation equal to .3491508 when the "in-control" mean and standard deviation of the process characteristic (x_i) are known. Therefore v_i has mean zero and standard deviation one. In addition, he showed the distribution of v_i is approximately normal when x_i has a normal or a symmetric heavy tailed distribution.

If the standard deviation of the process characteristic x_i increases by a multiplicative factor γ , then the standard deviation of the standardized value $y_i = (x_i - \mu)/\sigma$ increases to γ , and the expected values of $\sqrt{|y_i|}$ and v_i will correspondingly increase.

If y_i follows a normal distribution with mean zero and its standard deviation has increases from 1 to γ , the expected value of $\sqrt{|y_i|}$ can be found as shown in Eq. [2].

$$E\left[\sqrt{|y_i|}\right] = \frac{1}{\gamma\sqrt{2\pi}} \int_{-\infty}^{\infty} \sqrt{|y_i|} \exp\left(\frac{-y^2}{2\gamma^2}\right) dy, \quad [2]$$

Using numerical integration, the expected values of $\sqrt{|y_i|}$ and v_i for four possible values of γ were found and are shown in Table 1.

Therefore, a 50 percent increase in the process standard deviation will result in a 0.592298 standard deviation increase in the mean of v_i , since v_i has mean zero and standard deviation one when the process mean and standard deviation have not changed. Based on this relationship between the mean of v_i and the process standard deviation, Hawkins proposed to monitor individual values of v_i with a Cusum chart to detect small shifts (on the order of 50 percent) in the process standard deviation.

He approximated the average run length (ARL) for detecting changes from a known in-control value of the process standard deviation with a Cusum chart of v_i using simulations of 1000 sets of random data. He also showed, with one set of simulated data, that the Cusum chart reduced the number of observations to detect a 53 percent increase in the process standard by 9.2 percent compared to an R chart with subgroups of size 5. This result cannot be generalized from one set of random data. Additionally, he did not discuss the situation in which the in-control process standard deviation is estimated from Phase I data.

Although the performance of Cusum charts (including accurate tables of ARLs) for monitoring the process mean with individual values has been well studied, the literature still lacks an in depth study of the performance of the Cusum or EWMA chart of v_i for monitoring the process standard deviation with individual values. For that reason, I will focus on the ARL and ATS (average time to signal) of Cusum and EWMA charts for monitoring v_i for changes in the process standard deviation. Specifically, this article shows a more accurate method of computing the ATS of Hawkins's Cusum chart of v_i and a nearly equivalent EWMA chart of v_i . It also compares the ATS for the Cusum and EWMA chart of v_i to the ATS of the R chart and S^2 chart with subgroups of size 5. It

discusses the problem with computed ATS values when the in-control standard deviation is unknown. Finally, it demonstrates a method for adjusting the control limit multiplier for a EWMA chart of v_i to guarantee a target ATS when the in-control standard deviation is estimated from Phase I.

The rest of the article is organized as follows. The next section defines Hawkins (1981)'s Cusum chart for monitoring process variability with individual values and compares his simulated ARL values to more accurately computed values. The performance of this Cusum chart is then compared to the R and S^2 charts based on subgrouped data. The following section compares performance of a similar EWMA chart of individual values for monitoring the process variability to the Cusum chart. Examples are given, and the final section shows how the control limits of the EWMA should be adjusted if the in-control process mean and standard deviation are estimated from a Phase I study.

Comparison of simulated and actual ARL of a Cusum chart of individual observations of v_i

The Cusum chart that Hawkins (1981) recommends for monitoring the process standard deviation is defined by Eqs. [2] and [3].

$$C_i^+ = \max[0, v_i - k + C_{i-1}^+] \quad [3]$$

$$C_i^- = \max[0, -k - v_i + C_{i-1}^-], \quad [4]$$

where $C_0^+ = 0$, and $C_0^- = 0$. The Cusum with reference value $k = .25$ will provide good ARL properties for detecting a one-half standard deviation increase in the mean of v_i .

The mean of $\sqrt{|y_i|}$ and v_i not only change when the process standard deviation changes, but they will also change if the process mean (μ) changes. In addition, if the process standard deviation (σ) increases, not only do the means of $\sqrt{|y_i|}$ and v_i increase, but the chart for monitoring the mean is also sensitive to this kind of shift, in the same way that the \bar{X} chart is sensitive to an increase in the process standard deviation.

For this reason, Hawkins (1993) recommends keeping a Cusum chart of individual values y_i together with a Cusum chart of v_i when monitoring individual values in Phase II. Using both charts simultaneously in this way changes the false alarm rate, and the increased sensitivity doesn't come free. However Hawkins (1993) showed that if the process standard deviation increases, the Cusum chart for y_i used for monitoring the process mean may briefly cross its

control or decision limit (h), but that the Cusum for v_i will cross and stay above its control or decision limit. Therefore, the pattern on the two charts viewed together will help identify a potential false alarm on the chart for the process mean.

Hawkins (1981) approximated the ARL for detecting various values of the percent change in the standard deviation of x_i by simulating 1000 Cusums of v_i (with $k = .25$, and control or decision limit $h = 6$). More accurate values of the ARL can be obtained by numerically solving the ARL integral equation. This numerical solution has been incorporated into the `xcusum.arl` function in the R package `spc` (Knoth 2017).

Table 2 compares the ARL values obtained by `xcusum.arl` and the approximate ARL values obtained by Hawkins' simulations. The R code is shown in the appendix.

