Monitoring Processes With Highly Censored Data

STEFAN H. STEINER and R. JOCK MACKAY

University of Waterloo, Waterloo, Canada N2L 3G1

The need for process monitoring in industry is ubiquitous. By monitoring process output, process changes may be rapidly detected and problems corrected. However, in many industrial and medical applications, observations are censored due to either 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.

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 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 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. For simplicity we will always refer to the variable of interest as strength, although it may be any other censored response.

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 weight (CEV) 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%.

This article is organized in the following manner. We first introduce the CEV charting procedure that allows for the rapid detection of deterioration in the process quality when the monitored output is cen-

Dr. Steiner is an Associate Professor in the Department of Statistics. He is a member of ASQ. His email address is shsteine@uwaterloo.ca

Dr. MacKay is an Associate Professor in the Department of Statistics.

sored. Figures needed to determine control limits are given. The use of the procedure is then illustrated with the first example described above. Finally, we determine the power of the proposed procedure and compare it with more traditional approaches.

CEV Control Charts for Censored Data

In this section, control charts are derived for detecting mean and dispersion shifts in the process that produces censored data. We shall assume the observations are right-censored, though similar results may be obtained for other forms of censoring. With right-censored data, the goal of the CEV control charting is to detect decreases in the process mean and/or increases in the process standard deviation. In other words, the two control charts have only one-sided control limits. As will be shown, it is feasible to detect such process changes surprisingly well because they lead to decreases in the censoring proportion, which in turn means that each subgroup provides more process information. On the other hand, with right-censored data, it is very difficult to detect increases in the process mean because such changes increase the censoring proportion. Similarly, if the censoring proportion is greater than 50%, decreases in the process dispersion also lead to more censored observations. Subgroups with all censored observations provide very little information about shifts in the process parameters. 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.

To derive the charts, suppose that when the process is in control, the strength T can be represented by a normal random variable with mean, μ, and standard deviation, σ. Other distributional assumptions such as exponential and Weibull are also possible and do not change the procedure markedly. Denote the censoring level as C, that is, the exact strength is not observed for units with strength greater than C. Then, the probability of censoring equals

$$p_c = 1 - F(C) = Q\left(\frac{C-\mu}{\sigma}\right),$$

where $Q(C)$ is the survivor function of the standard normal distribution.

CEV Weights

The proposed control charts are based on the simple idea of replacing each censored observation with its conditional expected value. Based on these CEV weights, the subgroup averages and sample standard deviations are plotted in a manner similar to the traditional $\overline{X}$ and S charts. It can be shown (Lawless (1982)) that the conditional expected value of a censored observation is

$$w_c = E(T \mid T \geq C) = \mu + \sigma\left(\frac{\phi(z_c)}{Q(z_c)}\right), \qquad (1)$$

where $\phi(z)$ is the probability density function of the standard normal and $z_c = (C-\mu)/\sigma$. We define the conditional expected value (CEV) weight, w, of each unit as:

$$w = \begin{cases} t & \text{if } t \leq C \text{ (not censored)} \\ w_c & \text{if } t > C \text{ (censored)} . \end{cases} \qquad (2)$$

The CEV control charts consist of run charts of the subgroup average and standard deviation of the CEV weights; we call the resulting control charts the CEV $\overline{X}$ and the CEV standard deviation (S) chart, respectively. This method of deriving sample averages and standard deviations has a Bayesian flavor since the calculation of the CEV weights (for censored observations) depends on the in-control parameter values μ and σ. In applications, μ and σ are estimated from in-control process data in the initial implementation phase of the monitoring procedure. See the section on initial implementation for more details of how to estimate μ and σ in practice when observations are censored. For now, we assume the in-control values are known.

