Monitoring Processes with Highly Censored Data

Stefan H. Steiner and R. Jock MacKay
Univeristy of Waterloo

RR-98-06
August 1998

Monitoring Processes with Highly Censored Data

Stefan H. Steiner and R. Jock MacKay
Dept. of Statistics and Actuarial Sciences
University of Waterloo
Waterloo, N2L 3G1 Canada

The need for process monitoring in industry is ubiquitous. By monitoring process output, problems may be rapidly detected and corrected. However, in many industrial and medical applications observations are censored either due to inherent limitations or cost/time considerations. For example, when testing breaking strengths or failure times often a limited stress test is performed and only a small proportion of the true failure strengths or failure times are observed. With highly censored observations a direct application of traditional monitoring procedures is not appropriate. In this article, Shewhart type $\overline{X}$ and S control charts based on the conditional expected value weight are suggested for monitoring processes where the censoring occurs at a fixed level. We provide an example to illustrate the application of this methodology.

Keywords: Process Control; Scores; Type I Censoring

1. Introduction

In many industrial applications censored observations are collected for process monitoring purposes. For example, in the manufacture of material for use in the interior trim of an automobile, a vinyl outer layer is glued to an insulating foam backing. The strength of the bond between the layers is an important characteristic. To check the bond strength, a rectangular sample of the material is cut and the force required to break the bond is then measured. A pre-determined maximum force is applied to avoid tearing the foam backing. Most samples do not fail so it is known only that the bond strength exceeds the pre-determined force. That is, the bond strength data are censored. The process is monitored by selecting samples across the width of the material at a given frequency based on the amount of material produced. The purpose of the monitoring is to ensure that the bond strength does not deteriorate. Deterioration includes decreases in the average strength or increased variability. A second example, which we do not consider in more detail here, is the use of plug gauges to monitor hole size. To measure hole diameter, two plugs machined to have diameters at the upper and lower specification of the hole diameter respectively are applied. If the larger plug enters the hole, then the diameter exceeds the upper specification. If the smaller plug does not enter the hole, then the hole size is below the minimum specification. For the purpose of process monitoring, the actual diameter of the few holes that fail are measured. Here all diameters within the specification limits are censored. Similar situations that result in censored data occur in life testing and other areas of application such as medicine. For simplicity we will always refer to the variable of interest as a strength although it may just as well be a lifetime.

In these examples, a direct application of an $\overline{X}$ and S control chart on the observed strength, where we ignore the censoring, is reasonable if the censoring proportion is not large, say less than 50%. On the other hand, when the censoring proportion is very high, say greater than 95%, it is feasible to use a traditional np chart where we record only the number of censored observations. In this article, we propose conditional expected value (CEV) weight control charts appropriate for monitoring processes that produce censored observations. The proposed charts

are superior to traditional methods, especially when the censoring proportion lies between 50-95%.

As with tradition monitoring procedure, such as $\overline{X}$ and S control charts, the CEV weight charts are derived using a two step process. The first stage, often called the initial implementation, involves collecting a setup sample from an in-control process. With $\overline{X}$ and S charts the setup sample usually consists of around 20 subgroups. From the values observed in these initial subgroups the control chart(s) are established. If there is any evidence of instability in the initial sample, i.e. points plotting outside the control limits, the offending subgroups are closely examined and removed if the cause of the instability is determined. If any subgroups are removed, the control limits are re-established.

This article is organized in the following manner. In Section 2 methods for characterizing samples that contain censored observations are discussed. In this article we assume an underlying normal distribution. Section 3 introduces the CEV weight control charting procedure that allows for the rapid detection of deterioration in the process quality when the monitored output is censored. Finally, Section 4 illustrates the use this control charting procedure in the first example described above.

2. Characterizing Samples with Censored Observations

In a sample with observations censored at a fixed level, information about the underlying process is provided by both the observed failure strengths and the number of censored observations. In this situation we should try to use all the provided information to characterize censored samples, and a likelihood based approach seems appropriate. In Section 2.1 the likelihood of samples censored at a fixed level is derived when an underlying normal distribution is assumed for the response of interest.

Using the likelihood function we may derive maximum likelihood estimates (MLEs) for the underlying process mean and standard deviation. Computational methods for determining MLEs are well documented for many distributions, and have well known properties. For the

normal distribution, the derivation of the MLE is illustrated in Section 2.2. Maximum likelihood estimation is useful for the estimation of the process mean and standard deviation from the initial in-control sample. This is further discussed in Section 3.3. However, the MLE approach has a number of drawbacks. First, it works well only when sample sizes are fairly large, since otherwise the sampling variability is large, and there may be a non negligible probability that all observations are censored, and thus that the MLEs do not exist. In most process monitoring applications the sample subgroup sizes are small, n = 5 or 10 is typical, which may cause problems for the MLE approach. Also, the calculation of the MLEs is iterative thus requiring a fairly substantial computational effort that may be onerous on the shop floor.

An alternative idea is to use score weights. The score is defined as the derivative of the log likelihood function evaluated at the parameter values of the stable process. Consideration of the score weights in the normal case is provided in Section 2.3. The score weights are easily calculated, always exist, and are more stable than the MLEs since the calculation of the scores requires assumed process parameters. This is an advantage when monitoring a process for changes using small subgroups. One disadvantage of scores is that they do not always have a physical interpretation. However, for some distributions, such as the normal distribution, it can be shown that the scores are equivalent, under a rescaling, to the expected strengths given the censoring level. The expected strengths are physically interpretable, and are called the conditional expected value (CEV) weights. In Section 2.4, the information content of a censored sample in terms of Fisher information is derived. Fisher information can be used to determine theoretically how much information regarding either the mean or standard deviation is lost due to the censoring. Using the Fisher information we may determine the approximate sample sizes needed to match the power of the corresponding procedure using uncensored samples.

2.1 Likelihood of Samples with Censored Observations

Let T be a random variable whose distribution represents the failure strengths. The density and survivor function of T are $f(t;\theta)$ and $S(t;\theta)$ respectively. The censoring value is C.

