Detecting Changes in the Mean from Censored Lifetime Data

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

RR-98-07
August 1998

Detecting Changes in the Mean from Censored Lifetime Data

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

In many industrial and medical applications observations are censored either due to inherent limitations or cost/time considerations. For example, with many products their lifetimes are sufficiently long that it is infeasible to test all products until failure even using accelerated testing. As a result, often a limited stress test is performed and only a proportion of the true failure times are observed. In such situations, it may be desirable to monitor the process quality using repeated lifetesting on samples of the process output. However, with highly censored observations a direct application of traditional monitoring procedures is not appropriate. In this article, Shewhart type control charts based on the conditional expected value weight are developed for monitoring processes where the censoring occurs at a fixed level. An example is provided to illustrate the application of this methodology.

Keywords: Accelerated Testing; Extreme Value Distribution; Process Control; Scores; Type I Censoring; Weibull Distribution

1. Introduction

In industrial applications censored observations are collected for process monitoring purposes. As an example, consider the manufacture of metal electrical boxes used to cover transformers in residential neighbourhoods. These boxes are outside and are thus subjected to ravages of nature, and it is undesirable and dangerous if they rust. During the production process the boxes are painted to inhibit rust formation. It is desired to monitor the painting process to ensure that the rust resistant capabilities of the paint process does not deteriorate. Of course, even under poor regular conditions the boxes do not rust quickly and it would typically take years at a minimum before rust is visible. To speed up the rusting process an accelerated salt fog endurance test is performed. From each test box a panel is cut, scratched in a prescribed manner, and then put in a 30° Celsius salt spray chamber. Test units in the chamber are checked daily for rust, and stay in the chamber until either rust appears or a maximum of 20 days. The maximum is necessary to limit the cost of testing, and to allow more rapid feedback. In this way, censored observations of the time to rust under the accelerated test conditions are obtained. We may monitor the paint process by analyzing the observed censored data. Since our goal to detect deterioration in process, the accelerated nature of the test may be ignored so long as the testing is performed in the same manner each time, assuming the relationship between accelerated and normal conditions does not change.

Similar situations that result in censored life time data occur in medical applications. In addition, analogous situations may occur with failure strength data. For simplicity, in this article we refer to the variable of interest as a lifetime although it may just as well be a strength. We assume that the distribution of the failure times can be modeled with a Weibull distribution. We consider the Weibull distribution because it is very flexible and a popular choice in life testing (Nelson, 1990).

Monitoring the process quality in situations where there is a substantial amount of censoring can be challenging. However, if the censoring proportion is not large, say less than 25%, monitoring the process quality using a direct application of an $\overline{X}$ control chart with probability limits on the observed lifetime is reasonable. Also, when the censoring proportion is very high, say greater than 95%, it is feasible to use a traditional np chart where only the number of censored observations is recorded (Montgomery, 1991). In this article, we explore a number of possible monitoring procedures that are superior to the either of these methods in detecting mean shifts when the censoring proportion lies between 25-95%. These proposed methods take into account both the censoring and the fact that the underlying process output distribution is Weibull. The proposed charts include two likelihood based approached and three approaches based on the conditional expected value.

To set notation, let W be a random variable whose distribution represents the failure times. Assuming a Weibull distribution, the probability density and survivor function of W, $f(w;\alpha,\beta)$ and $S(w;\alpha,\beta)$ respectively, are given by (1), where α and β are the scale and shape parameters of the Weibull respectively.

$$\begin{aligned} f(w;\alpha,\beta) &= \frac{\beta}{\alpha^\beta} w^{\beta-1} \exp\left(-\left[\frac{w}{\alpha}\right]^\beta\right) \\ S(w;\alpha,\beta) &= \exp\left(-\left[\frac{w}{\alpha}\right]^\beta\right) \end{aligned} \tag{1}$$

With this parameterization of the Weibull, the mean and variance of W, μ and σ^2, are given as

$$\begin{aligned} \mu &= \alpha\Gamma\left(1+\frac{1}{\beta}\right) \\ \sigma^2 &= \alpha^2\left\{\Gamma\left(1+\frac{2}{\beta}\right) - \Gamma^2\left(1+\frac{1}{\beta}\right)\right\} \end{aligned}$$

where $\Gamma(x)$ is the Gamma function (Abramowitz and Stegun, 1972). The standard form of the Weibull is given by $W^* = W/\alpha$ and has probability density function $f(w;1,\beta)$. Note that the standardized Weibull still depends on β. The Weibull distribution is commonly used to model lifetimes because it is very flexible. The shape parameter, β, determines whether the failure rate increases with time (β>1) or decreases with time (β<1). When β = 1 the distribution of W is exponential.

