Quality Technology &
Quantitative Management
Vol. 3, No. 1, pp. 111-125, 2006

QTQM

A Semi-Bayesian Method for Shewhart Individual Control Charts

M. B. (Thijs) Vermaat and Ronald J. M. M. Does

Institute for Business and Industrial Statistics,
IBIS UvA, University of Amsterdam,
Amsterdam, The Netherlands
(*Received* February 2004, *accepted* June 2005)

Abstract: Shewhart control limits for individual observations are traditionally based on the average of the moving ranges. The performance of this control chart behaves quite well if the underlying distribution is normal and the sample size is greater than 250. Under non-normality it is recommended to use control charts based on non-parametric statistics. The drawback of these individual control charts is that at least 1,000 observations are needed to obtain appropriate results. In this paper we propose an alternative individual control chart which behaves quite well under non-normality for moderate sample sizes in the range of 250 through 1,000 observations. To apply this control chart one starts with an initial guess for the density function of the characteristic under study. Based on this initial guess and the observed data a density function can be derived by means of an approximation with Bernstein polynomials. The in-control and out-of-control performance of the proposed control chart and the traditional control charts are studied by simulation. If the initial guess is appropriate, then for non-normal data and moderate sample sizes in the order of 250 through 1,000 observations, the new method performs better than the individual control charts based on the average of the moving ranges or based on non-parametric statistics. So for these sample sizes we have tried to close the gap.

Keywords: Bernstein polynomials, non-parametrics, quality control, semi-Bayesian, statistical process control.

1. Introduction

1.1. Background

In 1924 Shewhart developed the control chart which is an outstanding instrument to facilitate process control (cf. [25]). At least 20-25 initial samples of about five units each are needed before we can estimate with sufficient confidence the control limits. Much larger data sets (at least 300) are needed if the sample size equals one (cf. [19]). For this situation the so called individual control chart has developed (cf. [10]). For the estimation of the parameters (μ and σ) of the control limits (cf. Section 2), Duncan [10] uses the sample mean and the average of the moving ranges, respectively. In the literature also other estimators for the control limits are proposed, cf. [7] and [23]. Up-to-date books on statistical process control are [9], [17], [20], and [32].

1.2. Non-Normality and Control Charts

Usually, one assumes that the underlying distribution function of the quality characteristic is normal. In practice we have often characteristics under study which are not normally distributed (e.g. skewed distributed). A number of authors has pointed out that Shewhart charts for subgroup means work well irrespective of the measurements are normally distributed or not ([24] and [31]). However, the behavior of the traditional control chart for individual measurements is seriously affected by departures from normality (see e.g. [3], [27], and [35]). Furthermore, non-normal data are quite common in SPC applications. In situations with a large number of observations it may be possible to subgroup the data to avoid an individual chart. If subgrouping is not possible and non-normality is evident, two alternatives are:

(1) Transform the data to normality or

(2) Modify the control limits based on a suitable (parametric or non-parametric) model for the data (see [21]).

Possibilities for the first alternative are to use the Box-Cox transformation or Johnson transformation. Statistical software programs like Minitab may carry out such a transformation almost automatically. In this paper we will concentrate on the second alternative. If large data sets are available, an attractive approach to parametric statistical inference is to use these large data sets to study the distributional form. There is a huge statistical literature available in the areas of goodness-of-fit and parametric modelling (cf. [8]).

In a recent paper of Vermaat *et al*. [30] individual control charts based on empirical quantiles, kernel estimators, and extreme-value theory are studied. It turned out that these alternative individual control charts were quite robust against deviations from normality if the number of observations is at least 1,000. Control charts for non-normal data are further studied by [3], [6], [26], and [27] among others.

Bayesian methods for the estimation of control limits are also studied in the literature of control charts. In [28] a discussion of the application of empirical Bayes for the estimation of a distribution of the characteristic under study is given (see also [29], [34], and [16]. Other papers which use Bayesian methods in quality control are [5], [11], [12], [13], and [14].

1.3. Semi-Bayesian Approach

For sample sizes till 1,000 there is no good alternative under non-normality (cf. [19] and [30]). In this paper we propose a new method for the estimation of the control limits for moderate sample sizes (i.e. for sample sizes in the range of 250 through 1,000). This method makes use of an initial guess of the underlying distribution function and Bernstein polynomials. The Bernstein polynomial is a linear function of order statistics with smooth weight functions and may be used to estimate quantile functions (cf. [18]). In [4] Bernstein polynomials are used for non-parametric density estimation. In [2] some improvements are studied. We compare the individual control charts based on Bernstein polynomials with the traditional Shewhart control chart for individuals with control limits based on the average of the moving ranges and with the empirical quantile control chart studied in [30]. The statistical performance of these methods will be studied by an extensive simulation study.

The paper is organized as follows. In the next section the different methods for the estimation of the control limits are given. We illustrate the different methods by means of a real life example. To compare the different methods a simulation study is presented. Finally, we conclude with our findings.