We shall assume the censoring is in the right tail, although similar results may be obtained for other censoring patterns. For samples with strength greater than C, the exact strength is not observed. It is well known that the log-likelihood is

$$\log L(\theta) = (n-r)S(C;\theta) + \sum_{i \notin D} f(t_i;\theta)$$

where D represent the set of all observations that were not censored, and r equals the number of uncensored observations. In what follows we assume normality of the actual strength distribution. Other distributional assumptions such as exponential and Weibull are also possible and do not change the procedure markedly. Assuming T ~ $N(\mu, \sigma^2)$, and thus $\theta = (\mu, \sigma)$, the log-likelihood is

$$\log L(\theta) = (n-r)\log\left[Q\left(\frac{C-\mu}{\sigma}\right)\right] + \sum_{i \notin D} \log\left[\phi\left(\frac{t_i-\mu}{\sigma}\right)\right] \qquad (2)$$

where $\phi(z) = e^{-z^2/2}/\sqrt{2\pi}$ and $Q(z) = \int_z^\infty \phi(x)dx$, are the probability density function and the survivor function respectively of the standard normal. Also, the probability of censoring equals

$$p_c = 1 - F(C) = Q((C-\mu)/\sigma) \qquad (3)$$

2.2 Maximum Likelihood Normal Parameter Estimates from Censored Data

Methods for determining maximum likelihood estimates (MLEs) for θ from (2) for many distributions common in the analysis of lifetime data are considered in Lawless (1982). For normal data, a number of approaches have been suggested. We present an iterative approach due to Sampford and Taylor (1959) since it also relates to the CEV weights discussed in Section 2.3. The Sampford and Taylor method is an application of the expectation-maximization (EM) algorithm discussed by Dempster et al. (1977). The procedure is iterative and involves replacing each censored observation with its conditional expected value given the current best guess for the process mean and standard deviation, $\hat{\mu}$ and $\hat{\sigma}$. It can be shown (Lawless, 1982) that assuming

a normal distribution the conditional expected value, evaluated at $\hat{\mu}$, $\hat{\sigma}$, of all censored observations is

$$w_C = E(T \mid T \geq C) = \hat{\mu} + \hat{\sigma}\frac{\phi(z_C)}{Q(z_C)} \tag{4}$$

where $z_C = (C - \hat{\mu})/\hat{\sigma}$. We define the conditional expected value (CEV) weight w of each sample unit as:

$$w = \begin{cases} t & \text{if } t \leq C \text{ (not censored)} \\ w_C & \text{if } t > C \text{ (censored)} \end{cases} \tag{5}$$

Based on the CEV weights given by (5) we may estimate the process mean and variance as:

$$\begin{aligned} \hat{\mu} &= \sum_{i=1}^{n} \frac{w_i}{n}, \\ \hat{\sigma}^2 &= \frac{\sum_{i=1}^{n}(w_i - \hat{\mu})^2}{[r + (n-r)\lambda(z_C)]} \end{aligned} \tag{6}$$

where $\lambda(z) = \frac{\phi(z)}{Q(z)}\left[\frac{\phi(z)}{Q(z)} - z\right]$. Note that $\lambda(z)$ always lies between 0 and 1. $\lambda(z)$ is near 1 when the censoring proportion is small, and near 0 when the censoring proportion is large. As a result, the term $r + (n-r)\lambda(z_C)$ can be thought of as a sample size adjusted for the number of censored observations.

To find the MLE values, we iteratively apply equations (4)-(6) to the data until the values for $\hat{\mu}$ and $\hat{\sigma}$ converge. The iterations rapidly converge (less than 10 iterations) to the MLEs so long as good initial values are employed. In most cases, good initial values are the sample mean and sample standard deviation obtained when ignoring the censoring.

Generally, maximum likelihood estimates work well for large samples, and are thus ideal for process capability studies (Kotz, 1993), and the initial implementation of a control chart where large samples are required. For ongoing monitoring when there is no censoring $\overline{X}$, S charts use estimates of μ and σ from small samples to detect changes in the process. With high censoring proportions the MLEs of μ and σ are poorly behaved. In the extreme case with all

observations censored unique MLE's do not exist. If the censoring proportion p_c, given by (3), is greater than 50% and the subgroup size n is small, the probability of complete censoring p_c^n is non-negligible. The issue of the sample sizes required to estimate the process mean and standard deviation based on censored samples with various censoring proportions is discussed further in Sections 2.4 and 3.3.

2.3 Score and Conditional Expected Value Weights

An alternative approach that works better than MLEs for small subgroup sizes is a score based approach. Since typically the process mean is of primary concern we work with mean or location scores. In the normal case, the mean score, denoted m, is defined as the first derivative of the log likelihood given by (2) with respect to μ evaluated at the process mean and standard deviation estimated from the stable process.

$$m = \left.\frac{\partial \log L}{\partial \mu}\right|_{\substack{\mu=\mu_0 \\ \sigma=\sigma_0}} = \begin{cases} \dfrac{t-\mu_0}{\sigma_0^2} & \text{if } t \le C \text{ (not censored)} \\ \dfrac{\phi\left((C-\mu_0)/\sigma_0\right)}{\sigma_0 Q\left((C-\mu_0)/\sigma_0\right)} & \text{if } t > C \text{ (censored)} \end{cases} . \qquad (7)$$

Comparing the mean score (7) with the CEV weights for the normal distribution given by (5) shows that w_i equals $\mu_0 + \sigma_0^2 m$ for both censored and uncensored observations. Thus, for normal data, the mean scores are a linear translation of the conditional expected value weights. Similar relations between CEV weights and scores exist for other distributions. For example, for the exponential distribution $w = \theta^2 m + 1/\theta$, where θ is the mean.

Note that when calculating CEV weights estimates of the current process mean and standard deviation values μ_0 and σ_0 are required. In applications, these values are estimated from the process data used in the initial implementation of the monitoring procedure. The estimation of the mean and standard deviation may be difficult when the process data are highly censored. The question of initial estimation is covered in more depth in Section 3.3.

Note that if $\hat{\mu} = \mu_0$ and $\hat{\sigma} = \sigma_0$, i.e. the current process parameters equal the estimates, and the process is in-control then the expected CEV weight equals μ_0. This follows because the expected score equals zero. However, when $\hat{\mu} = \mu_0$ and $\hat{\sigma} = \sigma_0$, the standard deviation of the CEV weights is less than the standard deviation of the underlying process. As the censoring proportion increases the standard deviation of the CEV weights decreases. Note, as well, that the distribution of the CEV weights is highly skewed when the censoring proportion is large.

