Appl. Statist. (2001)
50, *Part* 3, *pp.* 293–302

Monitoring processes with data censored owing to competing risks by using exponentially weighted moving average control charts

Stefan H. Steiner and R. Jock MacKay
University of Waterloo, Canada

[Received September 1999. Final revision November 2000]

Summary. In industry, process monitoring is widely employed to detect process changes rapidly. However, in some industrial applications observations are censored. For example, when testing breaking strengths and failure times often a limited stress test is performed. With censored observations, a direct application of traditional monitoring procedures is not appropriate. When the censoring occurs due to competing risks, we propose a control chart based on conditional expected values to detect changes in the mean strength. To protect against possible confounding caused by changes in the mean of the censoring mechanism we also suggest a similar chart to detect changes in the mean censoring level. We provide an example of monitoring bond strength to illustrate the application of this methodology.

Keywords: Competing risks; Process control; Scores; Type I censoring

1. Introduction

In some industrial situations, the value of a product or process characteristic may be censored. This censoring may arise for a variety of reasons. For example, the censoring may be due to time constraints, as in some life testing applications, or due to design, as in proof loading lumber to test breaking strengths, or due to competing risks, as in some strength testing applications where there are two or more failure modes. The observed response values may be right censored, left censored or interval censored. For example, in the testing of switches, a sample from production is subjected to a life test. However, owing to time constraints, the testing continues either until all the switches have failed or a set time has been reached. This yields right-censored responses. Another example that yields interval-censored data is the use of plug gauges to monitor sizes of holes. To measure the diameter of a hole, 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 is below the minimum specification. For process monitoring, the actual diameters of the few holes that fail are measured. Here all diameters within the specification limits are censored. An example of competing risks censoring occurs in the automotive industry and provides the main motivation for this work. An adhesive is used to attach a vinyl fabric to polyvinyl chloride foam backing for use in the interior of a car. The strength of the adhesive bond is a key characteristic and it is of interest to monitor the mean of the bond strength. However, during

Address for correspondence: Stefan H. Steiner, Department of Statistics and Actuarial Sciences, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada.
E-mail: shsteine@uwaterloo.ca

 0035–9254/01/50293

testing, the breaking strength of the bond is not always observed since the foam backing may fail first. For some units, the bond strength (in pounds per square inch) of the adhesive is observed, whereas for other units we only know that the adhesive bond strength is greater than the load level at which the foam failed. Similar situations that result in censored data occur in many other areas of application.

It is often desirable to monitor the mean (and/or variability) of a critical response. However, when responses are censored as in the examples described above, using statistical process monitoring methods that are naïvely based on the observed values (i.e. ignoring the censoring) is not effective. If the censoring is ignored, the resulting control chart will not have the desired operating characteristics. In particular, the chart will typically have a large false alarm rate. This problem arises because the distribution of the observed values is highly skewed owing to the censoring. However, even if we design the control chart using probability limits (based on the skewed distribution) to ensure the desired maximum false alarm rate, the resulting naïve chart will have low power to detect changes in the process mean.

In this paper we consider the case of competing risks censoring. The problem of monitoring using observations censored at a fixed level has been explored in conjunction with various distributional assumptions for the strength of the primary failure mode in Steiner and MacKay (1999, 2001). In the case of censoring at a fixed level, we set the desired censoring level on the basis of cost or time considerations. However, in the case of competing risks, the picture is more complicated because the censoring level varies. This means that the properties of any monitoring procedure will change if the behaviour of the competing risk(s) changes, and that changes in the primary failure mode may be masked by changes in the behaviour of the other failure mode(s). For example, in the adhesive strength example, we wish to use the results of strength tests on samples of units collected over time to monitor the performance of the production process. If the strength of the bond produced shows a substantial decrease we would like to know as soon as possible. Decreases in the bond strength signify serious problems since the adhesive may no longer successfully perform its function. However, changes in the strength of the foam (the competing risk) can make this monitoring difficult, especially if the mean foam strength decreases. In this paper we propose a monitoring procedure for a censored response that arises from competing risks. The procedure is designed so that it can also detect changes in the behaviour of the competing risk.