2. The Traditional Individual Control Chart based on Moving Ranges

The traditional Shewhart individual control chart has control limits defined by

$$\text{UCL} = \mu + \Phi^{-1}(1 - \frac{\alpha}{2})\sigma \quad \text{and} \quad \text{LCL} = \mu + \Phi^{-1}(\frac{\alpha}{2})\sigma \,,$$

where Φ^{-1} is the standard normal quantile function, and where μ is the mean and σ is the standard deviation of the normal distribution function F. Level α is the false alarm rate (e.g. for α = 0.0027 we obtain $\Phi^{-1}(1-\frac{\alpha}{2})=3$). Typically, μ and σ are unknown in practice. However, we assume that they can be estimated via a Phase I sample X_1, ..., X_k of independently and identically distributed random variables. Recall from [33], that Phase I refers to the retrospective analyze phase and Phase II refers to the monitoring phase. Estimators of μ and σ are usually the sample mean $\bar{X}_k = \sum_{i=1}^{k} X_i/k$ and the average of the moving ranges

$$\overline{\mathrm{MR}}_k = \frac{1}{k-1}\sum_{i=2}^{k} | X_i - X_{i-1} |$$

respectively, see [10]. The last estimator is scaled by $d_2(2) = 2/\sqrt{\pi}$ to obtain an unbiased estimator for σ. With these estimations for μ and σ we obtain the traditional Shewhart individual control chart with control limits

$$\widehat{\mathrm{UCL}} = \bar{X}_k + \Phi^{-1}(1-\frac{\alpha}{2})\frac{\overline{\mathrm{MR}}_k}{d_2(2)} \quad \text{and} \quad \widehat{\mathrm{LCL}} = \bar{X}_k + \Phi^{-1}(\frac{\alpha}{2})\frac{\overline{\mathrm{MR}}_k}{d_2(2)}. \tag{1}$$

2.1. The Average Moving Range Control Chart

In [23] a more exact version of the traditional Shewhart individual control chart has derived. With $V_k = \overline{MR}_k/d_2(2)$ they approximate the distribution of $W_k^2 = (V_k/\sigma)^2$ by $\tau^2\chi^2(\nu)/\nu$, where

$$\tau=\sqrt{\mathrm{Var(W}_k)+1}\,,$$

$$\nu = \frac{1}{2}(1+\frac{1}{\mathrm{Var(W_k)}}),$$

and $\chi^2(\nu)$ is the chi-square distribution with ν degrees of freedom. Using this approximation it follows that for $t > k$, $(X_t - \bar{X}_k)/(\overline{\mathrm{MR}}_k/d_2(2))$ is approximately distributed as $\frac{\sqrt{1+1/k}}{\tau}$ times a Student's t-distribution with ν degrees of freedom. The Average Moving Range (AMR) control chart has control limits defined by

$$\widehat{\mathrm{UCL}}_{\mathrm{AMR}} = \bar{X}_k + \frac{\sqrt{1+1/k}}{\tau} t(1-\tfrac{\alpha}{2};\nu)\frac{\overline{\mathrm{MR}}_k}{d_2(2)} \tag{2}$$

and

$$\widehat{\mathrm{LCL}}_{\mathrm{AMR}} = \bar{X}_k + \frac{\sqrt{1+1/k}}{\tau} t(\tfrac{\alpha}{2};\nu)\frac{\overline{\mathrm{MR}}_k}{d_2(2)}, \tag{3}$$

where $t(p,\nu)$ denotes the p-quantile of a t-distribution with ν degrees freedom and $d_2(2) = 2/\sqrt{\pi}$. This AMR control chart is more accurate than the traditional Shewhart individual control chart in (1): i.e. the rates of the traditional Shewhart individual control chart of falsely signalling an out-of-control situation are much larger than intended, see [22].

2.2. Empirical Quantile Control Chart

In this subsection the control limits of the empirical quantile control chart are defined as in [30]. A natural estimator of the q-quantile of the distribution function F is the empirical quantile $\hat{F}_k^{-1}(q)$, which is defined as

$$\hat{F}_k^{-1}(q) = \inf\{x \mid \hat{F}_k(x) \geq q\},\ 0 < q < 1,$$

where $\hat{F}_k$ is the empirical distribution function that puts mass $1/k$ at each X_i, $1 \leq i \leq k$, i.e.

$$\hat{F}_k(x) = \frac{1}{k}\sum_{i=1}^{k} I_{\{X_i \leq x\}}, -\infty < x < \infty,$$

With I the indicator function, i.e. $I_{\{x \leq y\}}$ equals 1 if $x \leq y$ holds and 0 otherwise. Hence, an obvious estimator of the upper control limit based on the empirical quantile (EQ) is

$$\widehat{\mathrm{UCL}}_{\mathrm{EQ}} = \hat{F}_k^{-1}(1-\tfrac{\alpha}{2}) = X_{\left(\left\lceil (1-\frac{\alpha}{2})k \right\rceil\right)}, \tag{4}$$

with $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(k)}$ denoting the order statistics of the initial sample $X_1, \ldots, X_k$ and $\lceil y \rceil$ the smallest integer not smaller than y. The lower control limit estimated by the empirical quantile is defined by

$$\widehat{\mathrm{LCL}}_{\mathrm{EQ}} = X_{\left(\left\lfloor \frac{\alpha}{2}k+1 \right\rfloor\right)}, \tag{5}$$

where $\lfloor y \rfloor$ denotes the largest integer not larger than y.

2.3. A Control Chart based on a Bernstein Approximation

In this subsection we introduce a semi-Bayesian method for estimating a density function f. From this we derive the corresponding cumulative distribution function. This method is introduced in [2, 5]. For technical details we refer to Appendix A.