2.4 Expected Information in Censored Samples

Censored samples and uncensored samples may be compared using statistical (Fisher) information. The inverse of the Fisher information gives the asymptotic variance of the maximum likelihood parameter estimates.

Fisher information is defined as minus the second derivative of the log likelihood function. For purposes of investigation, and without loss of generality, we shall assume that the mean and standard deviation of the failure strength assumed in the derivation are $\mu_0 = 0$ and $\sigma_0 = 1$. In the censored normal case we can derive the information matrix I from the log-likelihood expression (2). We get

$$I = \begin{bmatrix} E\left(-\frac{\partial^2 \log L}{\partial \mu^2}\right), & E\left(-\frac{\partial^2 \log L}{\partial \mu \partial \sigma}\right) \\ E\left(-\frac{\partial^2 \log L}{\partial \mu \partial \sigma}\right), & E\left(-\frac{\partial^2 \log L}{\partial \sigma^2}\right) \end{bmatrix} \tag{8}$$

where
$$E\left(-\frac{\partial^2 \log L}{\partial \mu^2}\right) = n(1-p_c) + np_c\lambda(C)$$
$$E\left(-\frac{\partial^2 \log L}{\partial \sigma^2}\right) = n(1-p_c)\left[3E\left(t^2 \mid t < C\right) - 1\right] + np_c\left[2C\phi(C)/Q(C) + C^2\lambda(C)\right]$$
$$E\left(-\frac{\partial^2 \log L}{\partial \mu \partial \sigma}\right) = 2n(1-p_c)E\left(t \mid t < C\right) + np_c\left[\phi(C)/Q(C) + C\lambda(C)\right]$$

and
$$E\left(t \mid t < C\right) = -\phi(C)/(1-Q(C))$$
$$E\left(t^2 \mid t < C\right) = -C\phi(C) + 1 - Q(C)$$

If there is no censoring, i.e. C = ∞ then $p_c = 0$, $I = \begin{bmatrix} n & 0 \\ 0 & 2n \end{bmatrix}$ and $I^{-1} = \begin{bmatrix} 1/n & 0 \\ 0 & 1/2n \end{bmatrix}$. Thus, the asymptotic variance of the maximum likelihood mean and standard deviation estimates are $1/n$ and $1/2n$ respectively. This asymptotic variance of the MLEs for uncensored samples may compared with the diagonal elements of the I^{-1} matrix for samples censored at different levels of C.

The increased sampling variability in estimation of the mean and standard deviation can be substantial when the censoring proportion is large. Figure 1 shows the sampling size required to match the sampling variability in the uncensored case for the mean and standard deviation. Note that for small censoring proportions we can estimate the mean quite well. However, when the censoring proportion increases it becomes increasingly difficult to estimate the process mean and standard deviation. Also, our ability to estimate the process mean degrades more quickly than our ability to estimate the process standard deviation as the proportion censored increases.

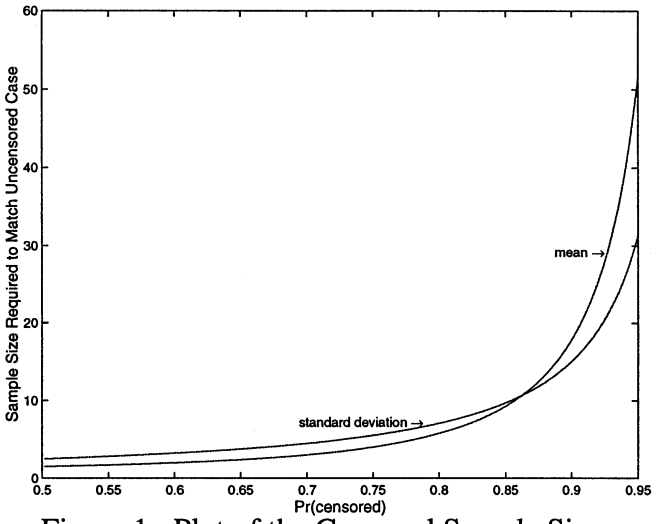

Figure 1: Plot of the Censored Sample Sizes Needed to Match an Uncensored Sample

For example, with 50% censoring we need only 1.5 and 2.5 times the uncensored sample size to estimate the mean and standard deviation respectively as well as in the uncensored case.

However, when the censoring rate is 95% the required sample size multiples are 51 and 31 respectively.

3. CEV Weight Control Chart for Censored Data

In this section the CEV weight control charts are derived for detecting mean and dispersion shifts in the process. Using a CEV weight control chart each censored observation is replaced with its condition expected value (CEV) weight, and the subgroup averages and sample standard deviations are plotted in a manner similar to the traditional $\overline{X}$ and S charts. With right censored data the goal of the CEV weight control chart is to detect decreases in the process mean and/or increases in the process standard deviation.

The reason for this relatively limited goal is that when the censoring proportion is large the sample provides only limited information regarding possible changes in the process parameters. In particular, it is very difficult to detect increases in the process mean or decreases in the process standard deviation. This is because if the process mean shifts upward we expect more censored observations. Similarly, if the censoring proportion is greater than 50%, decreases in the process dispersion also lead to more censored observations. Unfortunately samples with all or almost all censored observations provide very little information about the process parameters. As a result, it is not reasonable to expect that when the in-control process produces many censored observations that we will be able to detect process changes that result in even more censored observations without extremely large subgroup sizes. On the other hand, it is feasible to detect decreases in the process mean and/or increases in the process standard deviation surprisingly well. This is because decreases in the process mean or increases in the process dispersion lead to decreases in the censoring proportion which in turn means that each sample provides more process information. Fortunately, in most situations where we obtain right censored values decreases in the mean and increases in the dispersion are the types of process changes we are most concerned with since they represent a degradation of the process performance. As a result the two proposed CEV control charts are both one-sided. The CEV $\overline{X}$

chart has only a lower control limit, while the CEV standard deviation (S) chart has only an upper control limit. Note that for both charts using a centre line is not of much value since the distributions of the sample average and sample standard deviation of the CEV weights are highly skewed.