The paper is organized in the following manner. In Section 2 we define the notation and derive the likelihood for a censored sample. The fundamental idea underlying the methodology proposed is the conditional expected value (CEV) weight. These weights are defined in Section 3. Section 4 addresses the question of how to design control charts based on the CEV weights. It also compares the effectiveness of exponentially weighted moving average (EWMA) CEV control charts and Shewhart-type CEV charts. Section 5 discusses the implementation of the CEV chart in the context of competing risks. In that case, the effectiveness of the CEV chart depends on the probability of censoring. As a result, in addition to the EWMA CEV chart for the process mean, an EWMA CEV chart for the mean of the censoring mechanism is recommended. The use of EWMA CEV charts in the adhesive example is described in more detail in Section 6.

2. Preliminaries

To fix the notation, let T and C represent the strength of the primary failure mode and the strength of the censoring mechanism respectively. The censoring mechanism represents the effects of all other possible failure modes. We shall refer to the strength of the primary failure mode and the strength of the censoring mechanism as the process strength and censor strength respectively. We shall assume that the censoring is in the right-hand tail, although similar results may be obtained for other censoring patterns. If t and c represent the actual process and censor strengths for any measurement, then we observe (y, δ), where

$$y = \min(t, c)$$

and

$$\delta = \begin{cases} 1 & \text{if } t \leqslant c \text{ (failure due to the primary failure mode, i.e. not censored),} \\ 0 & \text{otherwise (failure due to some competing risk, i.e. censored).} \end{cases} \tag{1}$$

To model the process, assume stability for the moment, and denote the probability density functions of the process and censor strengths as $f(t)$ and $g(c)$ respectively. Similarly, denote the corresponding cumulative probability functions as $F(t)$ and $G(c)$ respectively. We assume that f and g are continuous distributions, although the methodology may be adapted to handle discrete observations. More importantly, we assume that T and C are statistically independent, an assumption that cannot be assessed with observations given by expression (1) (Lawless, 1982). Then, for an observation with observed strength y, the likelihood contribution can be written $g(y)\{1 - F(y)\}$ if the observation is censored and $f(y)\{1 - G(y)\}$ if the observation is not censored. This likelihood contribution can be summarized as

$$f(y)^{\delta}\{1 - F(y)\}^{1-\delta}\{1 - G(y)\}^{\delta} g(y)^{1-\delta}.$$

Thus, the likelihood for an independent sample of n observations is

$$\left[\prod_{i=1}^{n} f(y_i)^{\delta_i}\{1 - F(y_i)\}^{1-\delta_i}\right] \prod_{i=1}^{n} \{1 - G(y_i)\}^{\delta_i} g(y_i)^{1-\delta_i}. \tag{2}$$

Note that the likelihood can be written as the product of two factors, one dependent on f and one dependent on g. In what follows, we assume normality of the process strength and censor strength distributions, i.e. $T \sim N(\mu_t, \sigma_t^2)$ and $C \sim N(\mu_c, \sigma_c^2)$. Other distributional assumptions such as exponential, log-normal and Weibull are possible and do not change the procedure markedly.

As we shall see later, applying the monitoring procedure proposed requires estimates of the process parameters derived from some initial data. This corresponds to the usual set-up phase for $\bar{X}$- and R-charts where a set of 20 or so subgroups of data are collected to construct the control limits. Assuming normality, we may derive maximum likelihood estimates (MLEs) for the underlying process parameters by using expression (2) (Lawless, 1982). In a similar manner, assuming independence of failure modes, the censor mean and standard deviation can be estimated. For our example, using previously collected process data, and assuming normality, the MLEs for the in-control bond failure mean and standard deviation are 17.1 and 2.3 respectively. Similarly, the MLEs for the in-control mean and standard deviation of the strength of the foam (censor strength) are 18.9 and 3.9 respectively. Since these values are estimates of the in-control process we denote them $\mu_0 = 17.1$, $\sigma_0 = 2.3$, $\mu_{c0} = 18.9$ and $\sigma_{c0} = 3.9$. Of course, when the data are censored, less information about the process mean and variance than usual is available. The theoretical sample sizes needed to match the estimation precision obtained with uncensored data can be determined by using expected (Fisher) information (Steiner and MacKay, 1999).