Hawkins (1993) compared an R chart with subgroups of size 5 to a Cusum chart of individual values of v_i with respect to their ability to detect a 53 percent increase in the process standard deviation. He found the Cusum reduced the number of observations needed to detect the increase by 9.2 percent. However, this comparison was made with only one set of simulated data and can't be generalized. In the next subsection, I make a more general comparison between the Cusum chart of individual values of v_i and the Shewhart R chart and S^2 -chart of subgrouped data with respect to their ARL and their ATS for detecting an increase or a decrease in process variability. All the ARL and comparisons in this article are based on the assumption that the original data are normally distributed.

Comparing performance of Cusum charts of individual values of v_i to R and S^2 charts for Phase II monitoring of process variability

The average run length (ARL) is the average or expected number of samples before a control chart produces an out-of-control signal. It is often used to compare the performance of different types of control charts. However, when comparing a control chart based on individual values to a control chart based on subgroup statistics, the ARL is not a proper measure to compare the charts. The control chart based on individual values would have a smaller time interval between samples than a chart based on subgroups when both charts have the same number of observations per unit of time. Therefore, the charts should be compared based on the ATS rather than ARL. The

Table 2. Comparison of calculated and simulated ARL's.

% Change in σ	$E(v_i)$	ARL from <code>cusum.arl</code>	ARL From Hawkins' simulations
-20	-0.24860	50.64	61
+50	0.592298	19.38	17
+70	0.715482	13.13	12

Table 3. Performance comparison of R, S^2 and Cusum charts of individual values of v_i for detecting changes in the process standard deviation.

Performance	% Change	R	S^2	Cusum
	in σ	($n = 5$)	($n = 5$)	$k = .25$
Measure				$h = 8.117839$
ARL	-20	∞	63.1	80.29
ATS	-20	∞	315.3	80.29
ARL	0	217.3	74.5	370
ATS	0	1084	370	370
ARL	+50	7.198	5.594	20.29
ATS	+50	35.99	27.97	20.29

steady state ATS is a function of the steady state ARL that is given by $ATS = (ARL - 0.5) \times t$, where t is the time between samples.

The $ATS_1 = (ARL_1 - 0.5) \times t$ is the average time to signal when the process has shifted from the control state and is a measure of how quickly a control chart can detect process shifts. ATS_1 is positively related to the ATS_0 or the average time to signal when the process is in the control state. Therefore, to make a fair comparison of ATS_1 values between different control charts, their ATS_0 values should be equal.

When the in-control process standard deviation (σ) is known, the upper control limit for an R chart (of the standardized process characteristic y_i) that is frequently used in practice for monitoring the process standard deviation in Phase II is

$$UCL = d_2(n) \times D_4(n) \times \sigma, \quad [5]$$

where $d_2(n)$ and $D_4(n)$ are taken from the Appendix D Table of Factors for Control Charts in Christensen, Coombs-Betz, and Stein (2013) or Montgomery (2013). The lower control limit is zero unless the subgroup sizes are greater than 6.

Assuming independent subgroups of size n , the ARL for an R chart (in the in-control or out-of-control state) can be determined from the mean of the geometric distribution: $ARL = 1/(1-\beta)$, where β is the probability of the range R_i falling between zero and the upper control limit given in Eq. [5]. The probability of a signal is $(1-\beta)$. β can be calculated using the `oc.curves.R` function in the R package `qcc` (Scrucca 2017) based on the conventional assumption that the distribution of the observations is normal. The R code for doing this is shown in the appendix.

Table 4. Comparison of ATS or ARL for Cusum and EWMA charts of v_i .

	EWMA ATS	Cusum ATS
% Change	$\lambda = .05$	$k = .25$
in σ	$L = 2.504241$	$h = 8.117839$
-20	72.46	80.29
0	370	370
50	20.56	20.29

The ARL for an S^2 chart can be determined from the fact that when the process is in-control, and the process measurement is normally distributed with variance σ^2 , then $(n-1)s^2/\sigma^2$, where s^2 is sample variance of a subgroup of n follows a Chi-square distribution with $n-1$ degrees of freedom. Probability control limits for the S^2 chart can be determined from the percentiles of the Chi-square distribution, and the Chi-square distribution is used to calculate the probability of an out-of control signal.

The ATS for comparing the the R and S^2 charts to a chart based on individual observations is then calculated as $ATS = (ARL - 0.5) \times t$, where t is the time between subgroup samples expressed as a multiple of the time between individual samples for the chart based on individuals. For example, if one subgroup sample was taken every hour for the R or S^2 chart and one individual sample was taken every 12 minutes for the chart based on individual values, then $t = 5$.

The ARL and ATS of the R and S^2 charts based on subgroups of size $n = 5$ are compared to the ARL and ATS of Cusum charts of individual values of v_i in Table 3. The steady state ARL of the Cusum charts was calculated with the `xcusum.ad` function in the R package `spc`.

In Table 3, the Cusum charts used a small value of $k = .25$ to be sensitive to detecting small changes in the mean of v_i or standard deviation of y_i . The value of the control limit (or decision limit h) for the Cusum charts was chosen to produce the same ATS_0 as the S^2 chart. These values of h were found by modifying the value found using the `xcusum.crit` function in the R package `spc` as shown in the appendix. This function can be used to find the decision limit required to produce a specified ARL_0 . The ARL and ATS for the Cusum chart is the same since the time between samples is $t = 1$. The time between samples for the R and S^2 charts is 5, assuming the Cusum, R and S^2 charts have an equal number of measurements per unit of time.