This section of the article is organized in the following manner. Subsection 3.1 gives a procedure for the initial implementation of the control charts, including the initial estimation of the process parameters. In Subsection 3.2 the determination of the appropriate control limits for the CEV control charts is discussed, and figures are given that can be used to choose the control limits for a variety of sample sizes and censoring proportions. Subsection 3.3 shows the power of the CEV weight control charts for detecting certain type of process changes. Based on these results it is evident that when the censoring proportion is very large the $\overline{X}$ type chart only suffices to detect both mean and standard deviation shifts in the process. Finally, in Subsection 3.4 the performance of the CEV control chart is compared with more traditional control charts like an np chart based on the number of censored observations, and a Shewhart $\overline{X}$ chart based on the observed data where censoring is ignored.

3.1 Initial Implementation

In process monitoring, a distinction is made between the initial implementation phase and the on-going control phase of process monitoring. More emphasis is typically given to the ongoing control phase where we assume the process mean and variance are fixed and known. During the initial implementation phase the in-control process mean and variance are estimated and appropriate control limits are determined. In addition the assumption that the initial data comes from an in-control process is verified by examining the resulting control chart of the initial data. When working with uncensored data guidelines suggest that a minimum of 100 observations (often 20 subgroups of size 5) is required for the initial implementation of $\overline{X}$ and S charts. This sample size restriction ensures that the initial process parameter estimates are estimated reasonably accurately and that any estimation errors can be ignored.

Based on the results of Section 2.4 is evident that with high levels of censoring it is difficult to estimate the process mean and standard deviation unless the overall sample size is large. As a result, during the initial implementation step of control charting some changes to the traditional procedure must be made. For example, to obtain more process information we may either increase the total sample size used, or decrease the censoring proportion (at least temporarily for the initial implementation sample). Decreasing the censoring proportion is usually feasible, though perhaps expensive, since often the level of censoring is under our control. Figure 1 can be used as a guide to the sampling requirements. Note that to determine the required sample size from Figure 1 we must know the censoring proportion à priori. Likely we will not have accurate information on the censoring proportion, but typically we will have some idea. If very little prior information is available we may adapt the initial data collection procedure if the censoring proportion observed is much larger or smaller than anticipated.

If out-of-control subgroups are detecting in the initial sample the traditionally accepted practice is to explore the offending subgroups and if reasons for an out-of-control condition can be established the subgroup is removed and the initial implementation procedure is redone. The following step by step algorithm illustrates this practice for CEV $\overline{X}$ and S charts.

General Procedure for the Initial Implementation of the CEV Control Charts

1. Collect q subgroups of size n , where the total sample size and the censoring proportion are chosen so that the sampling variability of the process parameters is reasonable. Figure 1 can be used as a guide
2. Estimate the mean and standard deviation for all qn units using MLEs given by EM procedure (4)-(6).
3. Determine the censored CEV weight w_c using (4) based on the mean and standard deviation estimates derived in step 2 and replace all censored observations in the sample with w_c.
4. Create one-sided CEV $\overline{X}$ and S charts plotting the subgroup averages and standard deviations with control limits determined using the design figures in Section 3.2.

5. Look for any out-of-control signals (points outside the control limits) on the charts. Examine process conditions at the time these subgroups were collected. Repeat the procedure from step 2 if any out-of-control subgroups are removed from the sample.

The procedure described above is relatively robust to imprecise initial estimation of the process mean and standard deviation. In some sense, the CEV chart design procedure is self correcting since if the process mean is underestimated, the resulting control limit on the CEV $\overline{X}$ chart will be lower, but the CEV weight assigned to all censored observations is also lower.

In the derivation presented we suggested the use of mean scores. However, if changes in the process standard deviation are of great concern, it would be possible to use standard deviation scores when deriving the CEV S chart.

3.2 CEV Control Chart Design

An important question related to CEV weight control charts is how to choose appropriate control limits. The position of the appropriate control limits depend on both the sample size and the in-control probability censored. Due to effect of different degrees of censoring there is no generally applicable formula, such as the traditional plus or minus three standard deviation limits, that gives the appropriate control limits for CEV weight control charts. However, Figures 2 and 3 are provided to aid in the choice of control limits for the CEV $\overline{X}$ and S control charts. The figures are based on the assumption that the in-control proportion censored is known, and are derived through simulation with a million trials to estimate each point. Figure 2 gives the standardized lower control for the CEV $\overline{X}$ chart that has a theoretical false alarm rate of .0027. This particular false alarm rate was chosen to match the false alarm rate aimed for with the traditional Shewhart $\overline{X}$ control chart. Similarly Figure 3 gives the standardized upper control limit for a CEV S chart that yields a false alarm rate of .0027. Note that the horizontal axes in both Figures 2 and 3 are on a log scale.

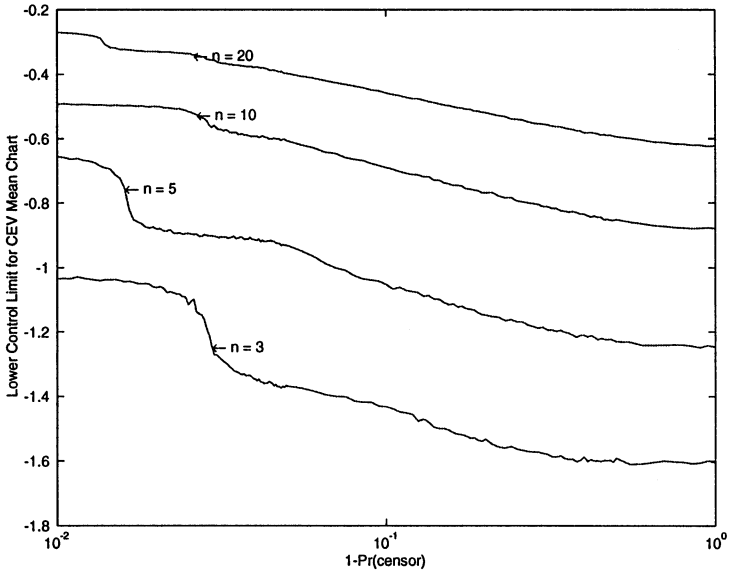

Figure 2: Plot of the Standardized Lower Control Limit for the CEV $\overline{X}$ chart