3. Conditional expected value weights

The key idea used in the subsequent monitoring procedures is to replace all censored observations with their CEV. Using the notation from expression (1), we define the CEV weights for the process strength as

$$w = E(t|\mu_0, \sigma_0, y, \delta) = \begin{cases} y & \text{if } \delta = 1 \text{ (i.e. } t \leqslant c\text{)}, \\ \mu_0 + \sigma_0 \dfrac{\phi\{(y-\mu_0)/\sigma_0\}}{Q\{(y-\mu_0)/\sigma_0\}} & \text{if } \delta = 0 \text{ (i.e. } t > c\text{)} \end{cases} \tag{3}$$

where $\phi(z) = \exp(-z^2/2)/\surd(2\pi)$ and $Q(z) = \int_z^\infty \phi(x)\,\mathrm{d}x$ are the probability density function and the survivor function respectively of the standard normal distribution and μ_0 and σ_0 equal the in-control process mean and standard deviation respectively, estimated using maximum likelihood as discussed in the previous section. In the calculation of w, all uncensored observations are unchanged, whereas all censored observations are replaced by their corresponding expected value conditional on their observed censor strength (Lawless, 1982). For example, in the adhesive strength problem, suppose that when testing a particular unit the foam breaks at 15.1 lb in^{-2}. In other words, we observe $y = 15.1$ lb in^{-2}, and $\delta = 0$. Using expression (3) we replace this censored value with the adhesive strength that we would have expected if the observation had not been censored. In this case, since we estimated $\mu_0 = 17.1$ and $\sigma_0 = 2.3$, from expression (3) we obtain $w = 17.9$ lb in^{-2}. The values of expression (3) can be tabulated for any application, thus easing implementation of the proposed control chart.

The CEV weights have intuitive appeal since they have a direct physical interpretation. By design, when the process is in control, the expected value of CEV weights, as defined in expression (3), is μ_0, the process mean. In Appendix A we show that they can also be motivated through the score function. In fact, the CEV weights are a simple linear transformation of the scores. This direct connection with the score function shows that the sample average of the CEV weights forms a good test statistic to monitor the process. However, the distribution of the average of the CEVs, denoted $\bar{w}$, is complex and is affected by the number of censored observations that are expected, which in turn depends on μ_t, σ_t, μ_c and σ_c. The distribution of $\bar{w}$ is skewed, becoming more skewed as the proportion censored increases.

The idea behind the CEV weights works equally well for other distributions, such as Weibull or log-normal distributions, but in general they are not a simple linear transformation of the scores. For more details see Steiner and MacKay (2001).

Although the MLEs are recommended for estimating the in-control process parameters, they are not a good alternative to the CEV weights for the small subgroup sizes that are typically used in on-going process monitoring applications. This is because the MLEs have a large sampling variability for small samples, and if all observations in the sample are censored the MLEs are not defined. The probability that all observations are censored equals p^n, where n is the sample size. Assuming normality, $p = \Pr(\delta = 0) = Q\{(\mu_c - \mu_t)/\surd(\sigma_t^2 + \sigma_c^2)\}$. When p is large and the sample size is small, this is a non-negligible probability. In addition, the calculation of the MLEs requires iteration that may be considered onerous.

4. Control charts based on conditional expected value weights

We establish a monitoring procedure for the mean of a process that produces observations censored owing to competing risks in the usual manner, i.e. we follow the five steps given below.

Step 1: determine the subgroup size (denoted n), frequency of sampling, statistic to be monitored, type of control chart, etc.
Step 2: collect an initial sample of m subgroups. For this initial sample, derive maximum likelihood estimates for μ_0, σ_0, μ_{c0} and σ_{c0}. See Section 2.
Step 3: determine trial control limits. This can be accomplished by simulating the distribution of the average subgroup CEV weight.