The ATS_0 for the R chart with 3σ limits, which is frequently used in practice, is not matched to the Cusum chart of v_i and is not directly comparable to the S^2 or Cusum charts shown in the Table 3. The ARL values in the table are not comparable between

the S^2 chart based on subgroups and the Cusum chart based on individuals, but the ATS values are comparable. It can be seen that the time to detect either a 20 percent decrease or 50 percent increase in the standard deviation using the Cusum chart is reduced by 74.5 percent and 27.4 percent respectively over the S^2 chart. Though not directly comparable, the ATS or average time to signal a 50 percent increase in the standard deviation for the Cusum chart is reduced by 43.5 percent over the R chart with 3σ limits that is frequently used in practice.

Comparing performance of EWMA and Cusum charts of individual values of v_i for Phase II monitoring of process variability

EWMA charts can also be used to monitor changes in individual values like v_i . The EMMA at time point i is defined as:

$$z_i = \lambda v_i + (1-\lambda)z_{i-1}, \quad [6]$$

where $z_0 = 0$, and according to Montgomery (2013) popular values for the smoothing constant λ are .05, .10 and .25. For detecting small shifts in the mean of v_i (corresponding to small shifts in the process standard deviation), the value $\lambda = 0.05$ is used.

The control limits for the EWMA chart at time point i are given by:

$$LCL = -L\sqrt{\frac{\lambda}{2-\lambda} [1 - (1-\lambda)^{2i}]} \quad [7]$$

$$UCL = +L\sqrt{\frac{\lambda}{2-\lambda} [1 - (1-\lambda)^{2i}]}, \quad [8]$$

and the value of the control limit multiplier L is chosen to establish the ARL or ATS. The `xewma.crit` function in the R package `spc` can be used to determine the value of the multiplier L . R code, in the appendix, shows how the value obtained from this function was modified to find the multiplier needed to produce an $ARL_0 = ATS_0 = 370$ to match that of the Cusum chart shown in Table 3.

Once the values of L were found, the steady state ARL of the EWMA charts were found using the `xewma.ad` function in the R package `spc`. This function determines the average run length by numerically solving the ARL integral equation. This function was used to produce the $ATS = ARL$ values in column 2 of Table 4, and the R code to do this is shown in the appendix.

Comparing EWMA ATS with that of the Cusum chart, it can be seen that the EWMA chart with $\lambda =$

Table 5. Spacer hole diameters.

Observation	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Diameter	.25	.25	.251	.25	.252	.253	.252	.255	.259	.261	.249	.25	.25	.25	.252

.05 is slightly better than the Cusum chart. The time to detect a 20 percent decrease in the standard deviation using the EWMA chart is reduced by 9.75 percent over the Cusum chart, and the time to detect 50 percent increase is about the same.

Examples

As an example of a Cusum chart to monitor process variability with individual values, consider the data in Table 5. This data comes from Summers (2000) and represents the diameter of spacer holes in surgical tables.

The specification limits are 0.25 ± 0.01 . If the mean was $\mu = 0.25$ and the standard deviation was $\sigma = 0.0025$ then $C_p = C_{pk} = 1.33$. A shift of the mean away from $\mu = 0.25$ by more than 0.0025 in either direction would result in $C_{pk} < 1.0$, and an increase in the standard deviation by more than 50 percent would also result in $C_p < 1.0$. Therefore, Phase II monitoring could be used to detect changes in the process mean and standard deviation from the target values of $\mu = 0.25$ and $\sigma = 0.0025$.

Figure 1 shows the resulting Cusum chart of v_i . It can be seen that the Cusum exceeds the decision interval at observation 10, but falls back within the limits at observation 11. It is unlikely that the process standard deviation would only temporarily change, and it is known that v_i is sensitive to changes in the process mean as well as the standard deviation. To get a better understanding of what has changed in the process, the Cusum chart of the standardized values y_i should also be examined.

Specifying $h = 4.77$ or and $k = .5$ will result in an ARL_0 of 370, and an ARL for detecting a one standard deviation in the mean of approximately 10 (See Tables 9.3 and 9.4 Montgomery 2013). The resulting Cusum chart is shown in Figure 2, and the R code to produce Figures 1 and 2 is shown in the appendix.

In Figure 2, it can be seen that a sustained increase in the mean of the standardized values y_i is detected at observation number 9. Since the mean of v_i increases when the mean of y_i increases, this is most likely the reason why the Cusum chart of v_i also briefly showed an out-of-control signal at observation 10. Making the Cusum charts of both y_i and v_i makes this clear.

A second example of monitoring the process variability with individual observations is illustrated using

a randomly generated data set containing 25 values. The code for generating the data is shown in the appendix, and with the `set.seed(25)` command fixed, re-running the code will regenerate the same data.

The sample mean of the randomly generated data is $\bar{x} = 48.99193$, and the sample standard deviation is $s = 8.280726$. Assuming the known in-control mean of and standard deviation of the process to generate this data was $\mu = 50$, and $\sigma = 5$, the data are monitored with EWMA charts. If the in-control process standard deviation was 5.0, and the in-control mean was 50.0, then it is estimated from the random data that a more than 50 percent increase in the standard deviation has occurred with about a $\frac{1}{5}$ th standard deviation decrease in the mean.