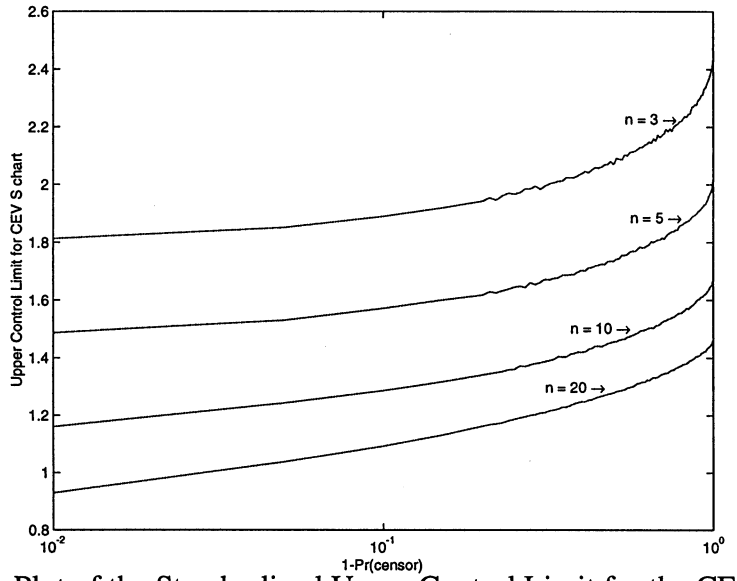

Figure 3: Plot of the Standardized Upper Control Limit for the CEV S chart

For sample sizes between the given values interpolation between the curves on the plot can be used. For example, when designing a CEV S chart with a sample size of 8, and an in-control probability of censoring equal to .9 using Figure 3 we choose a standardized upper control of 1.3. The irregular parts of Figure 2 are due to the discreteness inherent in the problem.

The control limits given in Figures 2 and 3 are standardized in the sense that they give the appropriate control limits given the sample size, the in-control probability of censoring, and

assuming the in-control process has mean zero and variance one. The control limits appropriate in any given example problem may be obtained using (9) where μ_0 and σ_0 are the actual in-control process parameters, and $lcl_{\bar{X}}$ and ucl_S are the standardized control limits given by Figures 2 and 3 respectively.

$$\begin{aligned} \text{lower control limit for CEV } \bar{X} \text{ chart} &= lcl_{\bar{X}}\sigma_0 + \mu_0 \\ \text{upper control limit for CEV S chart} &= ucl_S\sigma_0 \end{aligned} \tag{9}$$

3.3 CEV Control Chart Performance

It is of interest to compare the power of the CEV $\bar{X}$ and CEV S charts in detecting various types of process changes with different levels of in-control censoring. The following figures for subgroups of size 5 or 10 are based on control limits determined from the design figures in Section 3.2, and thus the false alarm rate of all the charts is set at .0027. The results presented here are based on simulation using 200,000 trials for each point. For comparison purposes the performance in the uncensored case is given with a dashed line in each plot. For each chart the horizontal axis corresponds to shifts in units of the standard deviation.

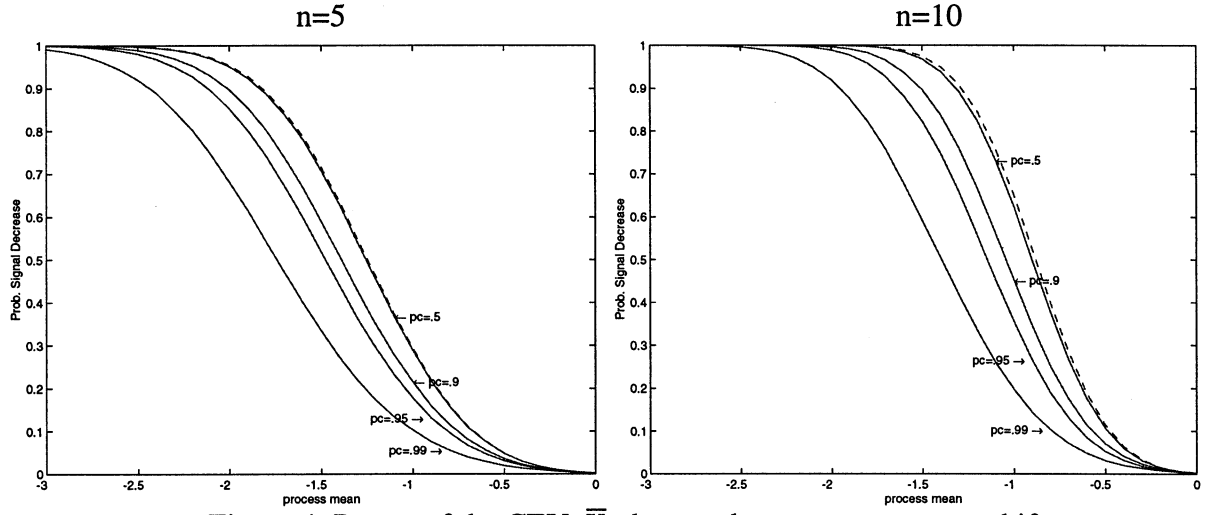

Figure 4: Power of the CEV $\bar{X}$ chart to detect process mean shifts
no censoring case given by the dashed line