In this article, we assume the censoring is in the right tail as is typical for lifetime data. However, similar monitoring procedures may be derived for other censoring patterns. Denote the fixed censoring time as C. Then, from (1), the probability of censoring, denoted p_c, equals

$$p_c = \exp\left(-(C/\alpha)^{\beta}\right) \qquad (2)$$

It is well known (Lawless, 1982) that the log-likelihood of a sample of n observations from a Weibull distribution censored at the fixed level C is

$$\begin{aligned} \log L_W(\alpha,\beta) &= (n-r)\log S(C;\alpha,\beta) + \sum_{i \notin D} \log f(w_i;\alpha,\beta) \qquad (3) \\ &= (n-r)\left[-\left(\frac{C}{\alpha}\right)^{\beta}\right] + \sum_{i \notin D}\left\{\log(\beta) - \log(\alpha) + (\beta-1)\left[\log(w_i) - \log(\alpha)\right] - \left(\frac{w_i}{\alpha}\right)^{\beta}\right\} \end{aligned}$$

where w_i equals an observed failure time, D represent the set of all observations that were not censored, n equals the sample size, and r equals the number of uncensored observations.

We shall design monitoring procedures to detect decreases in the mean lifetime caused by decreases in α, since process changes that cause changes to the shape parameter β seem less likely. Notice that this implies that decreases in the mean lifetime are accompanied by decreases in the process variability. In addition, with a high censoring rate, detecting decreases in the mean lifetime is possible, but increases will be very difficult to detect, since such changes would lead to more censoring. As a result, the proposed monitoring procedures are all one-sided.

It is well known that the Weibull distribution is related to the extreme value and exponential distributions through simple transformations. Since these other distributions will be used later we introduce them now. Letting $V = \beta\log(W/\alpha)$, V has a standardized minimum extreme value distribution with location and scale parameters η and ω equal to zero and unity respectively. The probability density and survivor function of the general extreme value distribution are

$$\begin{aligned} f(v;\eta,\omega) &= \frac{1}{\omega}e^{(v-\eta)/\omega}\exp\left(-e^{(v-\eta)/\omega}\right) \\ S(v;\eta,\omega) &= \exp\left(-e^{(v-\eta)/\omega}\right) \qquad (4) \end{aligned}$$

The mean and variance of V are $\eta + \gamma\omega$ and $\omega^2\pi^2/6$ respectively, where $\gamma = -.577$, the derivative of the Gamma function evaluated at unity, or minus Euler's constant.

Letting $X = (W/\alpha)^\beta$, X has a standardized exponential distribution with mean equal to unity. The probability density and survivor function of the general exponential distribution are given by (5). The mean and variance of X are θ and θ^2 respectively.

$$\begin{aligned} f(x;\theta) &= \exp(-x/\theta)/\theta \\ S(x;\theta) &= \exp(-x/\theta) \end{aligned} \tag{5}$$

This article is organized in the following manner. In Section 2, four distinct methods for monitoring the process mean from samples that contain censored observations from an underlying Weibull distribution for failure times are discussed. The methods compared include maximum likelihood and score based approaches, and two methods based on conditional expected value (CEV) weights. For each method the appropriate test statistic given in terms of the original Weibull scale is given. In addition, the determination of an appropriate level for the control limit in each case is discussed, and design figures are given where possible. The control limits are set to give a false alarm rate of .0027. The choice of .0027 comes from the false alarm rate of the standard Shewhart $\overline{X}$ chart. Section 3 compares the monitoring procedures in terms of their power to detect decreases in the process mean when the monitored output is censored at different rates. The best choice of procedure depends on the application, but we recommend the CEV extreme value approach as the best compromise choice. Finally, Section 4 illustrates the use the proposed CEV extreme value weight control charting procedure in the example described above.

2. Monitoring Procedures

The goal of the monitoring procedure is to quickly detect decreases in mean lifetime. To accomplish this the first step is to estimate the in-control process performance to provide a benchmark for comparison. In this initial step, called the implementation phase, the process parameters are estimated from a fairly large initial sample of the process output, a control chart is built, and we verify that the initial sample comes from an in-control process (Montgomery, 1991). With an initial censored Weibull sample, we may derive maximum likelihood estimates (MLEs) for the underlying Weibull parameters using the likelihood function given by (3). Computational methods for determining the MLEs from (3) are well documented, and have well known properties. See Lawless (1982). Of course, with censored observations the Weibull parameters are less well estimated than when all failure times are observed. The amount of lost information can be evaluated by comparing the expected (Fisher) information available in censored and uncensored samples (Escobar and Meeker, 1986). As a result, the usual recommendation of 20 subgroups of size 5, or 100 units total, must be increased. Using Fisher information appropriate sample sizes for any amount of censoring can be established. In the rest of the article we assume that initial Weibull parameters, denoted α_0 and β_0, have been determined from an in-control sample of sufficient size that the sampling variability can be ignored. This assumption corresponds to that made in the usual application of an $\overline{X}$ chart.

Now we turn attention to the ongoing monitoring procedure. An ideal procedure is able to detect decreases in the mean quickly, has an interpretable test statistic in the units of the original problem, and is easy to design. Note also that the procedure should work well with the small samples typical used in most process monitoring applications. We consider four distinct approaches: the Weibull CEV, MLE, exponential CEV (or score), and the extreme value CEV approaches.