The idea of using CEV weights is intuitive and may also be justified based on the likelihood function. It is well known (Lawless (1982)) that for a subgroup of size n, the log-likelihood for the right censored case is

$$\log L(\theta) = (n-r)\log\left[Q\left(\frac{C-\mu}{\sigma}\right)\right] + \sum_{i\in D}\log\left[\phi\left(\frac{t_i-\mu}{\sigma}\right)\right], \qquad (3)$$

where D represents the set of all observations that were not censored and r equals the number of uncensored observations. It is also known that the optimal test statistic to detect small changes from the in-control mean is based on the mean score (Cox and Hinkley (1974)). In the normal case, the mean score, denoted m, is defined as the first derivative of the log likelihood with respect to μ, that is,

$$m = \frac{\partial \log L}{\partial \mu}$$

$$= \begin{cases} \dfrac{t-\mu}{\sigma^2} & \text{if } t \le C \text{ (not censored)} \\ \dfrac{\phi\left(\frac{C-\mu}{\sigma}\right)}{\sigma Q\left(\frac{C-\mu}{\sigma}\right)} & \text{if } t > C \text{ (censored)} \end{cases} . \tag{4}$$

Comparing the mean score in Equation (4) with the CEV weights given by Equation (2) shows that w equals $\mu+\sigma^2 m$ for both censored and uncensored observations. Thus, for normal data, the mean scores are a linear translation of the conditional expected value weights, and control charts based on either will have equivalent operating characteristics. 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. For control charting, we recommend using the CEV weights since they have a direct physical interpretation in the original scale of measurement.

Determining CEV Control Limits

For CEV charts, control limits depend in a complex way on the subgroup size and the censoring proportion, p_c. There is no simple rule such as the traditional plus or minus three standard deviation limits. Instead, Figures 1 and 2, derived from simulations, are provided to aid in the construction of control limits for the CEV $\overline{X}$ and S control charts. The figures are based on the assumption that p_c, the in-control proportion censored is known. Figure 1 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 2 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 1 and 2 are on a log scale.

For subgroup 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 subgroup size of 8, and an in-control probability of censoring equal to .9, using Figure 2 we choose a standardized upper control of 1.3. The irregular parts of Figure 1 are due to the discreteness inherent in the problem.

The control limits shown in Figures 1 and 2 are standardized in the sense that they give the appropriate control limits, given the subgroup size, and

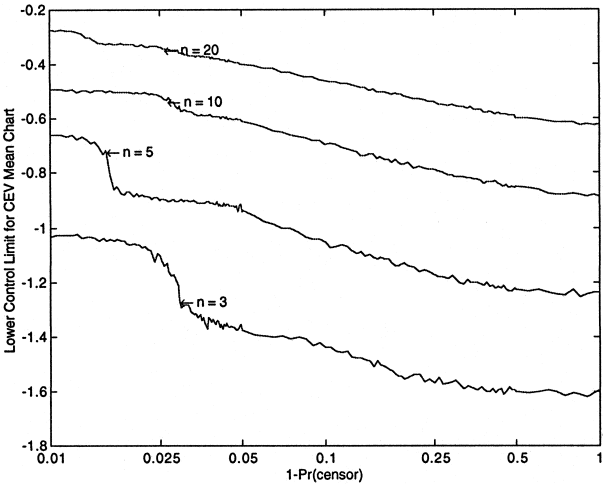

FIGURE 1. Plot of the Standardized Lower Control Limit ($LCL_{\overline{X}}$) for the CEV $\overline{X}$ Chart.

p_c, assuming the in-control process has mean of zero and standard deviation of one. The control limits appropriate in any given problem may be obtained using

$$\begin{aligned} &\text{lower control limit for CEV } \overline{X} \text{ chart } = LCL_{\overline{X}}\sigma + \mu \\ &\text{upper control limit for CEV } S \text{ chart } = UCL_S\sigma, \end{aligned} \tag{5}$$

where μ and σ are the in-control process parameters and $LCL_{\overline{X}}$ and UCL_S are the standardized control limits given by Figures 1 and 2, respectively. Note that for both charts using a center line is not of much

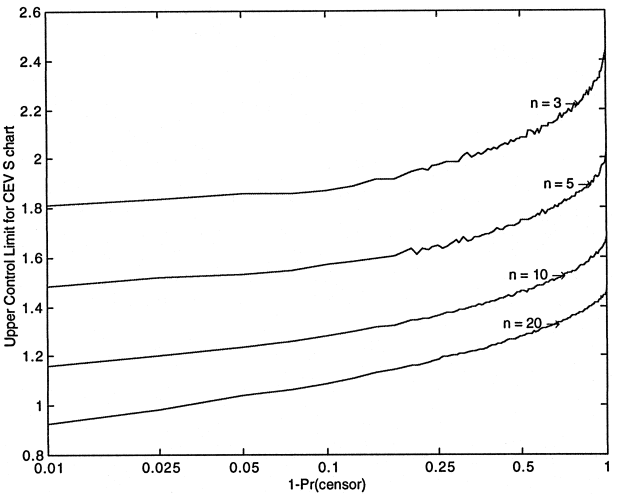

FIGURE 2. Plot of the Standardized Upper Control Limit (UCL_S) for the CEV S Chart.

value since the distributions of the sample average and sample standard deviation of the CEV weights are highly skewed when the censoring proportion is large.