Assume that the density function f is continuous and strictly positive on the interval (a,b) and that F is the corresponding cumulative distribution function. We choose an a priori density function ψ as an initial guess for the density function f. Based on the corresponding continuous cumulative distribution function Ψ, the order statistics $X_{(1)}, \ldots, X_{(k)}$ are transformed on $[0,1]$, by $Y = \Psi(X)$, such that

$$Y_{(0)} = 0, \quad Y_{(i)} = \Psi(X_{(i)}), \quad i = 1,\ldots,k, \quad Y_{(k+1)} = 1.$$

Because $P(Y \leq y) = P(\Psi(X) \leq y) = P(X \leq \Psi^{-1}(y)) = F(\Psi^{-1}(y))$ we have that $B = (F(\Psi^{-1}))^{-1}$ is the quantile function of the random variable $Y = \Psi(X)$. This quantile function is estimated by the so called Bernstein polynomials $B_k^{(m)}(p)$, where m is a smoothing parameter and $0 < p < 1$, see Appendix A. It follows that $\hat{F}_{k,BA}^{(m)} = B_k^{(m)\,-1}(\Psi)$ is an estimate for F.

The estimators for the upper and lower control limits based on the Bernstein approximation (BA) with a false alarm rate α are now given by

$$\widehat{\mathrm{UCL}}_{\mathrm{BA}} = \hat{F}_{k,BA}^{(m)^{-1}}(1-\tfrac{\alpha}{2}), \tag{6}$$

and

$$\widehat{\mathrm{LCL}}_{\mathrm{BA}} = \hat{F}_{k,BA}^{(m)^{-1}}(\tfrac{\alpha}{2}), \tag{7}$$

where $\hat{F}_{k,BA}^{(m)^{-1}}$ is defined by

$$\hat{F}_{k,BA}^{(m)^{-1}} = \Psi^{-1}(B_k^{(m)}(p)),$$

cf. formulas (9) and (10) in Appendix A. The control chart based on these control limits is denoted as the BA control chart.

The method described above needs an initial guess ψ. The initial guess has to be chosen a priori. This assumption is not necessarily unrealistic, cf. [1]. Based on process knowledge we can make an initial guess for the distribution of the process characteristic. The described method is to 'fine tune' the initial guess.

For the Bernstein approximation the smoothing parameter m has to be chosen. In Appendix A it is shown that we may take $m = 5.2\sqrt{k}$ and round this off to the nearest integer.

3. Real Life Example

The Bernstein approximation control chart seems to be rather cumbersome. However, with the use of computers and software, methods like this can be used by practitioners. To show this we apply the Bernstein approximation control chart to a real life example.

The collected data concern the part of a printer that squirts the ink. The ink is spurt through a groove. The relevant quality characteristic to be measured is the depth of the groove. We will study 535 observations of this depth. For reasons of confidentiality the measurements are multiplied by a constant. By studying the data, we found out that for 15 measurements the grooves were cut by a supplier who uses a different method to cut the grooves. In order to obtain stability of the process we have eliminated these data points. We will study the remaining data set, which consists of 520 measurements. For the calculation of the control limits we use the estimators described in the previous section. We found that the control limits based on the average moving ranges are (cf. formulas (2) and (3))

$$\widehat{\mathrm{UCL}}_{\mathrm{AMR}} = 0.210,$$

$$\widehat{\mathrm{LCL}}_{\mathrm{AMR}} = 0.148,$$

and for the empirical quantiles method we obtain (cf. formulas (4) and (5))

$$\widehat{\mathrm{UCL}}_{\mathrm{EQ}} = 0.201,$$

$$\widehat{\mathrm{LCL}}_{\mathrm{EQ}} = 0.131.$$

To obtain the control chart based on Bernstein polynomials we use the normal distribution as initial guess to estimate the control limits (cf. formulas (6) and (7)). From formulas (9) and (10) it is clear that we have to estimate the parameters of the initial guess and the smoothing parameter m. Because we have chosen the normal distribution as initial guess its parameters (μ and σ) can be easily estimated by the sample mean and sample standard deviation of the 520 observations. Because $5.2\sqrt{k} = 118.6$, we take $m = 119$. Hence the BA control limits are

$$\widehat{\mathrm{UCL}}_{\mathrm{BA}} = 0.207,$$

$$\widehat{\mathrm{LCL}}_{\mathrm{BA}} = 0.136.$$

In Figure 1 the data with the different control limits are drawn in one graph. Figure 2 shows that the distribution of the data is slightly skewed to the left. The autocorrelogram of the data in Figure 3 indicates that the data are not significantly correlated.

Note that because we use $\alpha = 0.0027$ and $n = 520$ the EQ control chart takes the largest and smallest observation as upper control limit and lower control limit, respectively. Hence, all observations are between these two control limits. The AMR control chart shows

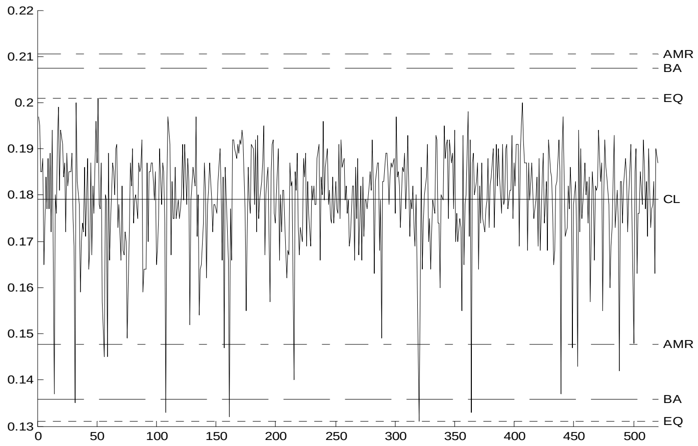

Figure 1. Depth of groove data and control limits based on the average moving ranges (AMR), the empirical quantiles (EQ), and the Bernstein approximation (BA). CL is the center line and represents the process average.