(a) Randomly generate a large number of subgroups (say 10000) of the size n for the *actual* process and censoring strengths using the parameter estimates from step 2.
(b) For all observations, determine the observed strength y and whether the observation is censored or not, δ, using expression (1).
(c) Replace all observations with their CEV weight, given by expression (3), and determine all the subgroup averages $\bar{w}$.
(d) Determine appropriate control limits from the empirical distribution of $\bar{w}$.

Step 4: check the initial sample against the trial control limits to determine whether the process was in control. Remove out-of-control subgroups, and iterate through steps 2–4 if there is evidence of instability. Otherwise move to step 5.
Step 5: continue on-going monitoring with the trial control limits.

4.1. Determining control limits

In step 3(d), we determine appropriate upper and lower control limits based on the distribution of $\bar{w}$. We set the control limits to yield any desired average run length (ARL) when the process is in control. The ARL is defined as the average number of subgroups required before the chart signals. While the process is in control the ARL before a signal should be long; however, when the process is out of control we would ideally have a short ARL.

The appropriate choice of control limit depends on the type of chart. For a Shewhart $\bar{X}$-type chart we choose appropriate percentiles of the simulated distribution of $\bar{w}$. For example, to give the standard false alarm rate of 0.0027 (which corresponds approximately to an in-control ARL of 370) we set control limits to the 0.135 and 99.865 percentiles.

Consider again the adhesive strength example. Say we wish to monitor the process for decreases or increases in the bond strength using subgroups of size 12. We used step 3 of the design algorithm described at the start of Section 4 to determine appropriate control limits. Simulating the distribution of $\bar{w}$ using 10000 subgroups, the appropriate percentiles yield Shewhart control limits equal to 15.2 and 18.8.

Using the distribution of $\bar{w}$, we may also design other types of control chart. For example, an EWMA chart for the process mean is given by $Z_i = \lambda\bar{w}_i + (1 - \lambda)Z_{i-1}$ where $\bar{w}_i$ is the average CEV weight of the n units in the ith subgroup, and $Z_0 = \mu_0$. Appropriate EWMA control limits depend on the subgroup size and both the in-control censor mean and standard deviation. Generally, EWMA limits can be obtained by using a suitable discretization of the distribution of $\bar{w}$ and a Markov chain approximation (Steiner, 1998). However, a reasonable approximation may be given directly from appropriate percentiles of the distribution of $\bar{w}$. For example, when the EWMA smoothing parameter λ equals 0.25 choosing the 1st and 99th percentiles results in in-control ARLs of between 300 and 450 approximately. Similarly choosing 0.5 and 99.5 percentiles gives ARLs of between 1000 and 1400, and choosing 1.5 and 98.5 percentiles gives ARLs between 175 and 225. The results are relatively insensitive to the proportion censored (assuming that it lies between 0.25 and 0.75) and the relative magnitude of the standard deviation of the censoring mechanism. Generally, as either the proportion censored or the standard deviation of the censoring mechanism increases, the ARL decreases slightly. Using this idea, in the adhesive strength example, control limits for an EWMA chart with $\lambda = 0.25$ would be 15.6 and 18.4 if we wished an in-control ARL of around 400.

4.2. *Comparing exponentially weighted moving average and Shewhart conditional expected value control charts*

In Section 4.1 we illustrated the design of both Shewhart and EWMA control charts. In the context of CEV charts for competing risks it is of interest to compare the power of EWMA and Shewhart CEV control charts to detect changes in the process mean. The results, shown in Fig. 1, are derived by using simulation. Clearly, both Shewhart and EWMA CEV charts are good at detecting decreases in the process mean. However, the EWMA CEV chart is substantially better than the Shewhart chart at detecting increases in the process mean. Note that with both types of charts the ARL depends greatly on the in-control censoring proportion. With 75% censoring, the Shewhart chart's ARL does not peak when the process mean is 0 because of skewness in the distribution of the average of the CEV weights.

In the adhesive strength example the estimated proportion censored (in control) is 35%. At this level of censoring the performances of Shewhart and EWMA CEV charts to detect changes in the process mean are comparable. However, the ARLs of the EWMA chart are somewhat better, especially for detecting increases in the process mean. As a result, in further discussion of the example in Section 6, we design and implement EWMA CEV charts.