Initial Implementation

As with traditional monitoring procedures, the establishment of CEV control charts requires a two-step process. The first step, often called the initial implementation phase, involves collecting a setup sample from the process while it is in control. When working with uncensored data, guidelines suggest that 100 observations or more are 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. See Quesenberry (1993) for a detailed analysis of the initial sample size question. From the initial subgroups, the appropriate control charts are established. If there is any evidence of instability in the initial sample, that is, 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.

The following algorithm illustrates the initial implementation procedure for CEV $\overline{X}$ and S charts for a specified censoring level C.

1. Collect q subgroups of size n.
2. Estimate the in-control mean and standard deviation, μ and σ, using maximum likelihood (see Appendix A).
3. Determine the censored CEV weight w_c using Equation (1) based on the estimates of μ and σ, and replace all censored observations with the value w_c.
4. Create one-sided CEV $\overline{X}$ and S charts by plotting the subgroup averages and standard deviations with control limits determined using the design Figures 1 and 2.
5. Look for any out-of-control signals (points outside the control limits) on the charts. Examine process conditions at the time any out-of-control 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 in-control process mean and standard deviation. The CEV chart design procedure is somewhat self-correcting since, for example, 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 smaller. Note that the procedure estimates the process variability using the complete sample rather than just the within-subgroup dispersion as is regularly done for control charts.

In Step 2, maximum likelihood estimation (MLE) is suggested because it works well for large samples, which are typical when considering all the data available in the initial implementation. Note that charting maximum likelihood estimates calculated from each subgroup is not a feasible alternative to the sample average and sample standard deviation of the CEV weights. In the extreme case that all observations are censored, finite MLE's do not exist. With small subgroups and substantial censoring this occurs with non-negligible probability. In addition, the calculation of the MLE's is iterative, thus requiring a fairly substantial computational effort that may be onerous.

The sample size needed to initially estimate the in-control parameters with precision (Step 2) can be determined through the expected information content of a censored sample in terms of Fisher information (see Appendix B). Fisher information determines theoretically how much information regarding either the mean or standard deviation is lost due to the censoring. Note also that a different censoring level can be used in the implementation phase to reduce the overall sample size required.

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 of the 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 previously undertaken. When out-of-specification bond strengths were detected, the process was investigated but typically no action was taken.

In the data, the first 125 observations of which are given in Table 1, there was an 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 in the Appendix A, we estimate the process mean

TABLE 1. First 125 Observations of the Example Data

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

and standard deviation as $\mu = 11.1$ and $\sigma = 1.24$. With a censoring level of 10, we get $w_c = 11.44$ from Equation (1). This is the weight assigned to all censored observations in the CEV monitoring procedure. Based on subgroups of size 5 and an 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 from Figures 1 and 2. Scaling the control limits by the estimated mean and standard deviation according to Equation (5) 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 in Figure 3.

Figure 3 shows that in the initial implementation there were no out-of-control points. Thus, the initial data appear to come from an in-control process. As a result, we may continue to monitor the process for deterioration using the CEV charts with the given control limits. To reduce the in-control out-of-specification rate from around 14%, the common cause of variation must be addressed.