2.1 Weibull CEV Method

Traditional $\overline{X}$ control charts use the sample average as test statistic, and set control limits at the estimated process mean plus or minus three standard errors $(\sigma/\sqrt{n})$. Thus, a naive application of a control chart in our application would ignore the censoring and set a lower control limit based on the three standard deviation unit rule. This is a poor choice in our application since the position of the control limit does not take into account the substantial skewness in the distribution of the sample averages caused by the underlying Weibull distribution and the censoring of the observations. The resulting naive control chart would have a much lower false alarm rate than desired, and thus also much lower power. This problem can be alleviated by taking into account the distribution of the sample average and using probability limits to set the control limit. However, we can do better still by not ignoring the censoring. One suggestion is to replace all censored observations with their conditional expected value (CEV) based on the initial parameter estimates. In other words, rather than work with the original sample that contains many observations that are censored at C, we use a sample of CEV weights where the censored observations are replaced by some weight larger than C.

It can be shown (Lawless, 1982) that assuming a Weibull distribution the conditional expected value of a observation censored at C, evaluated at α_0 and β_0, is

$$CE(W) \quad = \quad E(W \mid W \geq C) = \frac{\alpha_0 \Gamma^*(z_C, 1 + 1/\beta_0)}{\exp(-z_C)} \tag{6}$$

where $z_C = (C/\alpha_0)^{\beta_0}$, and $\Gamma^*(x,a) \; = \; \int_{y=0}^{x} y^{a-1} \exp(-y) dy$ is the incomplete Gamma function (Abramowitz and Stegun, 1972). The resulting sample of Weibull CEV weights is defined by

$$s(W) \quad = \quad \begin{cases} w & \text{if } w \leq C \text{ (not censored)} \\ CE(W) & \text{if } w > C \text{ (censored)} \end{cases} , \tag{7}$$

The conditional expected value (CEV) weight is a logical choice because then the expected sample average equals the process mean. Denoting the sample of n CEV weights as $s_1(W)$, $s_2(W)$, ..., $s_n(W)$, the test statistic is $\sum_{i=1}^{n} s_i(W) / n$, the sample average of the $s(W)$ weights. To distinguish this procedure from other procedures discussed later we call this approach the Weibull CEV method.

An appropriate lower control limit, one that gives .0027 probability of a false alarm, can be obtained through simulation for any values of α_0 and β_0. We may design the chart for the standardized Weibull to eliminate the dependence on α_0, but unfortunately, the design is complicated by a dependence on the shape parameter β_0 that can not be removed. As a result, no design figures are given for the Weibull CEV approach since a different figure would be needed for each value of β_0.

2.2 MLE Approach

Another way to monitor for decreases in the mean is to use a likelihood based approach. We are interested in detecting decreases in the mean caused by decreases in α. Based on the likelihood function (3), the maximum likelihood estimate (MLE) for α, denoted $\hat{\alpha}$, given β_0 is easily obtained. Of course, when β_0 is known w^{β_0} has an exponential distribution. To avoid possible division by zero (when all the observations in the sample are censored) we will use $1/\hat{\alpha}$ as test statistic. We get $\frac{1}{\hat{\alpha}} = \left[\frac{1}{r}\sum_{i=1}^{r} w_i + (n-r)C\right]^{-1/\beta_0}$. Using the notation introduced in the previous section the resulting sample of MLE based weights is defined by

$$s(M) = \begin{cases} w & \text{if } w \le C \text{ (not censored)} \\ \mathrm{C} & \text{if } w > C \text{ (censored)} \end{cases}, \qquad (8)$$

Based on (8), the appropriate test statistic from each sample of size n for the MLE approach is $\left(r/\sum_{i=1}^{n} s_i(M)\right)^{1/\beta_0}$.

Since the test statistic is $1/\hat{\alpha}$, decreases in α will be manifest as increases in the test statistic. As a result, the MLE approach requires an upper control limit (UCL). The position of the appropriate UCL for the MLE based control chart depends on the sample size, and the in-control probability censored. Figure 1 gives the appropriate UCL derived using simulation to have a theoretical false alarm rate of .0027, and assuming $\alpha_0 = \beta_0 = 1$.