5. Monitoring the competing risks process

The distribution of the sample average of the CEV weights ($\bar{w}$) depends on the mean and standard deviation of the competing risk. As a result, changes in the censoring distribution can mask critical changes in the actual failure distribution. For example, a CEV control chart for the process mean alone would not do a good job in detecting increases in the process mean that are coupled with decreases in the censor mean. We explore this problem by using a

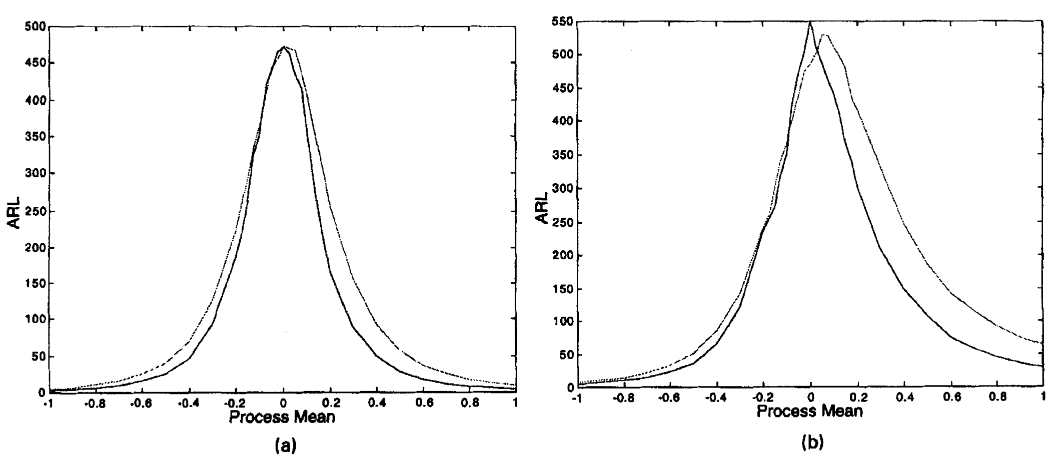

Fig. 1. Power of CEV process mean control charts (———, EWMA chart; ··········, Shewhart chart): (a) 25% censoring; (b) 75% censoring

standardized process where $\mu_0 = 0$ and $\mu_{c0} = \sigma_0 = \sigma_{c0} = 1$ in control. Fig. 2 shows contours of the ARL of the EWMA CEV chart for different combinations of process and censor means. The EWMA CEV chart is designed to have an ARL equal to 500 when the process is in control. In Fig. 2, we see that if the censoring mean decreases by 1 standard deviation unit to 0 and simultaneously the process mean increases by 0.5 standard deviation units to 0.5 the ARL is still close to 500. As a result, an EWMA CEV chart for the process mean will not be effective in detecting the change described.

For this reason, we suggest also monitoring the mean of the censoring mechanism. Of particular interest are decreases in the censor mean since then less information is available about the actual failure strength distribution. However, increases in the censor mean may also negatively affect the EWMA CEV control chart since, as the censoring rate decreases, the variability of the CEVs increases. This may result in larger false alarm rates.

Through symmetry, an EWMA CEV control chart for the censor mean can be designed using the procedure outlined in the previous section by reversing the roles of the process and censoring failure strengths. An example is given in Section 6.

The effectiveness of combining EWMA CEV control charts for both the process and the censor means is explored by showing their combined power in detecting simultaneous increases and decreases in the process mean and/or censor mean. For this illustration, consider CEVs determined from expression (3) in two cases. First, we look at a situation where the censoring rate is around 25% ($\mu_0 = 0$ and $\sigma_0 = \sigma_{c0} = \mu_{c0} = 1$). Second, we consider a case with a 75% censoring rate ($\mu_0 = 0$, $\sigma_0 = \sigma_{c0} = 1$ and $\mu_{c0} = -1$). In the 25% censoring case, using the control limit design algorithm suggested in Section 4.1, the lower and upper control limits for the EWMA CEV chart are set at -1.05 and 0.95 respectively for the process mean, and at -0.91 and 0.55 respectively for the censor mean.