n=5 n=10

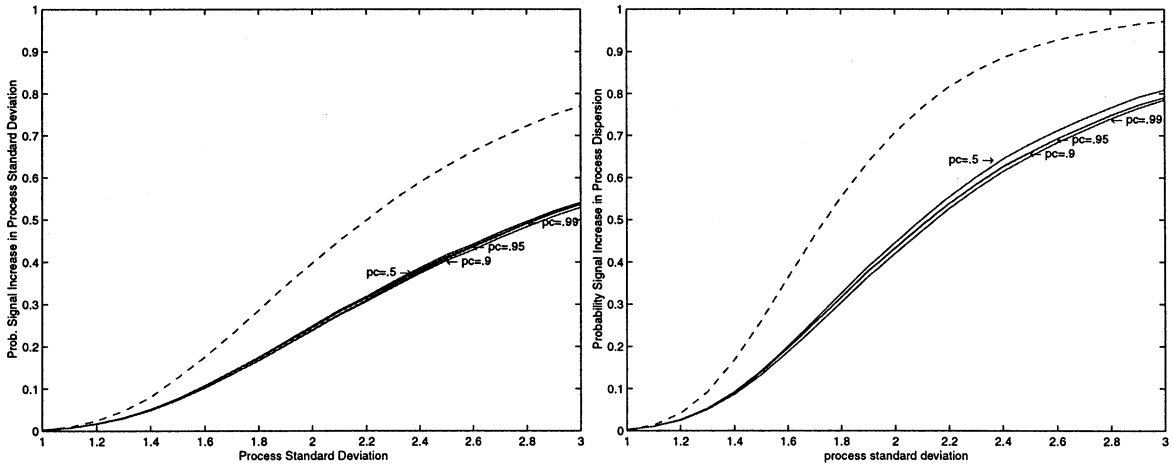

Figure 5: Power of the CEV S chart to detect process standard deviation shifts
no censoring case given by the dashed line

In the CEV S chart the power loss that results from using censored observations is fairly large for virtually any level of censoring. However, for censoring proportions between .5 and .99 the difference in power to detect process standard deviation shifts is small. This is because large increases in the process variability will result in some large negative values that will be observed even with a large amount of censoring. In the CEV $\overline{X}$ chart, on the other hand, the decrease in power as the censoring proportion increases is more gradual. For moderate censoring proportions, such as 50% censoring, there is almost no loss in power to detect process mean decreases.

Clearly, based on these results, there is a tradeoff between information content of the subgroup and the data collection costs. In many applications the censoring proportion is under our control through the censoring level C. Setting it so that there are few censored observations provides the most information, but will usually also be the most expensive. The optimal tradeoff point depends on the sampling costs and the consequences of false alarms and/or missing process changes.

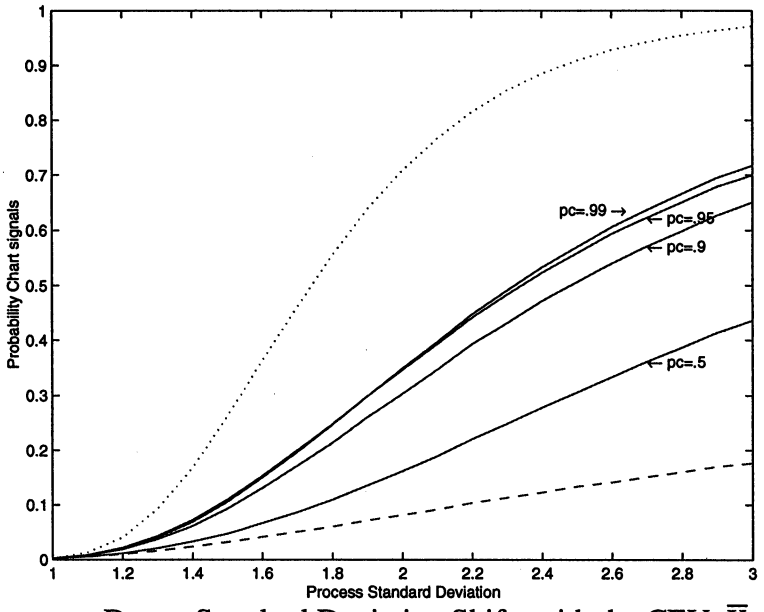

Figure 6: Power to Detect Standard Deviation Shifts with the CEV $\overline{X}$ chart, n = 10
$\overline{X}$ with no censoring chart given by the dashed line
S chart with no censoring given by the dotted line

The CEV $\overline{X}$ is also good at detecting changes in the process standard deviation. This is illustrated in Figure 6 for subgroups of size 10. Note that the detection of standard deviation shifts works only when the censoring proportion is large. This is because when the proportion censored is very large, say greater than 95% censoring, it is difficult to distinguish between increases in the process mean and decreases in the process variability. For highly censored data, increases in the process variability appears similar to decreases in the process mean since due to the censoring the large positive values are replaced by the CEV weight and this do not appear large. On the other hand, when there is no censoring, the large observations will be observed and tend to cancel the influence of the small observations in the calculation of the sample mean. As a result, when the in-control proportion censored is very large the process can be adequately monitored using only the CEV $\overline{X}$ chart.

3.4 Comparison of CEV Control Chart Performance to Traditional Charts

As a further comparison we may consider the use of a traditional control charts like the np chart for the number censored in each sample, and a Shewhart $\overline{X}$ chart of the data where we

ignore the censoring. A direct comparison between an np chart and the CEV $\overline{X}$ and S charts is difficult since changes in the proportion censored may be due to either changes in the process mean or the process dispersion or both. In addition, due to discreteness, the np chart can not necessarily be setup to have a particular false alarm rate. This is illustrated in Table 2 that gives the decision rules and corresponding probability of a false alarm for np charts that yields false alarm rates as close to .0027 as possible.

Table 2: np Chart Decision Rule when n = 5: if number censored < x then signal

p_c	x	Pr(false alarm)
.5	1	.031
.75	1	.001
.9	3	.009
.95	3	.001
.99	4	.001

The performance of the np chart in detecting decreases in the censoring proportion (caused by decreases in the process mean) is quite similar to the CEV $\overline{X}$ chart when p_c is very large. This is not surprising since when the censoring proportion is very large little additional information is available in knowing the few actually observed non-censored values. Figure 7 compares np charts and CEV $\overline{X}$ charts when the changes in the censoring proportion are due exclusively to mean shifts for in-control censoring proportions equal to .5 and .9. Figure 7 suggests that as the censoring proportion increases the performance of the two chart becomes more similar. In Figure 7 the control limit of the CEV $\overline{X}$ chart has been adjusted so that it yields approximately the same in-control false alarm rate as the np chart. Note that the power curves for the two different proportion censored are not directly comparable since they have different false alarm rates.