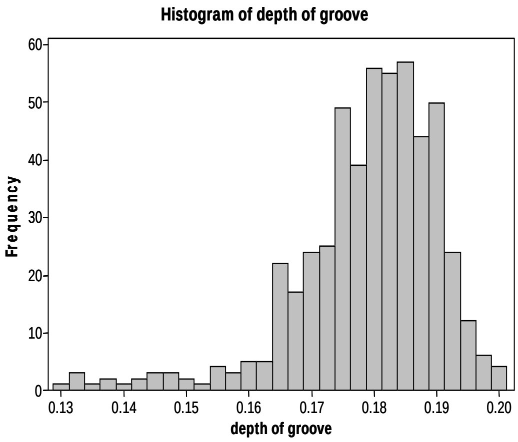

Figure 2. The histogram for the depth of groove data.

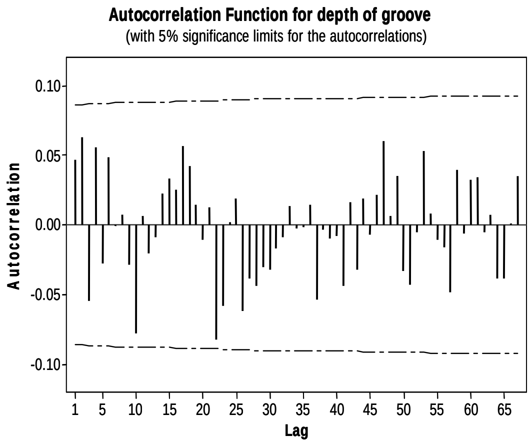

Figure 3. The autocorrelogram for the depth of groove series.

15 out-of-control signals and its control limits are symmetric with respect to the center line. The BA control chart signals only 5 out-of-control signals. To end with an in-control process, we eliminated these 5 out-of-control observations and we recalculated the control limits. Note that to recalculate the control limits we have to run the program once more. We iterated this process until we got no more out-of-control signals based on the BA control limits. Totally, we eliminated in this way 5 extra out-of-control points. For the remaining 510 observations we obtained as BA control limits $\widehat{UCL}_{BA} = 0.206$ and $\widehat{LCL}_{BA} = 0.145$. To obtain an in-control process based on the AMR control chart we have to eliminate in total 19 out-of-control points. The AMR control limits for the remaining 501 observations are $\widehat{UCL}_{AMR} = 0.207$ and $\widehat{LCL}_{AMR} = 0.154$.

In Phase II (monitoring phase) the control limits may be updated regularly as new data become available. This is done in the same way as in Phase I. In Chapter 5 of [8] an activity plan is given for the implementation and maintenance of a control chart.

4. Design of the Simulation Study

In order to study the performance of the control charts introduced in the previous sections we have conducted an extensive simulation study for some choices of the distribution function F and the sample size k. Based on the Phase I training sample X_1, …, X_k the UCL and LCL are estimated. The performance of each control chart is measured by calculating the average and standard deviation of the run length – ARL and SDRL respectively – in Phase II. We follow [30] in the calculations of the ARL and SDRL. They denote by $\hat{p}_k(X_1, \ldots, X_k)$ the conditional probability given the training sample, that a new independent random variable X from the same distribution F exceeds the upper control limit or is below the lower control limit. The average run length is

$$\text{ARL} = \text{E}\frac{1}{\hat{p}_k(X_1,\ldots,X_k)}$$

and the standard deviation of the average run length

$$\text{SDRL} = \left[2\text{E}\left(\frac{1}{\hat{p}_k(X_1,\ldots,X_k)}\right)^2 - \left(\text{E}\frac{1}{\hat{p}_k(X_1,\ldots,X_k)}\right)^2 - \text{E}\frac{1}{\hat{p}_k(X_1,\ldots,X_k)}\right]^{\frac{1}{2}} \quad (8)$$

Since these ARL and SDRL cannot computed directly, 10,000 training samples x_1, …, x_k are drawn. For each training sample we calculate $1/\hat{p}(x_1,\ldots,x_k)$ and $1/\hat{p}(x_1,\ldots,x_k)^2$. Averaging $1/\hat{p}(x_1,\ldots,x_k)$ over the 10,000 training samples gives the ARL. In more or less the same way we calculate the SDRL.

We study the control charts in an in-control and several out-of-control situations. For the comparison we distinguish different shifts in the mean of size $\delta\sigma$, where δ ranges from 0 through 5 and σ is the standard deviation of the studied distribution. With δ = 0 we have of course the in-control situation. Given the 10,000 training samples $x_1,\ldots,x_k$ we calculate for each shift the ARL and SDRL. Based on a false alarm rate α equal to 0.0027, which corresponds with the traditional 3σ-limits of Shewhart individual control charts, we expect that the ARL and SDRL are around 370.

The simulations are carried out for six different distributions, i.e. normal, Student's t, logistic, exponential, chi-square, and Weibull distribution and for sample sizes k equal to 250, 500, and 1,000. Seventeen shifts in the mean of size $\delta\sigma$ are used in the range of 0 (0.25) 3.5, 4, and 5.