The EWMA charts of y_i and v_i are shown in Figure 3. In Figure 3, it can be seen that an out-of-control signal is first shown on the EWMA chart of y_i at observation number 14, but the EWMA immediately falls back within the control limits at observation 15. However, the EWMA chart of v_i shows a sustained increase and an out-of-control signal at observation 14 and again at observations 19–22 and 24–25. Since the standardized values y_i are sensitive to changes in the standard deviation as well as the mean, in the same way the \bar{X} chart is sensitive to an increase in the process standard deviation, examination of both charts helps in determining what is out of control. The EWMA chart of v_i seems to show a clearer indication that the process standard deviation has increased, and the single out-of-control point on the EWMA of y_i was probably caused by that increase. The R code in the appendix confirms this fact, since the random data was generated from a process with mean 50 and standard deviation 7.5.

Effect of Phase I estimated process parameters on ARL

When a control chart is set up to detect a change in the process mean from a value estimated during a Phase I study, rather than a known value, (Saleh et al. 2016) recommend that a distinction be made between the conditional ARL and the marginal or unconditional ARL. They define the conditional ARL as the ARL of a chart given a specific set of Phase I parameter estimates, and the marginal or unconditional ARL as the conditional ARL averaged over all possible values of the Phase I estimates.

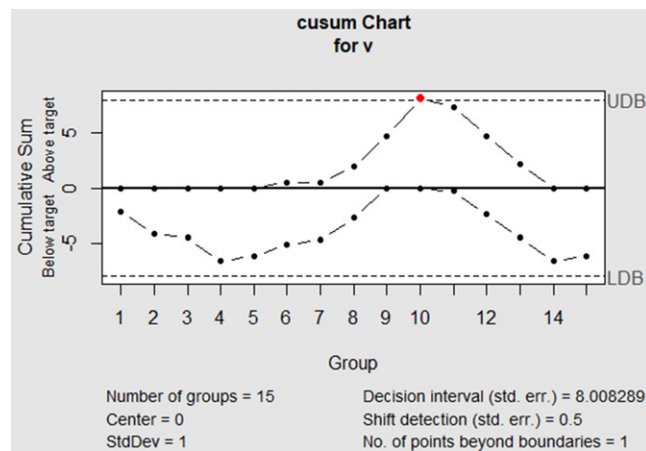


Figure 1. Cusum chart of v calculated from diameters of spacer holes.

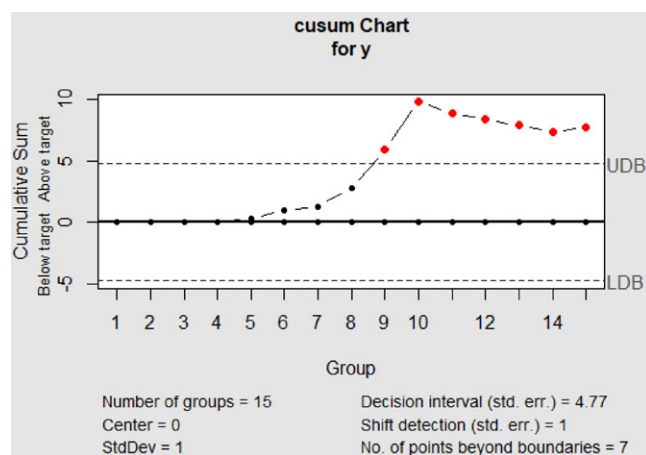


Figure 2. Cusum chart of y standardized diameter of spacer holes.

The distributions of the run length (RL) over all possible Phase I estimate for both R charts and time-weighted charts, like the Cusum or EWMA, are wider than the RL distribution when the parameters (i.e., process mean and standard deviation) are known. When Phase I estimated parameters are used to define the control chart limits, some extremely long ARL values can result in Phase II monitoring when the process is stable. This appears to be a desirable result of using Phase I estimated parameters to define the chart. However, a longer ARL_0 is associated with a longer ARL_1 and less sensitivity in detecting changes in the parameters. In addition, when using Phase I estimated parameters, ARL_0 values that are substantially smaller than the target ARL_0 can occur with non-negligible probability. ARL values less than the target ARL_0 result in more frequent false alarm signals and is another inevitable negative consequence of using Phase I estimated parameters.

(Chen 1998) found that when compared to R charts with a known σ , R charts with estimated

probability limits based on a Phase I estimate of σ will signal more often when the process is stable, and they do not signal as quickly when the process variability has shifted. He showed that when using subgroups of size $n=5$, and a Phase I sample of $m=20$ subgroups (typical numbers used in practice), the marginal ARL (as defined by Saleh et al. 2015) decreases 11.6 percent when compared to the case where σ is known. He also showed that when the standard deviation increases by 40 percent, there is a 26.3 percent increase in the marginal ARL_1 compared to the case where σ is known.

Diko et al. 2017 found corrected probability based control limits for the R chart. Using these limits, rather than the uncorrected probability based limits used by (Chen 1998), the unconditional in-control performance will achieve the desired nominal ARL_0 . However, the out-of-control ARL_1 deteriorates. This deterioration is most noticeable when the number of Phase I samples is small (i.e. 20–30 as normally recommended for R charts) and when the shift in the