The contour plots of the ARL in Fig. 3 show the effect of simultaneous changes in the process mean and censor mean. For example, when the censoring level is 25%, a process mean shift from 0 to 0.5 coupled with a censor mean shift from 1 to 0.5 yields an ARL of around 15. Fig. 3 suggests that the combined charts react relatively quickly to any change in

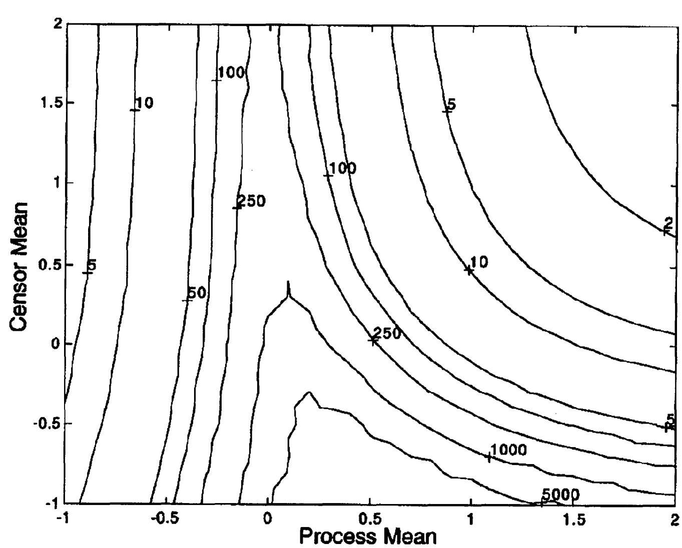

Fig. 2. ARL contour comparison for an EWMA CEV chart for the process mean ($\mu_0 = 0$; $\sigma_0 = \sigma_t = \sigma_c = 1$)

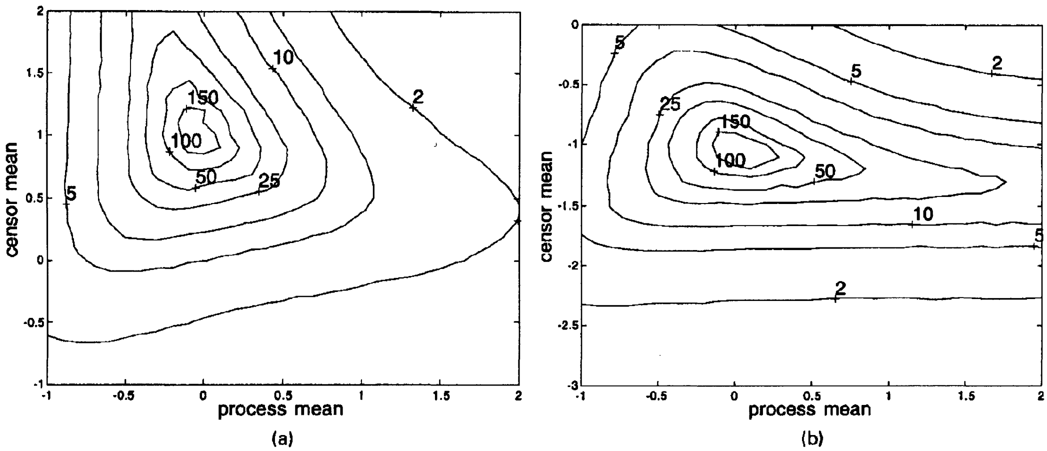

Fig. 3. ARL contour comparison for combined EWMA charts ($\sigma_t = \sigma_c = 1$): (a) $\mu_0 = 0$ and $\mu_{c0} = 1$ (approximately 25% censoring); (b) $\mu_0 = 0$ and $\mu_{c0} = -1$ (approximately 75% censoring)

the mean from the in-control condition. Of course, it is easier to detect changes in the process mean when the censoring rate is smaller. The ARLs in Fig. 3 are derived assuming independence between the two EWMA charts. This assumption is reasonable away from the in-control conditions where only one of the charts is likely to signal. The contour plots in Fig. 3 suggest that the combined EWMA charts quickly detect any changes in the process mean. The only exception is when the censoring proportion is large, where increases in the process mean that occur at the same time as small decreases in the censor mean will be difficult to detect.