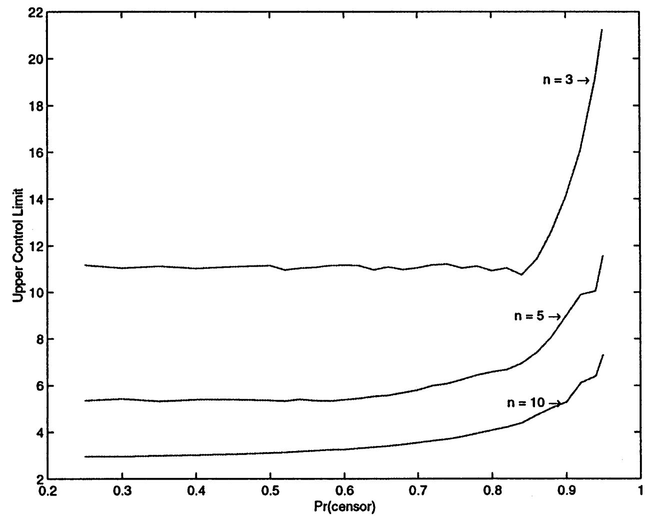

Figure 1: Standardized Upper Control Limit, ucl_0, for the MLE chart

Through the connection between the Weibull and the exponential distributions, the UCL appropriate in any given example problem equals $(\alpha_0\, ucl_0)^{-1/\beta_0}$, where α_0 and β_0 are the actual in-control process parameter, and ucl_0 is the standardized upper control limit given by Figure 1.

2.3 Score and Exponential CEV Method

Another likelihood based approach is to use α-scores. The α-score is defined as the derivative of the log likelihood function with respect to α evaluated at the parameter values of the stable (in-control) process. Due to their derivation, the α-scores form the basis of optimal tests for small changes in the process mean caused by changes in α. From (3), α-scores equal

$$\left.\frac{\partial \log L}{\partial \alpha}\right|_{\substack{\alpha=\alpha_0\\ \beta=\beta_0}} = \begin{cases} -\dfrac{\beta_0}{\alpha_0}+\dfrac{\beta_0}{\alpha_0}\left(\dfrac{w}{\alpha_0}\right)^{\beta_0} & \text{if } w \le C \text{ (not censored)} \\[2ex] \dfrac{\beta_0}{\alpha_0}\left(\dfrac{C}{\alpha_0}\right)^{\beta_0} & \text{if } w > C \text{ (censored)} \end{cases} \tag{9}$$

The corresponding test statistic is given by the sample average of the α-scores. One general disadvantage of using scores is that they do not usually have a physical interpretation. However, in this case, the α-scores are a simple linear translation of the conditional expected value weights derived on the exponential scale. To see this we determine the CEV weight for censored observations on the standardized exponential distribution scale. Based on the translating the censoring time to the exponential scale, we get

$$CE(X) = E\left(X \middle| X \geq (C/\alpha_0)^{\beta_0}\right) = (C/\alpha_0)^{\beta_0} + 1 \tag{10}$$

Thus the CEV weight on the exponential scale is $\begin{cases} x = (w/\alpha_0)^{\beta_0} & \text{if } w \leq C \text{ (not censored)} \\ \text{CE(X)} & \text{if } w > C \text{ (censored)} \end{cases}$.

Comparing the α-scores and the exponential CEV weight we see that one is a linear function of the other. This means that the CEV exponential method is optimal in the same sense as the α-score approach. Translating back to the Weibull scale, a sample average on the exponential scale is equivalent to using the test statistic $\alpha_0\left(\sum_{i=1}^{n}(s_i(X)/\alpha_0)^{\beta_0}/n\right)^{1/\beta_0}$ where $s(X)$ is given by (11).

$$s(X) = \begin{cases} w & \text{if } w \leq C \text{ (not censored)} \\ \alpha_0 CE(X)^{1/\beta_0} & \text{if } w > C \text{ (censored)} \end{cases} \tag{11}$$

The appropriate lower control limit derived using simulation for the exponential CEV weight is given in Figure 2 assuming $\alpha_0 = \beta_0 = 1$. The irregular parts of the figures are due to the discreteness inherent in the problem. The appropriate control limit given α_0 and β_0 is $\left(\alpha_0\, lcl_0^X\right)^{1/\beta_0}$, where lcl_0^X is read from Figure 2.

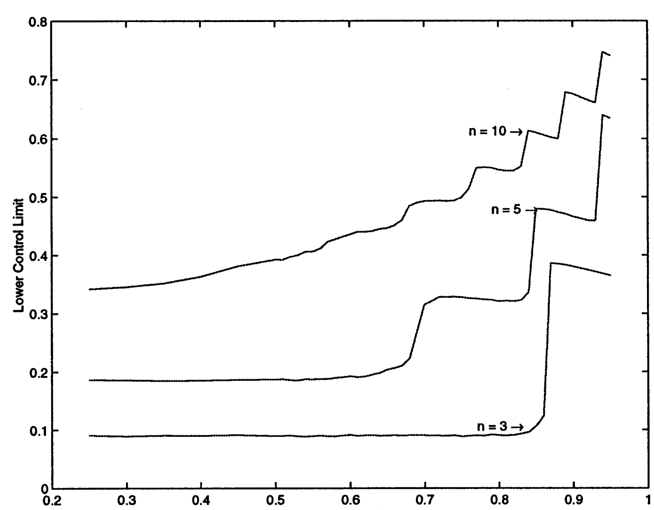

Figure 2: Standardized Lower Control Limit, lcl_0^X, for the CEV Exponential chart