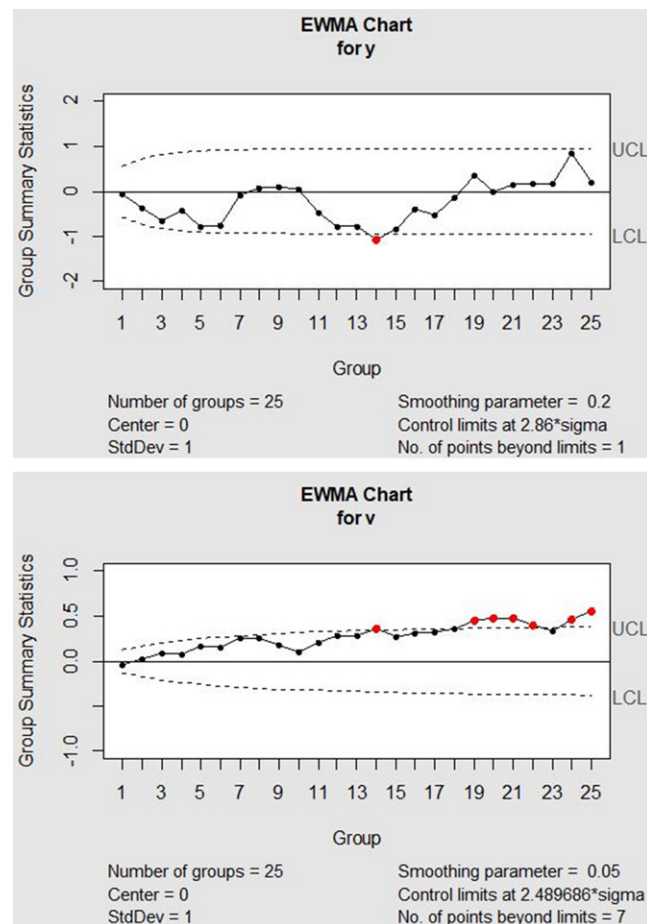


Figure 3. EWMA of random data y (top) v (bottom).

standard deviation is small ($\sigma/\sigma_0 < 2.0$). (Diko et al. 2017) provide R code in the appendix of their article to compute the adjusted control limits.

Epprecht, Loureiro, and Chakraborti 2015 discuss the effect of the amount of Phase I data on the performance of S^2 and S control charts, and (Goedhart et al. 2017) also present a method to determine adjusted control limits for Shewhart S^2 charts. Their proposed adjustment ensures that a minimum in-control performance is guaranteed with a specified probability.

Jones, Champ, and Rigdon 2004 found that the distribution of the conditional in-control ARL of a Cusum chart (over all possible values of the Phase I estimates) was extremely right skewed, and that there are many values much smaller than the desired ARL_0 . This type of ARL distribution will cause a higher false alarm rate when the process is in-control. Saleh et al. 2015 came to similar conclusions regarding the conditional in-control distribution of the EWMA chart.

To counteract this kind of problem, Gandy and Kvaloy 2013 proposed that a control chart should be designed to have at least a specified conditional ARL

Table 6. Percentiles of conditional ARL_0 distribution of the EWMA of v_i with $\lambda = .05$ and target $ARL_0 = 370$.

Percentile	Unadjusted $L = 2.489686$	Adjusted $L = 3.2774$
.10	129.86	477.73
.50	280.02	1762.17
.90	365.87	3064.26

with a specified probability. They proposed a method based on bootstrap samples to find the 90th percentile of the control limit multiplier to guarantee the conditional ARL to be at least equal to the target ARL_0 90 percent of the time. They found that this increase to the in-control ARL only slightly increased the ARL for detecting a one standard deviation shift in the process mean. Saleh et al. 2016 similarly found bootstrap adjusted control limits for the two-sided Cusum chart to guarantee conditional in-control ARL to be at least a specified value.

The R package `spcadjust` (Gandy and Kvaloy 2015) contains the function `SPCproperty` that computes the adjusted multiplier ($L \times \sqrt{\frac{\lambda}{2-\lambda}}$) shown in Eqs. [7] and [8] for the two-sided EWMA control limits. When the process parameters are estimated from a Phase I study, the `SPCproperty` function can be used in place

Table 7. Percentiles of conditional ARL_1 distribution of the EWMA of v_i for a 50 percent increase in σ with $\lambda = .05$ and target $ARL_0 = 370$.

Percentile	Unadjusted $L = 2.489686$	Adjusted $L = 3.2774$
.10	16.14	24.39
.50	20.92	33.54
.90	29.38	52.24

of the `xewma.crit` function in the R package `spc` that was illustrated in the section “Comparing performance of EWMA and cusum charts of individual values of v_i for Phase II monitoring of process variability”. An example of the use of `SPCproperty` is shown in the R code in the appendix.

Table 6 shows the percentiles of the in-control ARL_0 distribution determined from 1,000,000 simulations with 100 Phase I samples used to estimate the process parameters. In this table it can be seen that for the unadjusted EWMA chart with $L = 2.489686$, the conditional ARL is less than the target for more than 90 percent of the users or practitioners, which will result in many false alarm signals. However, when using the adjusted value of $L = 3.2774$, the conditional ARL is greater than the target value for at least 90 percent of the users. Thus, using the adjusted L will prevent the false alarm signals that result from basing the chart on Phase I estimated parameters.

Table 7 shows the percentiles of the ARL_1 for detecting a 50 percent increase in the standard deviation determined from 1,000,000 simulations with 100 Phase I samples used to estimate the process parameters. While the 50th percentile of the adjusted ARL_0 in Table 6 is 6.29 times as large as the unadjusted ARL_0 , the 50th percentile of the adjusted ARL_1 in Table 7 is only 1.60 times as large as the unadjusted ARL_1 .

To illustrate the use of an EWMA of v_i with corrected limits to detect a change in the process variability from the Phase I estimated values, consider again the randomly generated data for the second example in the section “Examples.” In that example, it was assumed that the in-control process mean and standard deviation were known. Here, it is assumed that the in-control mean and standard deviation are unknown and are estimated from a Phase I study. The Phase I data for this example consists $m = 100$ values randomly generated values from the in-control distribution with $\mu = 50$ and $\sigma = 5$. The Phase II data was the same data used in the second example in the section “Examples” and was randomly generated from the distribution with $\mu = 50$ and $\sigma = 7.5$ using `set.seed(25)`. This data represents a 50 percent increase in the process standard deviation over the distribution that generated the Phase I sample.