Another possible solution to the potential masking effect of changes to the censor mean is to monitor the proportion censored. This approach is not explored in detail in this paper. Monitoring the proportion censored could be accomplished by using a cumulative sum chart on the number of censored units in each sample. Increases or decreases in the censoring proportion signify that changes in the process mean may be masked. An appropriate cumulative sum chart that takes into account the inherent discreteness in the number of censored observations may be designed by using the methodology described in Steiner *et al.* (1996).

6. Example

Consider the adhesive strength example described in Section 1. To ensure the quality of the adhesive bond produced we shall monitor this process for decreases or increases in the bond strength and/or foam strength by using a combined EWMA CEV chart for the process and censor means. For the monitoring, subgroups of size 12 were obtained each day.

As discussed in Section 2, the MLEs for the in-control strength of the bond and foam are estimated to be $\mu_0 = 17.1$, $\sigma_0 = 2.3$, $\mu_{c0} = 18.9$ and $\sigma_{c0} = 3.9$. At these levels, the censoring rate of the strength test while the process is in control is approximately 35%, which matches our experience. On the basis of this information, expression (3) gives the CEV sample weights w used in the charts for the process mean. Reversing the roles of the bond and foam strengths, expression (3) also gives the weights for the censor mean. Denote these weights v. An example of the calculation of the weights for the 12 units in the first subgroup is given in Table 1. For this subgroup, we calculate $\bar{w} = 17.5$ and $\bar{v} = 18.4$.

Table 1. Example of the weight calculation

Parameter	*Results for the following units:*											
	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*
y_i	15.1	18.3	16.7	19.1	13.9	13.5	14.3	16.3	14.5	15.2	14.3	20.0
δ_i	0	1	1	1	0	0	0	1	1	0	1	1
w_i	17.9	18.3	16.7	19.1	17.5	17.4	17.6	16.3	14.5	17.9	14.3	20
v_i	15.1	21.6	20.7	22.2	13.9	13.5	14.3	20.6	19.8	15.2	19.8	22.2

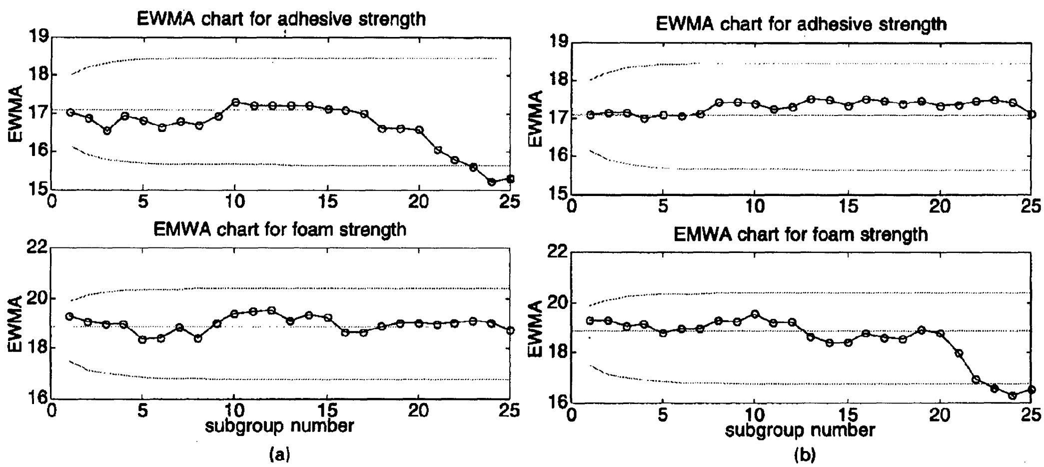

Fig. 4. Examples of EWMA CEV control charts for censored data: (a) decrease in adhesive strength after sample 20; (b) decrease in foam strength after sample 20

We shall use two EWMA charts: one designed to monitor the mean adhesive strength and the other designed to monitor the mean foam strength. On the basis of the simulation discussed in Section 4.1 we set the lower and upper control limits for the EWMA CEV chart for the mean adhesive strength at 15.6 and 18.4. Interchanging the roles of the adhesive strength and foam strength we can use the same design algorithm to determine appropriate control limits for an EWMA CEV chart for the foam strength. We obtain lower and upper control limits of 16.6 and 20.7.