2.4 Extreme Value CEV Method

The final approach we shall consider is the CEV extreme value method. With this approach the censored observations are replaced by their conditional expected values derived on the extreme value scale. Working with the standardized extreme value distribution, the conditional expected value for a censored observation is

$$CE(V) \;=\; E\big(V \,\big|\, V \geq \beta_0 \log(C/\alpha_0)\big) \;=\; \int_{x=\beta_0 \log(C/\alpha_0)}^{\infty} x e^{x} \exp(-e^{x})dx \Big/ S\big(\beta_0 \log(C/\alpha_0)\big) \qquad (12)$$

where $S(\;)$ is the survivor function of the standardized extreme value distribution given by (4). The value of $CE(V)$ can be determined through numerical integration. Given $CE(V)$ from (12) we get the weights given by (13) on the original Weibull scale.

$$s(V) \;=\; \begin{cases} w & \text{if } w \leq C \text{ (not censored)} \\ \alpha_0 \exp\big(CE(V)/\beta_0\big) & \text{if } w > C \text{ (censored)} \end{cases}, \qquad (13)$$

A sample average on the extreme value scale, namely, $\left(\sum_{i=1}^{r} \beta_0 \log(w_i/\alpha) + (n-r)CE(V)\right)\Big/n$, corresponds to the test statistic $\left(\prod_{i=1}^{n} s_i(V)\right)^{1/n}$, or the geometric sample average, on the original Weibull scale.

Figure 3 gives the appropriate lower control limit for the extreme value CEV approach derived through simulation to give a false alarm rate of .0027, and assuming $\alpha_0 = \beta_0 = 1$. The appropriate lower control limit for any values of α_0 and β_0 is $\alpha_0 \exp\big(lcl_0^V/\beta_0\big)$, where lcl_0^V is read from Figure 3.

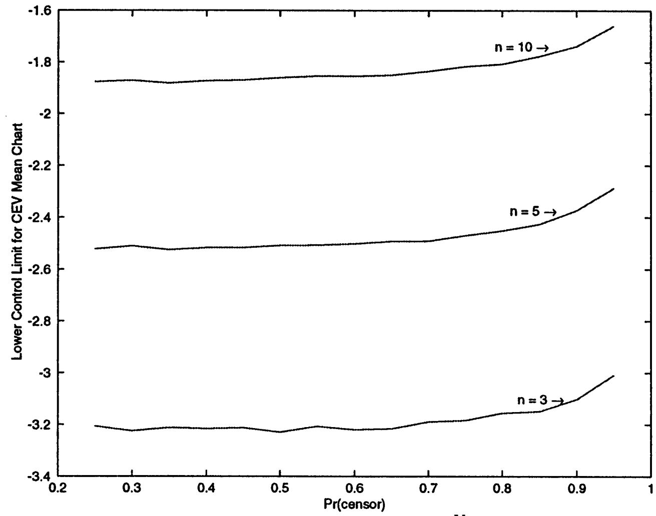

Figure 3: Standardized Lower Control Limit, lcl_0^V, for the CEV Extreme chart

2.5 Summary of Methods

As shown, the four distinct methods all correspond to replacing the censored observations with a different value. The MLE approach uses C and adjusted the denominator, while each of the other methods uses a CEV weight derived on a different scale. It is interesting to compare the different replacement values when they are all translated back to the original Weibull scale. Figure 4 shows the resulting replacement value plotted against the censoring time C.