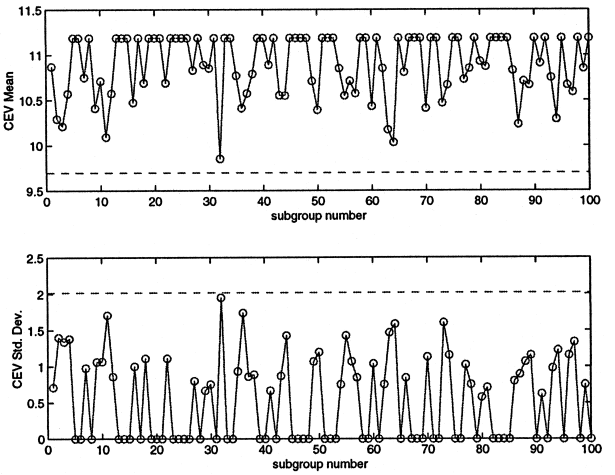

FIGURE 3. Example CEV $\overline{X}$ and S Charts With $n = 5$.

CEV Control Chart Performance

In this section the power of CEV $\overline{X}$ and S control charts to detect process changes is explored. It is shown that when the censoring proportion is very large, the CEV $\overline{X}$ chart alone suffices to detect both mean and standard deviation shifts in the process. In addition, we compare the performance of the CEV control chart with more traditional control charts, such as 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. Figures 4 and 5 give results for changes in the process mean and standard deviation, respectively, for different initial censoring proportions. For both Figures 4 and 5, the control limits of the charts are determined from Figures 1 and 2, and thus the false alarm rate of all the charts is set at .0027. The results are based on simulation using 200,000 trials for each point. For comparison purposes, the performance in the uncensored case ($p_c = 0$) is given with a dashed line in each plot. In Figure 4, the horizontal axis corresponds to shifts in units of the standard deviation.

In Figure 4 we see that for the CEV $\overline{X}$ chart, the decrease in power as the censoring proportion increases is quite gradual. In fact, for moderate censoring proportions, such as 50% censoring, there is almost no loss in power to detect process mean decreases. On the other hand, for the CEV S chart shown in Figure 5, 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

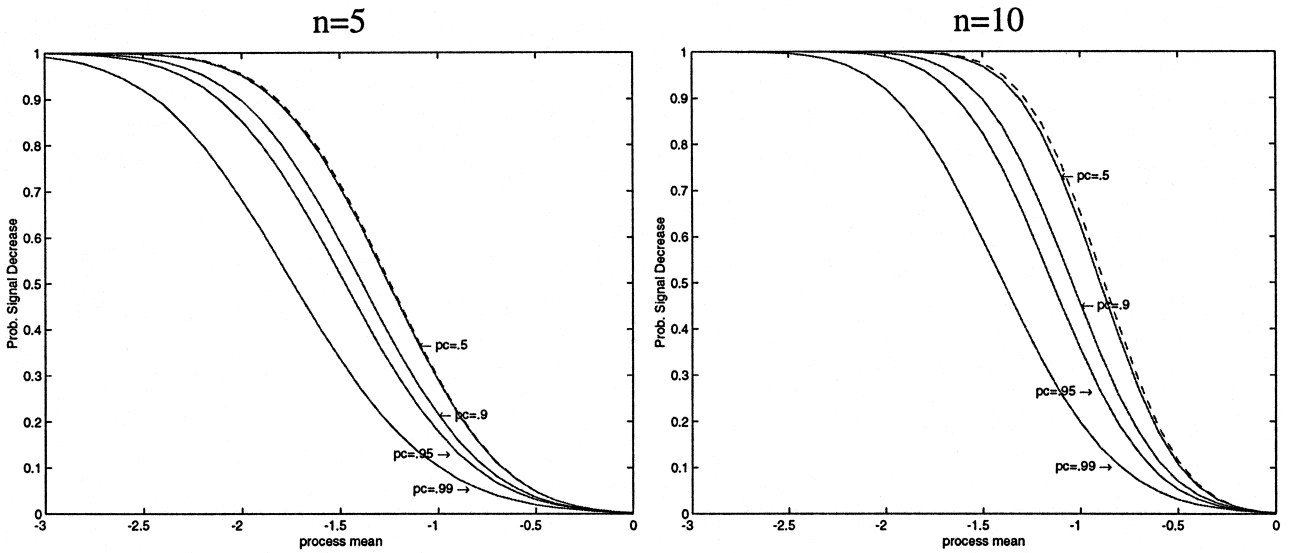

FIGURE 4. Power of the CEV $\overline{X}$ Chart to Detect Process Mean Shifts (No Censoring Case Given by the Dashed Line).

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.

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 missing process changes.