The standardized values y_i were created by subtracting the Phase I estimated mean, \bar{x} , and dividing by the Phase I estimated standard deviation, s . The v_i values were calculated from y_i . Next, the adjusted control limit multiplier L for the EWMA charts of y_i and v_i were determined using the `SPCProperty` function in the R package `spcadjust`. Finally, EWMA charts of y_i and v_i were made. The R code to do all this is shown in the appendix. Running the code produces EWMA charts that look very similar to those shown in Figure 3. The control limit multipliers are slightly larger than those used to create Figure 3, but the conclusions remain the same. Therefore, if the in-control process mean and standard deviation are unknown and are estimated from 100 or more Phase I samples, the EWMA charts of the individual values y_i and v_i can still detect a 50 percent increase in the standard deviation within 20–25 observations.

Conclusions and recommendations

The conclusions of this study show the Cusum and EWMA charts of individual values will detect changes (on the order of 50 percent) in the process standard deviation more quickly (based on the ATS) than either the standard R chart often used in practice or the more efficient S^2 chart or the equivalent S chart.

The two-sided EWMA chart with $\lambda = .05$ and $L = 2.489686$ was shown to be slightly better than the Cusum chart recommended by Hawkins (1981), and code is shown for adjusting the control limit multiplier L to guarantee a target value of ARL_0 when the in-control process mean and standard deviation must be estimated from a Phase I study.

In some processes like continuous chemical processes, it is inconvenient or impossible to collect rational subgroups of data needed for a standard R or S^2 chart. In other cases, the cost of sampling individual observations is no different from the cost of sampling subgroups of observations when the number of observations per unit of time is the same. Based on the results of this study, in either of these situations, it is recommended that Phase II monitoring be conducted using EWMA charts of individual standardized values y_i and v_i to reduce the time to detect changes in the process.

If it is more convenient or cost effective to sample in rational subgroups during Phase II monitoring, EWMA charts of the subgroup means and the EWMA chart of the log transformation of the subgroup sample variances recommended by Crowder and Hamilton (1992) will both detect small changes in

the process faster than the \bar{X} and S^2 charts. Therefore, in this situation it is recommended that EWMA charts again be used in Phase II monitoring.

About the author

John Lawson is a Professor in the Department of Statistics at Brigham Young University. His research focuses on the design and analysis of experiments and statistical quality control.

References

- ASQC Chemical and Process Industries Division Chemical Interest Committee. 1987. *Quality for the chemical and process industries: A manual for good practices*. Milwaukee, WI: American Society for Quality Control, ISBN 0-87389-035-3.
- Castagliola, P. 2005. A new s^2 -ewma control chart for monitoring the process variance. *Quality and Reliability Engineering International* 21 (8):781–794.
- Castagliola, P., G. Celano, and S. Fichera. 2006. Monitoring process variability using ewma. In *Springer handbook of engineering statistics*, 291–325. London: Springer.
- Champ, C. M., and S. P. Chou. 2003. Comparison of standard and individual limits phase I shewhart \bar{X} , r , and s charts. *Quality and Reliability Engineering International* 19 (2):161–170.
- Chang, T. C., and F. F. Gan. 1995. A cumulative sum control chart for monitoring process variance. *Journal of Quality Technology* 27 (2):109–119.
- Chen, G. 1998. The run length distributions of r , s , and s^2 control charts when σ is estimated. *Canadian Journal of Statistics* 26 (2):311–322.
- Christensen, E. H., K. M. Coombs-Betz, and M. S. Stein. 2013. *The certified quality process analyst*. 2nd ed. Milwaukee, WI: ASQ Quality Press.
- Corporate Quality Education and Training Center Ford Motor Company. 1987. *Continuing process control and process improvement*. Dearborn, MI: Corporate Quality Education and Training Center Ford Motor Company.
- Crowder, S. V., and M. D. Hamilton. 1992. An EWMA for monitoring a process standard deviation. *Journal of Quality Technology* 24 (1):12–21.
- Deming, W. E. 1986. *Out of the crisis*. Cambridge, MA: MIT Center for Advanced Engineering Study.
- Diko, M. D., R. Goedhart, S. Chakraborti, R. J. M. M. Does, and E. K. Epprecht. 2017. Phase II control charts for monitoring dispersion when parameters are estimated. *Quality Engineering* 29 (4):605–22.
- Epprecht, E. K., L. D. Loureiro, and S. Chakraborti. 2015. The effect of the amount of phase I data on the phase II performance of s^2 and s control charts. *Journal of Quality Technology* 47 (2):139–155.
- Gandy, A., and J. T. Kvaloy. 2013. Guaranteed conditional performance of control charts via bootstrap methods. *Scandinavian Journal of Statistics* 40 (4):647–668.
- Gandy, A., and J. T. Kvaloy. 2015. *spcadjust: Functions for Calibrating Control Charts*. R package version 0.5.4.
- Goedhart, R., M. M. da Silva, M. Schoonhoven, E. K. Epprecht, S. Chakraborti, R. J. Does, and A. Veiga. 2017. Shewhart control charts for dispersion adjusted for parameter estimation. *IISE Transactions* 49 (8): 838–848.
- Hawkins, D. M. 1981. A cusum for a scale parameter. *Journal of Quality Technology* 13 (4):228–235.
- Hawkins, D. M. 1993. Cumulative sum control charting: An underutilized spc tool. *Quality Engineering* 5 (3): 463–477.
- Jones, L. A., C. W. Champ, and S. E. Rigdon. 2004. The run length distribution of the cusum with estimated parameters. *Journal of Quality Technology* 36 (1):95–108.
- Knuth, S. 2017. *spc: Statistical Process Control – Collection of Some Useful Functions*. R package version 0.5.4.
- Montgomery, D. C. 2013. *Introduction to statistical quality control, Seventh edition*. New York: John Wiley and Sons.
- Ncube, M., and K. Li. 1999. An ewma-cuscore quality control procedure for process variability. *Mathematical and Computer Modeling* 29:181–192.
- Saleh, N. A., M. A. Mahmoud, L. A. Jones-Farmer, I. Zwetsloot, and W. H. Woodall. 2015. Another look at the ewma control chart with estimated parameters. *Journal of Quality Technology* 47 (4):363–382.
- Saleh, N. A., I. M. Zwetsloot, M. A. Mahmoud, and W. H. Woodall. 2016. Cusum charts with controlled conditional performance under estimated parameters. *Quality Engineering* 28 (4):402–415.
- Scrucca, L. 2017. *qcc: Quality Control Charts*. R package version 1.1.
- Summers, D. C. S. 2000. *Quality second edition*. Upper Saddle River, NJ: Prentice Hall.