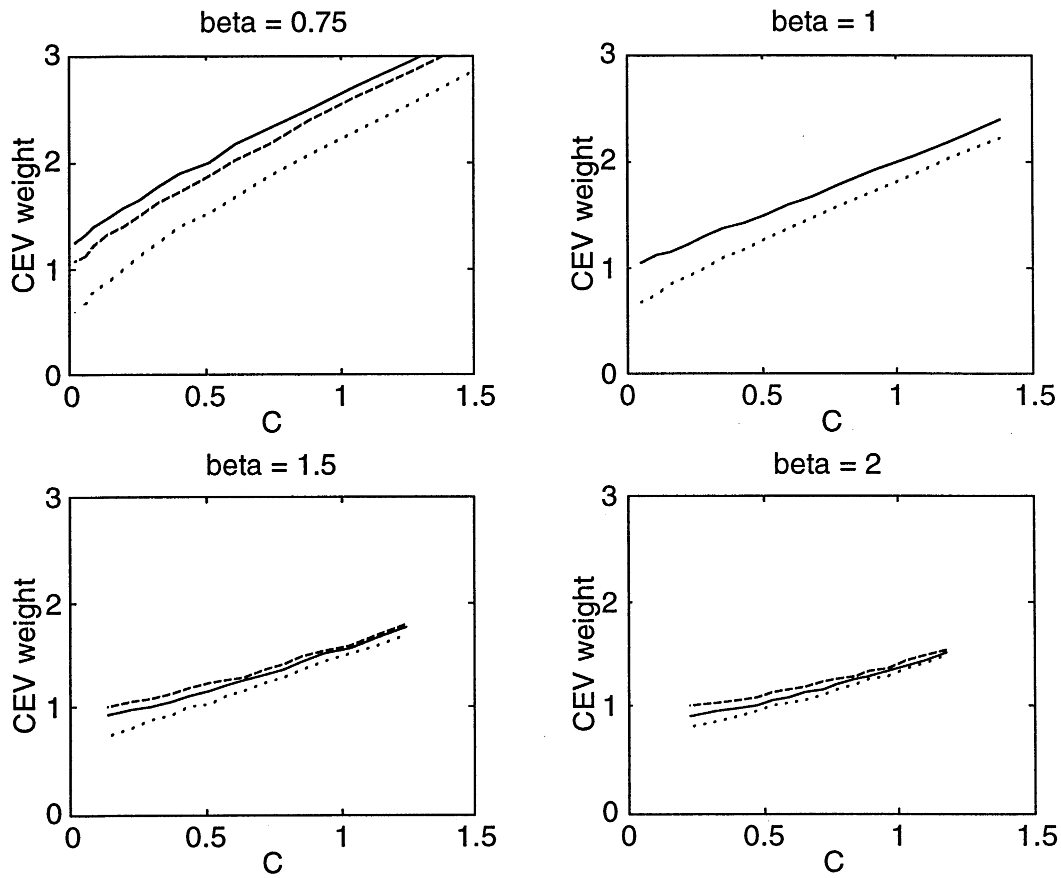

Figure 4: Conditional Expected Value Weights as a Function of the Censoring Time
solid line - Weibull CEV, dotted line - exponential CEV, dashed line - extreme value CEV

Figure 4 shows that the CEV weights are all larger than C, as expected. However, the values are similar. In fact, the major difference between the approaches lies more in the test statistic used to monitor for decreases in the mean, and the ease of determining the control limit. The test statistic for the Weibull CEV method is the easiest, since we use the (arithmetic) sample average, as is typically for $\overline{X}$ control charts. However, the determination of an appropriate control limit is complicated by a dependence on β_0 that can't be removed. This is a problem, since it means that it is not possible to give one figure from which the appropriate control limit can be determined for any situation. On the other hand, the test statistics for the exponential CEV and MLE approaches are not simple, since they are not interpretable in terms of lifetimes and thus may be difficult for production personnel to understand. The design of control limits for the exponential CEV and MLE approaches, on the other hand, is easier than for the Weibull

CEV approach, and generally applicable design figures are given in previous sections. From this perspective the best approach is the extreme value CEV method. The extreme value CEV control limits can also be determined independent of β_0, and the test statistic is the geometric sample average. This is test statistic is not as familiar to production personnel as the sample average, but it does retain interpretability. Of course, in each case, we could simply use the sample average as test statistic with what every replacement value was suggested. However, as seen with the Weibull CEV method, then the design of control limits can not be divorced from the value of β_0. Table 1 summarizes the comparison of the four methods in terms of the replacement value, or censored weight, test statistic and ease of design.

Table 1: Method Summary and Comparison

Method	Censored Weight	Test Statistic	Design
Weibull CEV	$CE(W)$	$\sum s_i(W)/n$	depends on β_0
Extreme value CEV	$\alpha_0 \exp(CE(V)/\beta_0)$	$\left(\prod s_i(V)\right)^{1/n}$	Figure 3
Exponential CEV	$\alpha_0 CE(X)^{1/\beta_0}$	$\alpha_0\left(\sum_{i=1}^{n}(s_i(X)/\alpha_0)^{\beta_0}/n\right)^{1/\beta_0}$	Figure 2
MLE	C	$\left(r/\sum_{i=1}^{n} s_i(M)\right)^{1/\beta_0}$	Figure 1

3. Power Comparison of the Control Charts

In this section, we compare the power of the four control charts, namely the Weibull CEV, exponential CEV (alpha score), MLE, and the extreme value CEV methods, for detecting mean shifts in the process. Since the rate of censoring plays a role, we compare the power of the charts with in-control censoring rates of 50% and 90%. Figure 5 gives results for subgroups of size 5 or 10, where control limits are determined through simulation to have false alarm rates equal to .0027. For all approaches other than the Weibull CEV method, the control limits can be determined from the design figures given in Section 2. For comparison purposes the performance in the uncensored case is given with a dashed line in each plot. The Weibull CEV, exponential CEV and MLE approaches are all virtually indistinguishable in term of power. The extreme value CEV approach has good power, but is less powerful than the other methods. The results are given in terms of standard deviation shifts in the extreme value location parameter because then the results can be generalized as explained below. Decreases in the extreme value mean correspond to decreases in $\log(\alpha)$ on the Weibull scale.