The CEV $\overline{X}$ chart 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,

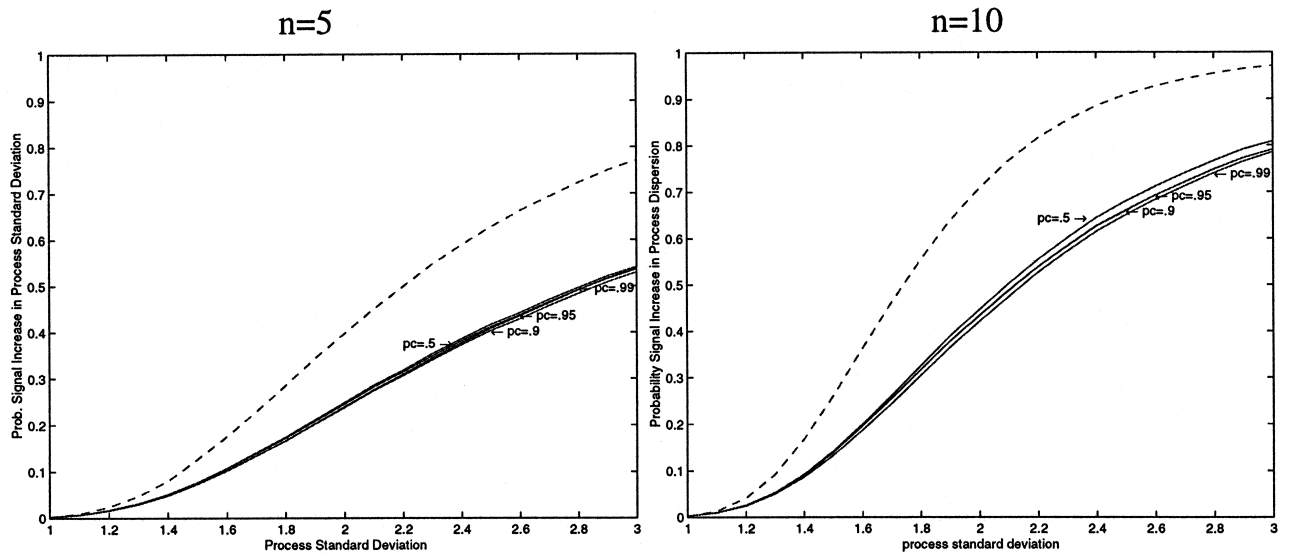

FIGURE 5. Power of the CEV S Chart to Detect Process Standard Deviation Shifts (No Censoring Case Given by the Dashed Line).

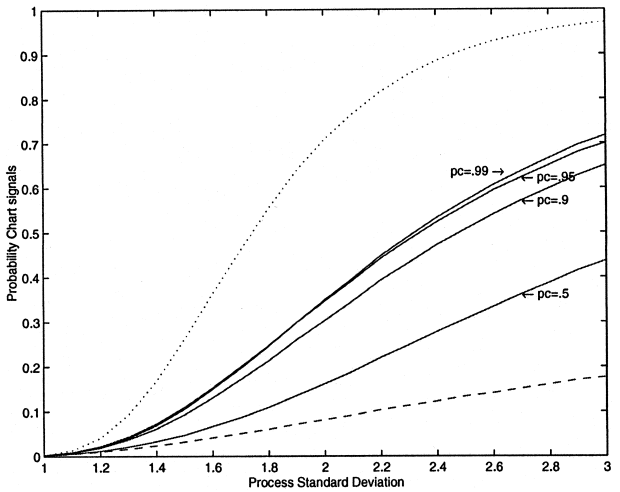

FIGURE 6. Power to Detect Standard Deviation Shifts With the CEV $\overline{X}$ Chart, $n = 10$ ($\overline{X}$ with No Censoring Chart Given by the Dashed Line and S Chart with No Censoring Chart Given by the Dotted Line).

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 thus 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.

Comparison of CEV Control Chart Performance to Traditional Charts

As a further comparison, we may consider the use of a traditional control chart 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 due to discreteness since the np chart can not necessarily be setup to have a particular false alarm rate. This is illustrated in Table 2 which 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.

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. 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 observed non-censored values. Figure 7 suggests that as the censoring proportion increases, the performance of the two charts converge. 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 proportions censored are not directly comparable since they have different false alarm rates.