Appendix – R code from the article

The block of R code below was used to calculate the ARL values for the Cusum chart shown in Table 2

```
library(spc)
xcusum.arl(k=.25,h=6,mu=-.2486,sided="two")
xcusum.arl(k=.25,h=6,mu=.52923,sided="two")
xcusum.arl(k=.25,h=6,mu=.71548,sided="two")
```

The block of R code below was used to calculate the ARL of the R chart with 3σ limits in Table 3.

```
library(qcc)
sdata<-rnorm(100,0,1)
# Get ARL for R chart for detecting 50
# change in sigma when n=5
robj<-qcc(sdata,type="R",sizes=5,nsigmas=3.0,plot=FALSE)
beta<-oc.curves.R(robj,n=5,c=1.5)
ARL0<-1/(1-beta[1,1])
ARL1<-1/(1-beta[11,1])
```

In this code `sdata` is simulated random normal data. It is needed as the first argument for the `qcc` function that creates a `qcc` object `robj`. A `qcc` object is needed as first argument to the `oc.curves.R` function (although the specific value of the first argument for the `qcc` function will not affect the result of the `oc.curves.R` call). The other two arguments are the subgroup size (5) and the multiplier (1.5) of the standard deviation to be detected. The ARL calculated by this code is 7.198. This code, with changes the multiplier, was used to produce the third column in Table 3. Note that the ARL for detecting a decrease in the standard deviation using an R chart with subgroups of size 5 is infinity because the range R_i can never fall below the lower control limit of zero.

The block of R code below was used to calculate the ARL of the S^2 chart with probability limits to produce an ARL = 74.5 resulting in an ATS = 370 in the fourth column Table 3.

```
# S^2 ARL0
ARL<-74.5
alpha<-1/74.5
ll<-qchisq(alpha/2,4)
ul<-qchisq((1-alpha/2),4)
# check
1/(1-(pchisq(ul,4)-pchisq(ll,4)))
# 20
mult<-(1/.8)^5E;2
ARL1<-1/(1-(pchisq(mult*ul,4)-pchisq(mult*ll,4)))
# 50
mult<-(1/1.5)^2
ARL1<-1/(1-(pchisq(mult*ul,4)-pchisq(mult*ll,4)))
```

The block of code below shows how the values of h for the Cusum chart in Tables 3 and 4 were found by modifying the value found using the `xcusum.crit` function in the R package `spc`, and how the Cusum ARL was calculated using the `xcusum.ad` function in the R package `spc`.

```
# Cusum ARL0
library(spc)
h0<-xcusum.crit(k=.25,L0=370,mu=0,sided="two")
h0<-h0+.10955
ARL0<-xcusum.ad(k=.25,h=h0,mu0=0,mu1=0,sided="two")
# Cusum ARL1
# 20
ARL1<-xcusum.ad(k=.25,h=h0,mu0=0,mu1=-.24860,sided="t")
# 50
ARL1<-xcusum.ad(k=.25,h=h0,mu0=0,mu1=.592258)
```

The block of code below shows how the control limits multiplier L for the EWMA chart in Table 4 was found using the `xewma.crit` function in the R package `spc` and how the ARL for the EWMA chart was found using the `xewma.crit` function in the R package `spc`.

```
# Get EWMA ARL=ATS for Table 4
library(spc)
cv<-xewma.crit(l=.05,L0=370,mu0=0,sided="two")
cv<-cv+.014555
# ARL0
xewma.ad(l=.05,c=cv,mu0=0,mu1=0,sided="two")
# 20
xewma.ad(l=.05,c=cv,mu0=0,mu1=-.2486026,sided="two")
# 50
xewma.ad(l=.05,c=cv,mu0=0,mu1=.592298,sided="two")
```

The block of code below was used to produce the graphs in the first example.

```
# Example 1
                                                                                   diameter<-
c(.25,.25,.251,.25,.252,.253,.252,.255,.259,.261,.249,.250,.250,.250,
  .252)
mu<-.25
sigma<-.0025
y<-(diameter-mu)/sigma
v<-(sqrt(abs(y))-0.822179)/0.3491508
library(qcc)
# Figure 1, se.shift=.5 specifies k=.25 cusum(v,center=0,std.-
dev=1,decision.interval=8.008289,se.shift=.5) library(qcc)
# Figure 2
  cusum(y,center=0,std.dev=1,decision.interval=4.77,se.shift=1)
```