4.1. Simulation Design for the Bernstein Approximation Control Chart

As mentioned before, the BA control chart has developed to 'fine tune' an initial guess for the underlying density function. Based on physical properties of a product and/or on experience it is generally known whether the characteristic under study has approximately a normal distribution or some skewed distribution. In our simulation we use this fact in order to differentiate between a normal and a skewed distribution.

To distinguish two moderately deviations from a normal distribution, we use Student's t-distribution with 30 degrees of freedom and the logistic distribution with scale parameter equal to 1. The BA control chart performance will be illustrated by using these two distributions and a normal distribution. The initial guess will be in these three cases the normal distribution. Based on the training sample $x_1,\ldots,x_k$ the μ and σ of the normal distribution (i.e. the initial guess distribution) are estimated by the sample quantities of the parameters.

Secondly, we study three deviations from a gamma distribution, i.e. the standard exponential, the chi-square distribution with one degree of freedom, and the Weibull distribution with shape parameter equal to 1 and scale parameter equal to 2. The initial guess ϕ will be in these three cases a gamma density. The two parameters of this gamma density are estimated from the first two moments of the training sample (cf. [15]).

4.2. Discussion of the Simulation Results

The results of the simulations are given in two complementary figures. In the first figure the ARL is depicted for the different estimation methods. In the second figure the SDRL is depicted.

The results of the simulations are given in Figures 4 through 12 showing the ARL and SDRL respectively. The shift in the mean is put on the horizontal axes and the ARL (SDRL) on the vertical axes. For sample sizes 250 and 500 only the results of the AMR and the BA control charts are given. For sample size equal to 1,000 the graphs are added with the results of the EQ control chart. For smaller sample sizes its behavior is not satisfactory (cf. [30]). Furthermore, we have added in all figures the theoretical (TH) ARL and SDRL (based on the exact exceedance probability p for the given distribution). This was easy because the exact distribution was known in the simulation study.

In Figure 4 the results for the normal distribution function are given for sample size $k = 250$ and with the normal density as initial guess. We see that the BA control chart is closer to the theoretical ARL and SDRL than the AMR control chart. If the sample size k increases, we observe that the performance lines of both control charts converge to the theoretical line. In Figure 5 the control limits of the EQ control chart are equal to the second largest respectively second smallest observation. The performance of all charts is comparable.

In Figure 6 the results for the t-distribution with 30 degrees of freedom and sample size $k = 500$ are given. We see that the BA control chart is slightly better than the AMR control chart.

Figure 7 shows the results for the logistic distribution. Although the logistic distribution looks like a normal distribution the performance of the AMR control chart and BA control chart are very bad for $k = 250$. This is due to the fact that the tail behavior of the logistic distribution is very different from the tail behavior of a normal distribution. Even for large values of k the behavior of both control charts is bad compared with the theoretical one. For $k = 1{,}000$ we observe a more regular behavior of the ARL (and SDRL) for the EQ control chart, see Figure 8.

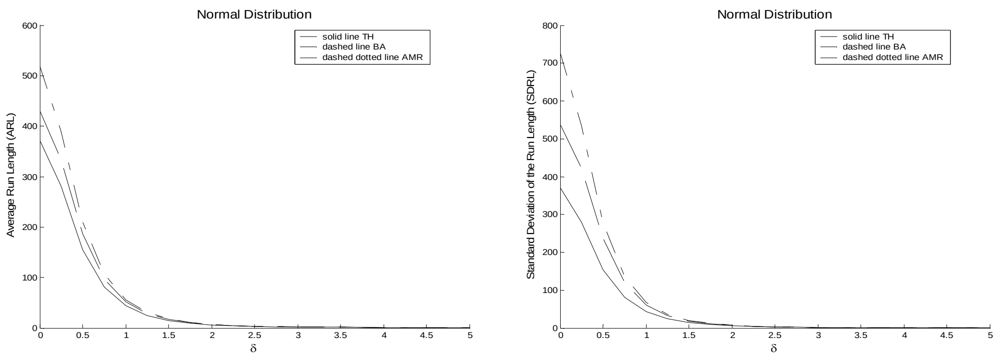

Figure 4. The ARL and the SDRL for k = 250 based on an underlying normal distribution.

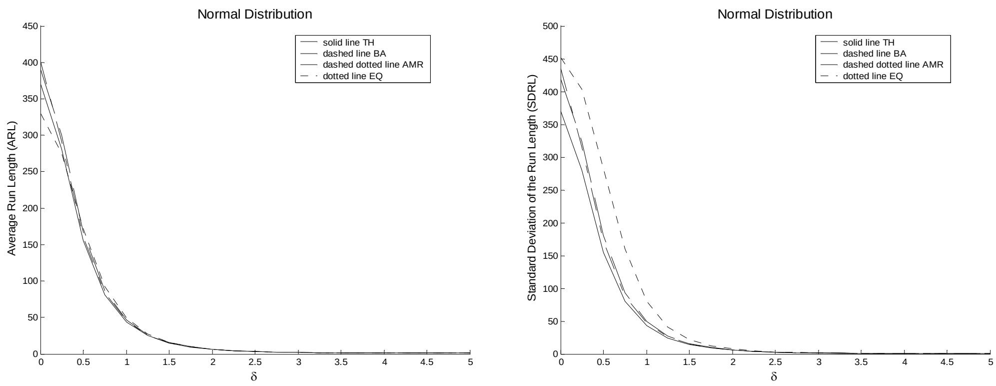

Figure 5. The ARL and the SDRL for k = 1,000 based on an underlying normal distribution.