Fig. 4 shows two examples of the combined EWMA CEV control charts for the process and censor means. In both cases, the first 20 samples represent an in-control condition. In Fig. 4(a) the effect of a decrease in the adhesive strength is simulated. After observation 20 the decrease in the adhesive strength is simulated through a decrease of 1 standard deviation unit in the adhesive strength mean. This decrease in the adhesive strength results in a number of signals in the EWMA CEV chart for the adhesive strength. In the two plots in Fig. 4(b) the effect of a simulated decrease in the breaking strength of the foam of 1 standard deviation unit at observation 20 is shown. This decrease in the censor strength is clearly evident in the signals for the EWMA CEV chart for the foam strength after observation 20.

7. Summarizing remarks and conclusions

In applications when observed data may be censored, traditional process monitoring approaches have undesirable properties such as large false alarm rates or low power. When the censoring comes from competing risks, adapted control charting procedures to monitor the process based on CEV weights are proposed. The CEV weights are equivalent to the likelihood scores and are thus optimal if the underlying distribution is normal. Since changes in the censoring level may mask changes in the process mean, we propose the simultaneous use of two EWMA CEV control charts to monitor both the process mean and the censor mean. The theoretical properties of such charts are shown to provide excellent protection from process changes.

Acknowledgements

This research was supported, in part, by the Natural Sciences and Engineering Research Council of Canada and General Motors of Canada. The authors would like to thank two referees and the Associate Editor for their valuable suggestions and comments that greatly improved the presentation of the material.

Appendix A: Justifying conditional expected value weights by using the score function

Utilizing the normality assumption, the CEV weights can also be justified by a likelihood argument. Consider the score, denoted s, which is defined as the first derivative of the log-likelihood (given by expression (2) with $n = 1$) with respect to μ evaluated at the in-control process mean and standard deviation, i.e.

$$s = \begin{cases} \dfrac{t - \mu_0}{\sigma_0^2} & \text{if } t \leqslant c \text{ (not censored)}, \\ \dfrac{\phi\{(c-\mu_0)/\sigma_0\}}{\sigma_0\, Q\{(c-\mu_0)/\sigma_0\}} & \text{if } t > c \text{ (censored)}. \end{cases} \tag{4}$$

These scores do not have a direct physical interpretation but are the basis for an optimal test statistic to detect small shifts in the process mean (Cox and Hinkley, 1974). Note that the scores given by expression (4) are equivalent, under a rescaling, to CEV weights given by expression (3). In particular, $w = \mu_0 + \sigma_0^2 s$ for both censored and uncensored observations. Thus, the scores are a linear translation of the CEV weights, and the CEV weights are also optimal. We know that

$$E[\partial\{\log(L)\}/\partial\mu] = 0$$

so it follows that the mean of w equals μ_0, independent of the mean and variance of the censor strength, when $\mu_t = \mu_0$ and $\sigma_t = \sigma_0$. However, the standard deviation of the CEVs depends on the mean and variance of the censor strength. As either the censor mean or variance changes in such a way that the probability of censoring increases, the standard deviation of w decreases.

References

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

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

Steiner, S. H. (1998) Grouped data exponentially weighted moving average control charts. *Appl. Statist.*, **47**, 203–216.

Steiner, S. H., Geyer, P. L. and Wesolowsky, G. O. (1996) Grouped data sequential probability ratio tests and cumulative sum control charts. *Technometrics*, **38**, 230–237.

Steiner, S. H. and MacKay, R. J. (1999) Monitoring processes with highly censored data. *J. Qual. Technol.*, to be published.

——— (2001) Detecting changes in the mean from censored lifetime data. In *Proc. Int. Wrkshp Intelligent Statistical Quality Control*. To be published.