TABLE 2. np Chart Decision Rule When $n = 5$: If Number Censored $< x$ Then Signal

p_c	x	Pr(false alarm)
.50	1	.031
.75	1	.001
.90	3	.009
.95	3	.001
.99	4	.001

As discussed above, 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 pro-

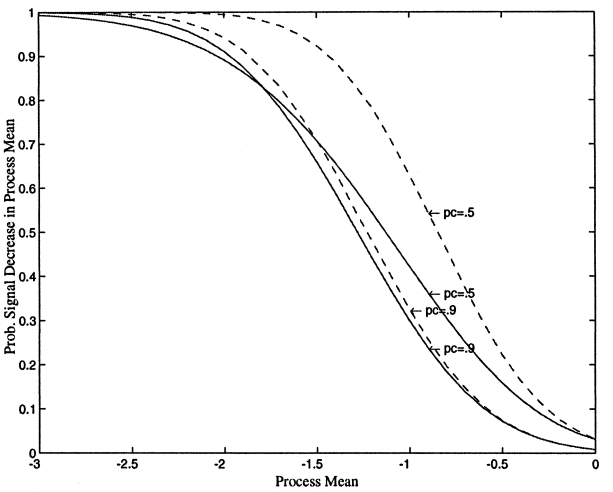

FIGURE 7. Comparison of Performance Between CEV $\overline{X}$ and np Charts, $n = 5$ (CEV $\overline{X}$ Chart Given by Dashed Line).

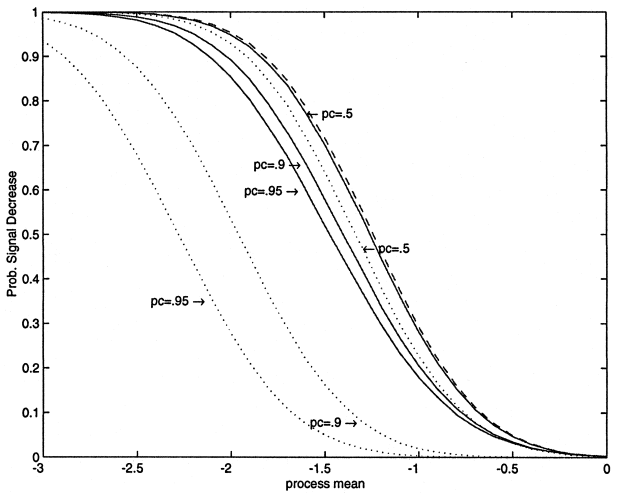

FIGURE 8. Comparison of Performance Between CEV $\overline{X}$ and Naive $\overline{X}$ Charts, $n = 5$ (No Censoring Given by Dashed Line, CEV Charts Given by Solid Lines, and Naive Charts by Dotted Lines).

portion 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% as 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.

The comparison between the CEV $\overline{X}$ chart and the traditional Shewhart $\overline{X}$ chart is also difficult. A naive application of an $\overline{X}$ chart would ignore the censoring and set a lower control limit at $\overline{\overline{X}} - 3\widehat{\sigma}/\sqrt{n}$, where the standard deviation estimate is given by either $\bar{s}/c_4$, or $\overline{R}\,/d_2$, where $\bar{s}$ 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 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 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 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.

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 fixed levels are proposed. The proposed charts are based on the idea of replacing all observations by their conditional expected values and then charting standard statistics of these CEV weights. The CEV weights are equivalent to likelihood-based mean scores if the underlying distribution is normal. Control limits for CEV control charts given are derived from simulation of the sampling distributions of the subgroup statistics assuming that the in-control distribution is known.

There are many variations of the proposed charts that we do not consider. The monitoring procedure we have constructed is derived assuming the process has an underlying normal distribution, but the same methodology is applicable to other distributions. In addition, since the amount of information in each subgroup to detect process changes is small when the censoring is severe, cumulative sum charts based on the average CEV can be constructed. The procedure defined here is a fixed one-sided (type I) censoring scheme. There are many other practical censoring schemes to be investigated. The plug gauge example mentioned in the introduction has two-sided censoring. In many applications, the censoring may not be fixed but be due to one or more competing risks. In the bond strength example, if the force required to break the bond or tear the foam backing, whichever occurred first, was measured, then a different charting methodology is required to monitor the bond strength process.