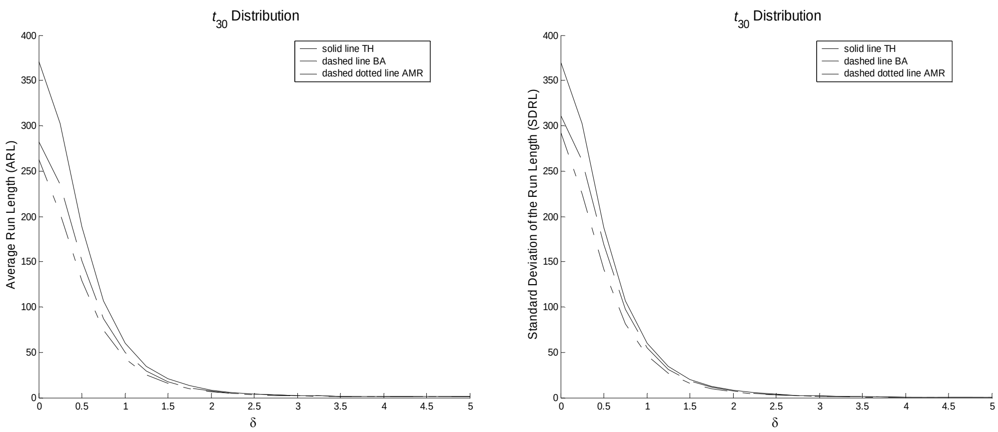

Figure 6. The ARL and the SDRL for k = 500 based on an underlying *t*-distribution with 30 degrees of freedom.

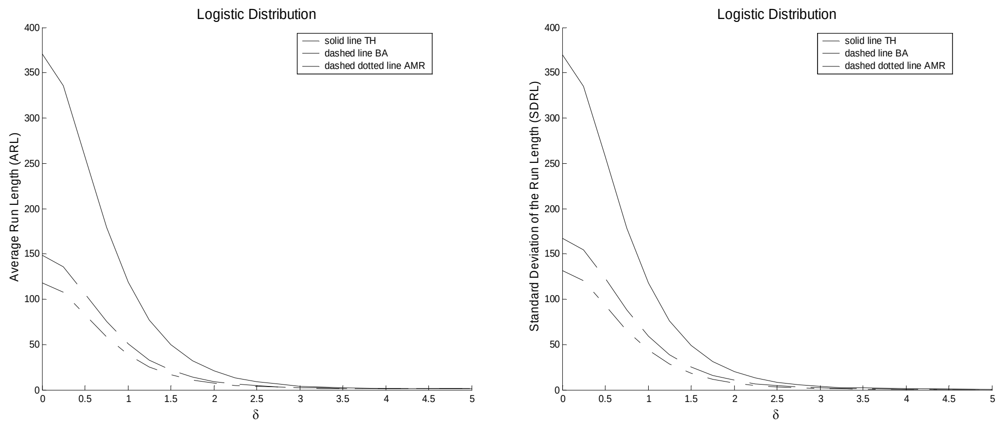

Figure 7. The ARL and the SDRL for k = 250 based on an underlying logistic distribution.

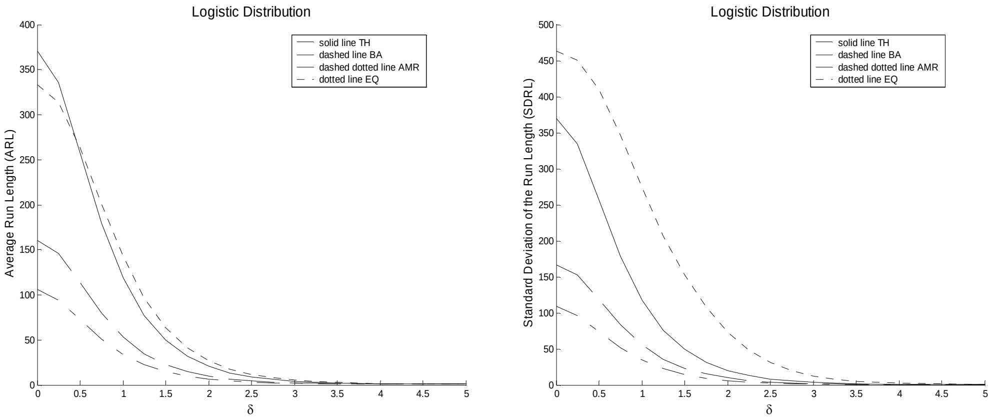

Figure 8. The ARL and the SDRL for k = 1,000 based on an underlying logistic distribution.