As discussed the ability of np charts to detect decreases in the process mean is comparable to the CEV $\overline{X}$ chart when the in-control proportion censored is large. However, when the changes in the proportion censored are due to increasing dispersion the np chart does not do as well as the CEV S chart. This is clearly evident, for example if the censoring

proportion is 50%, then increases in the process dispersion do not lead to changes in the proportion censored. In general, the np chart will perform poorly if the process changes do not lead to large decreases in the proportion censored since the np chart can not distinguish between changes to the process mean and standard deviation.

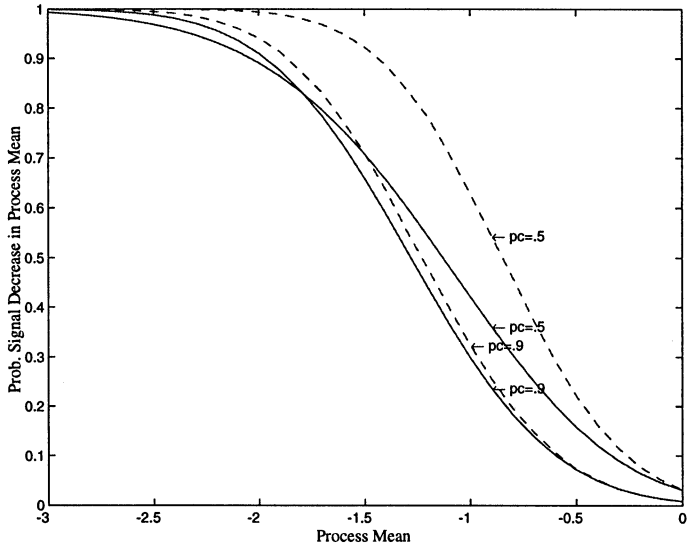

Figure 7: Comparison of Performance Between CEV $\overline{X}$ and np charts, n = 5
CEV $\overline{X}$ chart given by dashed line

The comparison between the CEV $\overline{X}$ chart and the traditional Shewhart $\overline{X}$ chart is also difficult since a naive application of an $\overline{X}$ chart would ignore the censoring and set a lower control limit at $\overline{\overline{X}} - 3\hat{\sigma}/\sqrt{n}$ where the standard deviation estimate is given by either $\overline{s}/c_4$, or $\overline{R}/d_2$, where $\overline{s}$ and $\overline{R}$ are the average subgroup standard deviation and average subgroup range respectively, and c_4 and d_2 are control chart constants (Ryan, 1989). By ignoring the censoring it is meant that the censored values are used as if they are actual observed failure strengths. This naive $\overline{X}$ chart would ignore the skewness of the observations introduced by the censoring, and thus would likely not have the desired false alarm rate. For example, assuming 90% censoring the naive method would yield an $\overline{X}$ chart with almost a 10% chance of signaling when the process is in-control. This is clearly unacceptable. However, using a procedure similar to that presented in Section 3.2 for the CEV charts we may derive a lower control limit for the Shewhart $\overline{X}$ chart where censoring is ignored that gives the desired false alarm rate. Figure 8 shows a

comparison between the power of the CEV $\overline{X}$ chart and the naive Shewhart $\overline{X}$ chart with adjusted control limits. The figure shows that the for highly censored data the CEV $\overline{X}$ chart has superior performance, substantially so for very high censoring rates. Note also that the CEV $\overline{X}$ chart is preferable to the naive Shewhart $\overline{X}$ chart because with the CEV chart the sample average can be interpreted as an estimate of the process mean.

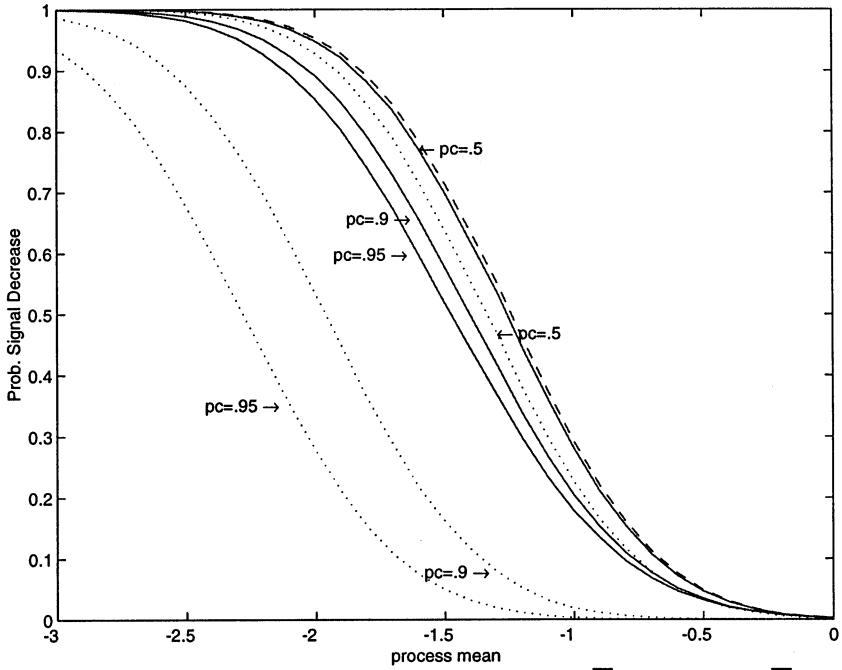

Figure 8: Comparison of Performance Between CEV $\overline{X}$ and Naive $\overline{X}$ charts, n = 5
no censoring $\overline{X}$ given by dashed line, CEV charts by solid lines, and the naive $\overline{X}$ by dotted lines

4. Example

In the glue bond strength example described in the introduction, an initial sample of 100 subgroups of size 5 was selected from historical monitoring records. The censoring point C had been set at the specification limit, here coded at 10 units. This was well below the tearing strength of the foam. No charting had been undertaken. When out-of-specification bond strengths were detected, the process was investigated but typically no action was taken.