Appendix A: Maximum Likelihood Estimation

Using the log-likelihood function in Equation (3) we may derive maximum likelihood estimates for the process mean and standard deviation from normal censored data. We present an iterative approach due to Sampford and Taylor (1959) since it uses CEV weights. The Sampford and Taylor method is an application of the expectation-maximization (EM) algorithm discussed by Dempster, Laird, and Rubin (1977). The procedure is iterative and involves replacing each censored observation with its conditional expected value, given by Equation (2) where we replace μ and σ with the current best guess for the process mean and standard deviation denoted $\widehat{\mu}_0$ and $\widehat{\sigma}_0$. Based on these CEV weights, we re-estimate the process mean and variance as:

$$\widehat{\mu}_1 = \sum_{i=1}^{n} \frac{w_i}{n}, \qquad \widehat{\sigma}_1 = \frac{\sum_{i=1}^{n} (w_i - \widehat{\mu}_0)^2}{r + (n-r)\lambda(z_C)}, \qquad \text{(A1)}$$

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, being near 1 when the censoring proportion is small and near 0 when the censoring proportion is large. As a result, the term $r + (n-r)\lambda(z)$ can be thought of as a sample size, adjusted for the number of censored observations.

To find the MLE's, we iteratively apply Equations (2) and (A1) to the data until the estimates converge. The iterations rapidly converge (less than 10 iterations) to the MLE's so long as good initial values are employed. In most cases, good initial values are the sample mean and sample standard deviation obtained by ignoring the censoring.

Appendix B: Expected Information in Censored Samples

The relative value of censored and uncensored samples of the same size may be compared using Fisher information. The inverse of the Fisher information gives the asymptotic variance of the maximum likelihood parameter estimates. Expected Fisher information is defined as the expected value of minus the second derivative of the log likelihood function. In the censored normal case, we may de rive the information matrix **I** from the log-likelihood expression in Equation (3).

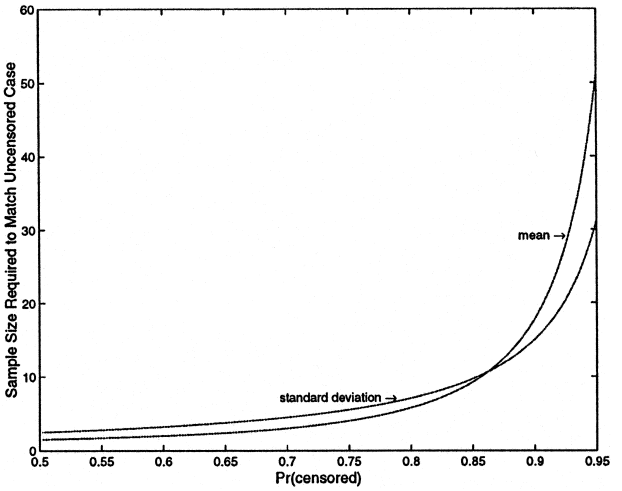

FIGURE B1. Plot of the Censored Sample Size Multiples Needed to Match the Estimation Precision of an Uncensored Sample.

Figure B1 shows the relative sampling size required to match the sampling variability in an uncensored sample for the mean and standard deviation. Note that for small censoring proportions we can estimate the mean and standard deviation relatively well. However, when the censoring proportion grows 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. 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.

Acknowledgments

This research was supported, in part, by the Natural Sciences and Engineering Research Council of Canada and General Motors of Canada. All simulations were executed using MATLAB, a registered trademark of MathWorks.

References

Cox, D. R. and Hinkley, D. V. (1974). *Theoretical Statistics*. Chapman and Hall, London, UK.

Dempster, A. P.; Laird, N. M.; and Rubin, D. B. (1977). "Maximum Likelihood from Incomplete Data via the EM Al-

gorithm" (with discussion). *Journal of the Royal Statistical Society* B 39, pp. 1–38.

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

Quesenberry, C. P. (1993). "The Effect of Sample Size on the Estimated Limits for $\overline{X}$ and X Control Charts". *Journal of Quality Technology* 25, pp. 237–247.

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

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

Key Words: *Censored Data, Process Control, Scores*.