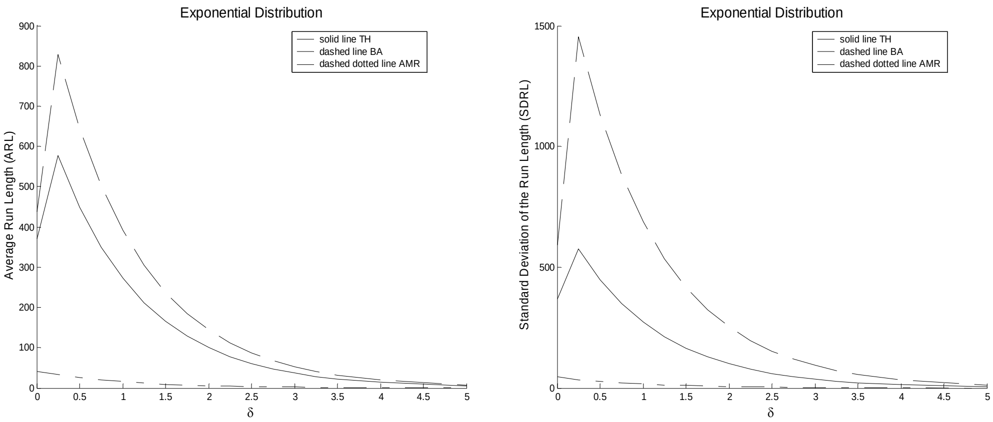

Figure 9. The ARL and the SDRL for k = 250 based on an underlying exponential distribution.

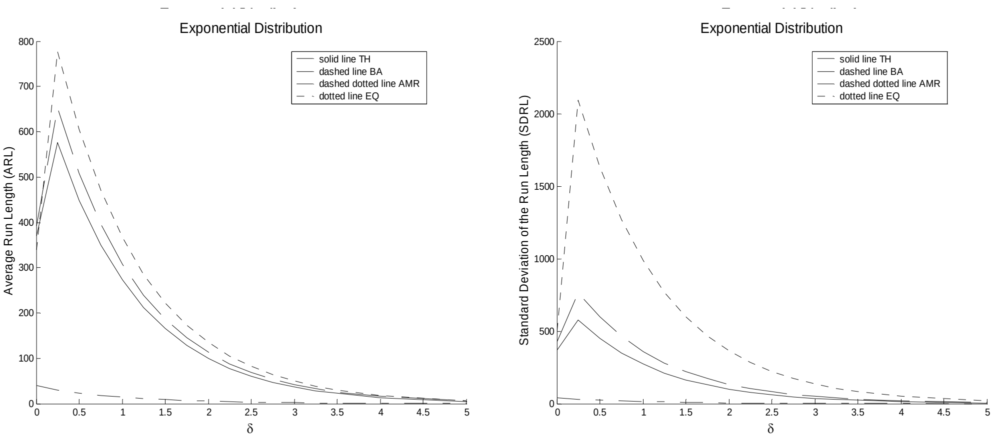

Figure 10. The ARL and the SDRL for k = 1,000 based on an underlying exponential distribution.

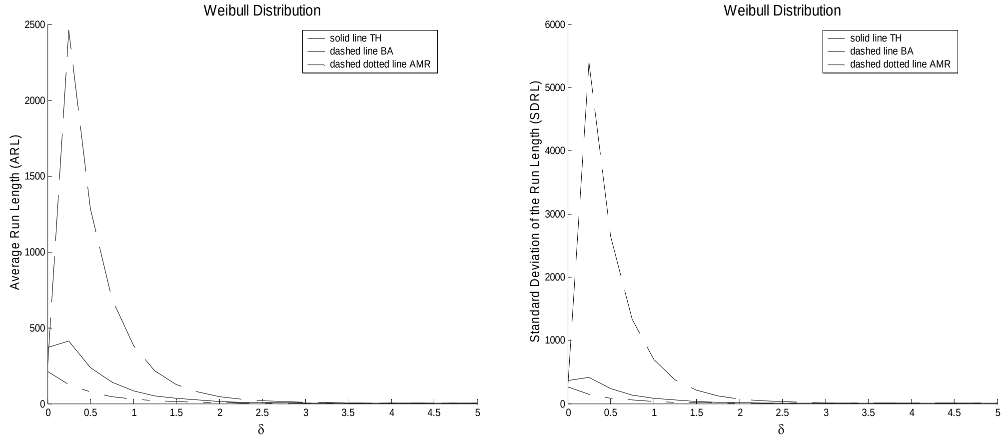

Figure 11. The ARL and the SDRL for k = 250 based on an underlying Weibull distribution.

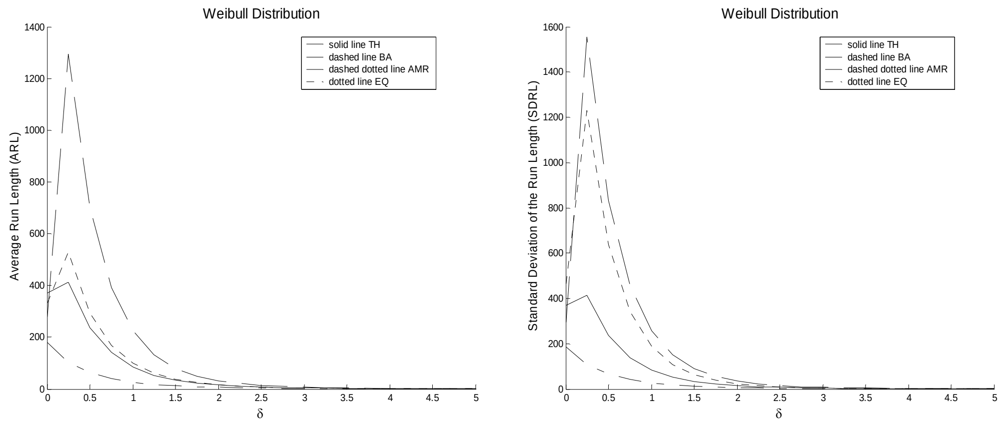

Figure 12. The ARL and the SDRL for k = 1,000 based on an underlying Weibull distribution.