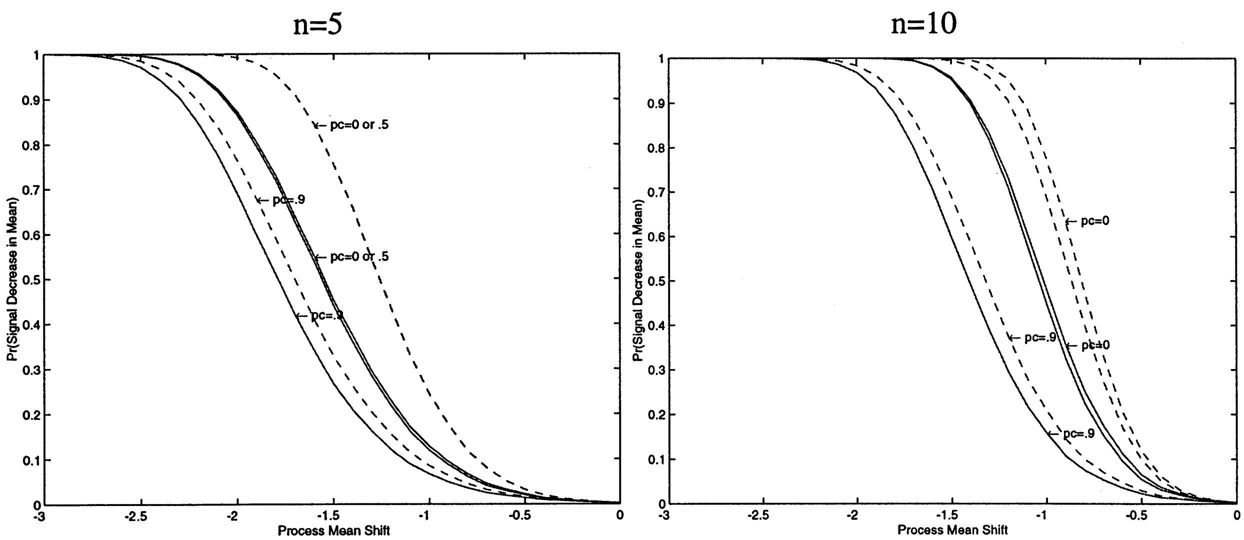

Figure 5: Power to Detect Shifts in the Mean of the Extreme Value Distribution, $\alpha_0=\beta_0=1$
dashed line - CEV Weibull, CEV exponential, MLE method
solid line - Extreme value method

From Figure 5 we see that for moderate censoring proportions, such as 50% censoring, there is almost no loss in power to detect process mean decreases. The loss in power becomes more pronounced when the censoring rate approaches 90%. One reason why we can detect decreases in the mean quickly even with high censoring rates is that as the mean decreases the censoring rate also decreases, and thus each sample contains more information. 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 be the most expensive. The optimal tradeoff point depends on the testing costs and the consequences of false alarms and/or missing process changes.

Figure 5 can be used to determine the approximate power to detect a decrease in alpha from α_0 to α_1 given β_0 for any of the methods. The power of the Weibull CEV approach actually depends on β_0, but the difference is small for β_0 in the usual range .5 to 2. The horizontal axis gives the mean shift in terms of the standard deviation units of the location parameter of the extreme value distribution. Using the relationship between the Weibull and extreme value distributions, the number of standard deviation units on the extreme value scale is given by $\beta_0\sqrt{6}\log(\alpha_0/\alpha_1)/\pi$. Thus, for example, if α_0 = 2.5 and β_0 = 1.25 a decrease to α_1 = .5 corresponds to a 1.57 standard deviation shift in the location parameter of the extreme value distribution. Thus, from Figure 5, our chance of detecting this shift with a single sample using the exponential CEV method with samples of size 5 and an in-control censoring rate of 50% is around 80%.

The control charts derived in this article were designed to detect decreases in mean lifetime caused by decreases to the scale parameter of the Weibull. However, it is interesting to explore the power of the various methods when decreases in the mean lifetime are due to changes in the mean holding the variance constant. This corresponds to simultaneous changes to both the scale and shape parameters of the Weibull. As explained in Section 1, this kind of process change is not as compelling since it implies the shape of the distribution changes. With simultaneous changes to α and β it is difficult to give general results because the performance will depend on β_0. For illustration purposes Figure 6 compares the performance of the monitoring procedures when β_0 = 1. Based on this limited comparison, the Weibull CEV, exponential CEV and MLE approaches again give very similar power. Now, however, the extreme value CEV approach is the most powerful. It is not clear why with the Weibull CEV, exponential CEV and MLE approaches we do better with 90% than with no censoring.