In the data, the first 100 observations of which is given in the Appendix, there was a 86% censoring rate. The high proportion of out-of-specification readings was the motivation for the implementation of the charting procedure. Using the MLE procedure given by (4)-(6) we

estimate the process mean and standard deviation as μ = 11.1 and σ = 1.24. With a censoring level of 10, from (4) we get w_c = 11.44. This is the weight assigned to all censored observations in the CEV monitoring procedure. Based on subgroups of size 5 and a 86% censoring rate the standardized control limits for the $\overline{X}$ and S charts are –1.13 and 1.62 respectively. These values may be determined approximately either from Figures 2 and 3. Scaling the control limits by the estimated mean and standard deviation according to (9) gives a lower control limit of 9.7 for the CEV $\overline{X}$ chart, and an upper control limit of 2.02 for the CEV S chart. The resulting CEV $\overline{X}$ and S charts for the example data are given by Figure 9.

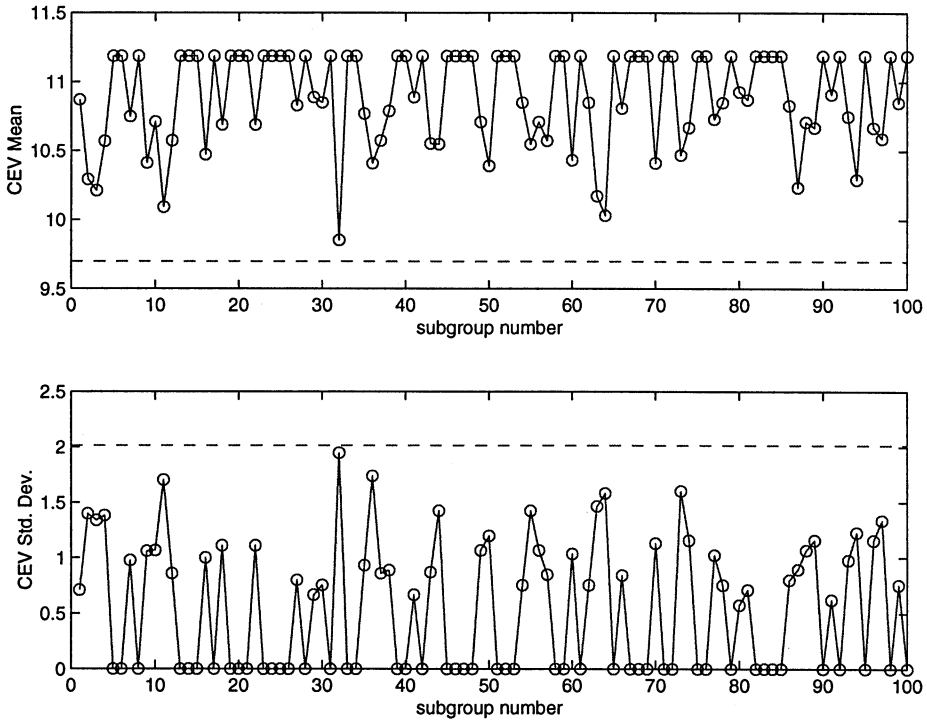

Figure 9: Example CEV $\overline{X}$ and S charts with n = 5

Figure 9 shows that in the initial implementation there were no out-of-control points. Thus, the initial data appears to come from an in-control process, and we should have obtained reasonably accurate estimates of the process mean and standard deviation. As a result, we may continue to monitor the process for deterioration using the CEV charts with the given control

limits. To reduce the out-of-specification rate from around 14% the common cause of variation must be addressed.

5. Summary and Conclusions

In applications where observed data may be censored, traditional process monitoring approaches, such as $\overline{X}$ and R charts, have undesirable properties such as large false alarm rates or low power. In this article, adapted control charting procedures to monitor the process mean and standard deviation applicable when observations are censored at a fixed levels are proposed. The proposed charts are based on the idea of replacing all observations by their conditional expected values (CEV) weights. The CEV weights are equivalent to the scores if the underlying distribution is normal. The monitoring procedure is derived assuming the process has an underlying normal distribution, but the same methodology is applicable to other distributions. The procedure is illustrated with an example from the automotive industry. Further articles will address other censoring schemes and the more complicated situation of competing risks.

Acknowledgments

This research was supported, in part, by the Natural Sciences and Engineering Research Council of Canada, and General Motors of Canada. MATLAB® is a register trademark of the MathWorks.

Appendix

Table A1: First 100 Observations of the Example Data

Subgroup #	Observations					Subgroup #	Observations				
1	10	9.6	10	10	10	14	10	10	10	10	10
2	10	8	10	10	9.9	15	10	10	10	10	10
3	10	8.9	8.6	10	10	16	10	9.7	9.1	10	10
4	8.1	10	10	10	10	17	10	10	10	10	10
5	10	10	10	10	10	18	10	10	10	8.7	10
6	10	10	10	10	10	19	10	10	10	10	10
7	10	9	10	10	10	20	10	10	10	10	10
8	10	10	10	10	10	21	10	10	10	10	10
9	9.3	10	9.2	10	10	22	8.7	10	10	10	10
10	10	10	10	8.8	10	23	10	10	10	10	10
11	7.3	10	9.6	10	10	24	10	10	10	10	10
12	10	10	10	9.4	9.9	25	10	10	10	10	10
13	10	10	10	10	10						

References

Dempster, A.P., Laird, N.M., Rubin, D.B., (1977), "Maximum Likelihood from Incomplete Data via the EM Algorithm" (with discussion), *Journal of the Royal Statistical Society, Series B,* Vol. 39, pp. 1-38.

Lawless, J.F. (1982), *Statistical Models and Methods for Lifetime Data*, John Wiley and Sons, New York.

Kotz, S. (1993), "Process Capability Indices," Chapman and Hall, New York.

Montgomery, D.C. (1991), *Introduction to Statistical Quality Control*, Second Edition, John Wiley and Sons, New York.

Ryan, T.P. (1989), *Statistical Methods for Quality Improvement*, John Wiley and Sons, New York.

Sampford, M.R., Taylor J. (1959), "Censored Observations in Randomized Block Experiments," *Journal of the Royal Statistical Society, series B*, 21, 214-257.

Steiner, S.H., Geyer, P.L., Wesolowsky, G.O. (1996), "Grouped Data Sequential Probability Ratio Tests and Cumulative Sum Control Charts," *Technometrics,* Vol. 38, 230-237.

Woodall, W. H. (1986), "The design of CUSUM quality control charts," *Journal of Quality Technology*, 18, 99- 102.