The block of code below was used to produce the graphs in the second example. In this example it is assumed the in-control process mean and standard deviation were known.

```
# Example 2 with simulated data
set.seed(25)
# random data from normal N(50, 7.5)
x<-rnorm(25,50,7.5)
# standardize assuming sigma=5, mu=50
y<-(x-50)/5
# calculate v
v<-(sqrt(abs(y))-0.822179)/0.3491508
library(qcc)
EWMA <- ewma(y, center=0, std.dev=1,lambda=.2, nsigmas=2.86,
plot=FALSE)
EWMA$statistics <- rep(NA,length(EWMA$statistics))
plot(EWMA,ylim=c(-2,2))
EWMA <- ewma(v, center=0, std.dev=1,lambda=.05, nsigmas=2.489686,
plot=FALSE)
EWMA$statistics <- rep(NA,length(EWMA$statistics))
plot(EWMA,ylim=c(-1,1))
```

The arguments `center` and `std.dev` represent the known in-control process mean and standard deviation of y_i . The EWMA parameters $\lambda = .2$ and $L = 2.86$ are defined in the first function call with the arguments `lambda=.2` and `nsigmas = 2.86`. The second EWMA with $\lambda = .05$ and $L = 2.489686$ are defined in the second function call with `lambda=.05` and `nsigmas = 2.489686`.

The block of code below illustrates the use of the `SPCProperty` function in the R package `spcadjust` to produce adjusted control limit multipliers for the two-sided EWMA chart.

```
set.seed(52)
X<-rnorm(100,50,5) set.seed(99)
X<-rnorm(100)
# This is simulated Phase I in-control data, it should be replaced with
actual Phase I data
# if available.
library(spcadjust)
```

```
chart <- new("SPCEWMA", model = SPCModelNormal(Delta=.0), lambda=.05)
xihat <- xiofdata(chart, X)
cal <- SPCproperty(data=X, nrep=1000, property="calARL", chart=chart,
params=list(target=370), quiet=TRUE)
```

In the code above, the argument $\Delta = 0$ is the standard deviation shift in the mean. In this case 0 means you are interested in adjusting the limit multiplier for the ARL_0 . The $\lambda = .05$ is the EWMA smoothing parameter. In the SPCProperty function call, $\text{data} = X$ specifies the vector of Phase I data, $\text{nrep} = 1000$ specifies the recommended number of bootstrap samples, and $\text{target} = 370$ specifies the target value of ARL_0 . The output of the function call is in the list variable cal which contains the adjusted value of $L \times \sqrt{\frac{\lambda}{2-\lambda}}$, from which the adjusted value of L can be determined. For example, running this code resulted in the adjusted value in $\text{cal} = .5248$, and solving for $L = .5248 / (\sqrt{\frac{.05}{2-.05}}) = 3.2774$. This code takes a few minutes to run because it is making 1000 bootstrap estimates based on the Phase I data.

The block of code below repeats the second example without assuming the in-control process mean and standard deviation were known. In this case random data from a normal distribution with mean $\mu = 50$, and standard deviation $\sigma = 5$ were simulated to represent in-control data from a Phase I study.

```
# Create simulated Phase I data
library(qcc)
set.seed(52)
x1<-rnorm(100,50,5)
muhat<-mean(x1)
sighat<-sd(x1)
y1<-(x1-muhat)/sighat
v1<-(sqrt(abs(y1))-.822179)/.3491508
library(spcadjust)
# Calculate the adjusted L for the EWMA chart of y; Note: it takes a few
minutes for this section
# of code to run
charty <- new("SPCEWMA", model = SPCModelNormal(Delta=.0), lambda=.2)
xihat <- xiofdata(charty, y1)
cy <- SPCproperty(data=y1, nrep
=1000, property="calARL", chart=charty, params=list(target=370), quiet =
=TRUE)
Ly<-cy@res/(sqrt(.2/(2-.2)))
# Calculate the adjusted L for the EWMA chart of v; Note: it takes a few
minutes for this section
# of code to run
chartv <- new("SPCEWMA", model = SPCModelNormal(Delta=.0), lambda=.05)
xihat <- xiofdata(chartv, v1)
cv <- SPCproperty(data=y1, nrep
=1000, property="calARL", chart=chartv, params=list(target=370), quiet =
=TRUE)
Lv<-cv@res/(sqrt(.2/(2-.2)))
# Here is the same simulated Phase II data used in the second example in Section
5
set.seed(25)
# random data from normal N(50, 7.5)
x<-rnorm(25,50,7.5)
# standardize using estimated mu and sigma
y<-(x-muhat)/sighat
# calculate v
v<-(sqrt(abs(y))-.822179)/.3491508
# adjusted charts, the value of L=nsigmas for these two charts was obtained
using the
# SPCProperty function
EWMA <- ewma(y, center=0, std.dev=1, lambda=.2, nsigmas=Lv, plot=FALSE)
EWMA$statistics <- rep(NA, length(EWMA$statistics))
plot(EWMA, ylim=c(-2,2))
EWMA <- ewma(v, center=0, std.dev=1, lambda=.05, nsigmas=Lv,
plot=FALSE)
EWMA$statistics <- rep(NA, length(EWMA$statistics))
plot(EWMA, ylim=c(-1,1))
```