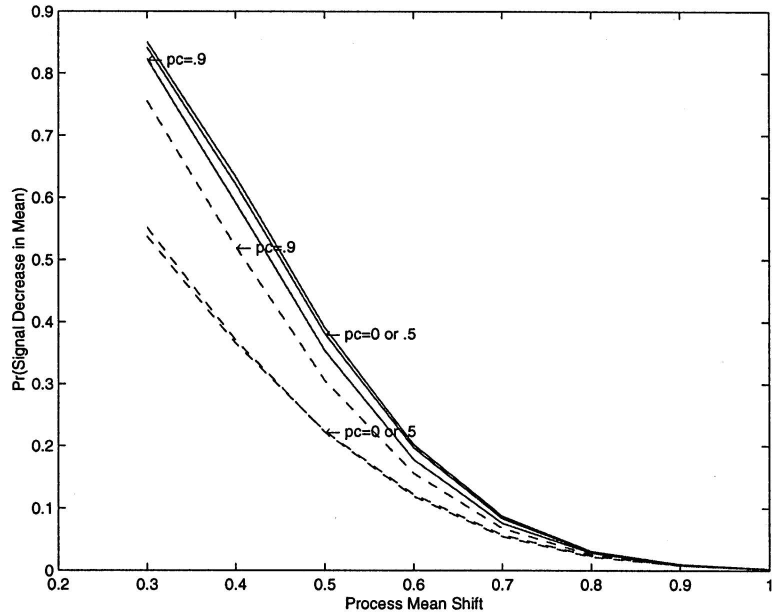

Figure 6: Power to Detect Shifts in the Mean with Fixed Variability, $\alpha_0=\beta_0=1$
dashed line - CEV Weibull/score/MLE method
solid line - Extreme value method

4. Example

In the rust test example described in the introduction, an initial sample of 100 subgroups of size 3 was selected from historical monitoring records. The censoring time was 20 days. In the data there was a 77% censoring rate. Using the MLE procedure described in Lawless, 1982, we estimate α_0 = 48.04 and β_0 = 1.51. It was also verified that even though our data are discrete a Weibull distribution was a reasonably good fit to the data. Based on the analysis presented in this article, the extreme value CEV method was used. With a censoring time equal to 20 and α_0 and β_0 we get $CE(V) = -.023$ from (12). Translating this into the original Weibull scale means that all censored observations were assigned a weight equal to 47.3. Since the salt spray chamber has a maximum capacity of 60 units a monitoring procedure based on subgroups of size 3 was utilized. Thus, with the estimated in-control Weibull parameters the standardized lower control limit for the extreme value CEV $\overline{X}$ chart is –3.2. This value was determined through simulation, but may be approximated from Figure 3. Scaling the control limits by the in-control parameter estimates gives a lower control limit of 5.9 for the extreme value CEV $\overline{X}$

chart. The resulting CEV $\overline{X}$ chart for the example data is given by the first 100 subgroups in Figure 7.

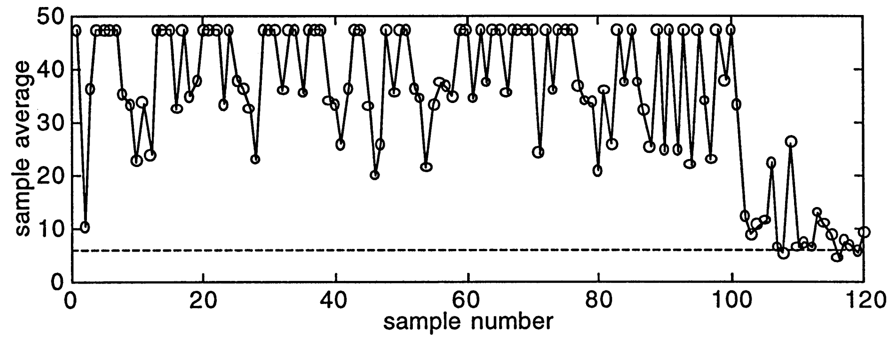

Figure 7: Example Extreme Value CEV $\overline{X}$ charts with n = 3

Figure 7 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 parameter estimates. As a result, we may continue to monitor the process for deterioration using the extreme value CEV charts with the given control limit. The final 20 subgroups are simulated assuming a process shift to α_1 = 12. These simulated observations show that the control chart would signals a decrease in the mean with fairly high probability.

5. Summary and Conclusions

In applications where observed data may be censored, traditional process monitoring approaches, such as $\overline{X}$ 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 applicable when observations are censored at a fixed levels are proposed. The monitoring procedure is derived assuming the process has an underlying Weibull distribution. A number of possible procedures are compared, including procedures based on maximum likelihood, score and conditional expected value. A chart based on the idea of replacing all observations by their conditional expected values (CEV) weights calculated on the extreme value scale is recommended. The procedure is illustrated with an example.

Acknowledgments

This research was supported, in part, by the Natural Sciences and Engineering Research Council of Canada, and General Motors of Canada.

References

Abramowitz, M., Stegun, I.A. (1972), *Handbook of Mathematical Functions,* Dover, New York.

Escobar, L.A., Meeker, W.Q., Jr. (1986), "Elements of the Fisher Information Matrix for the Smallest Extreme Value Distribution and Censored Data, *Applied Statistics*, 35, 80-86.

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

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

Nelson, W. (1990), *Accelerated Testing*, John Wiley and Sons, New York.

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