In Figures 9 and 10 we show the results for the exponential distribution function for sample sizes $k = 250$ and $k = 1{,}000$. We see that for a training sample with $k = 1{,}000$ the performance of the BA control chart is much better than for a training sample with $k = 250$. The same pattern holds for a chi-square distribution. This can be explained by the fact that the estimators of the parameters of the initial guess, i.e. the gamma distribution, need a lot of data to be accurate and unbaised, (cf. [15] p. 358). More data lead to a more accurate initial guess, which implies a better performance of the BA control chart. We also observe in Figures 9 and 10, that the BA and theoretical control chart have a maximum in the ARL and SDRL around $\delta = 0.25$. This maximum is due to the fact, that these control charts estimate an LCL within the support of the exponential distribution. If the process shifts towards the UCL when δ increases, the LCL is of course harder to violate and the UCL easier. Since the density of the exponential is larger near the LCL than near the UCL, this causes that the probability of an alarm decreases and hence the ARL increases for small δ. Because the estimates of the LCL in the AMR control chart typically fall outside the support of the exponential distribution, this control chart does not show this phenomenon. Note that the behavior of the AMR control chart for $\delta = 0$ is really bad.

In Figure 11 and 12 the results for a Weibull distribution have drawn for sample sizes $k = 250$ and $k = 1,000$. We see that the results for both the BA and AMR control charts for this distribution are bad. For the BA control chart this is due to the fact that the initial guess (i.e. a gamma distribution) differs too much from the Weibull distribution as theoretical distribution. In this case only the performance of the EQ control chart is satisfying for the ARL with $k = 1,000$.

5. Conclusion

In this study we have applied a semi-Bayesian method for the estimation of the control limits of a Shewhart individual control chart. The control chart based on the average of the moving ranges performs well for normally distributed observations. If the observations are no longer normally distributed the behavior becomes quite bad. In this situation non-parametric (e.g. empirical quantiles) control charts perform quite well for more than 1,000 sampled observations. In situations of moderate sample sizes (i.e. 250 to 1,000) we suggest to use the Bernstein approximation method. The purpose of this method is to 'fine tune' an initial guess of the underlying density function. With this method the control chart is robust against small violations of normality. We have also shown that if the initial guess is a reasonable approximation of the underlying distribution, the performance of the new control chart is better than the AMR control chart. If on the other hand the initial guess is a bad approximation of the underlying distribution the results of the new control chart are not good enough, but nevertheless better than the AMR control chart.

Acknowledgements

The authors thank the referees, the Associate Editor, and the Editor for their useful comments and suggestions, which led to substantially improvements of the paper.

References

1. Albers, C.J. (2003). *Distributional Inference: The Limits of Reason*, Ph.D. Thesis, University of Groningen, The Netherlands.
2. Albers, C.J. and Schaafsma, W. (2003). Estimating a density by adapting an initial guess. *Computational Statistics & Data Analysis*, 42, 27-36.
3. Borror, C.M., Montgomery, D.C., and Runger, G.C. (1999). Robustness of the EWMA control chart to non-normality. *Journal of Quality Technology*, 31, 309-316.
4. Bruin, de R., Salomé, D. and Schaafsma, W. (1999). A semi-Bayesian method for nonparametric density estimation. *Computational Statistics & Data Analysis*, 30, 19-30.
5. Carter, P.L. (1972). A Bayesian approach to quality control. *Management Science*, 18(11), 647-656.
6. Chou, Y., Polansky, A.M. and Mason, R.L. (1998). Transforming non-normal data to normality in statistical process control. *Journal of Quality Technology*, 30, 133-141.
7. Cryer, J.D. and Ryan, T.P. (1990). The Estimation of Sigma for an X Chart: $\overline{MR}/d_2$ or S/c_4? *Journal of Quality Technology*, 22, 187-192.
8. D'Agostino, R.B. and Stephens, M.A. (1986). *Goodness-of-Fit Techniques*, Marcel Dekker Inc., New York.
9. Does, R.J.M.M., Roes, K.C.B. and Trip, A. (1999). *Statistical Process Control in Industry*, Kluwer Academic, Dordrecht, The Netherlands.
10. Duncan, A.J. (1986). *Quality Control and Industrial Statistics*, Fifth Edition, Irwin,

Homewood, Illinois.

11. Feltz, C.J. and Shiau, J-J. H. (2001). Statistical process monitoring using an empirical bayes multivariate process control chart. *Quality and Reliability Engineering International*, 17, 119-124.

12. Girard, P. and Parent, E. (2001). Bayesian analysis of autocorrelated ordered categorical data for industrial quality monitoring. *Technometrics*, 43(2), 180-191.

13. Girshick, M.A. and Rubin, H. (1952). A Bayes approach to a quality control model. *The Annals of Mathematical Statistics*, 23(1), 114-125.

14. Hamada, M. (2002). Bayesian tolerance interval control limits for attributes. *Quality and Reliability Engineering International*, 18, 45-52.

15. Johnson, N.L., Kotz, S. and Balakrishnan, N. (1994). *Continuous Univariate Distributions*, Volume 1, Second Edition, John Wiley and Sons, New York.

16. Menzefricke, U. (2002). On the evaluation of control chart limits based on predictive distributions. *Communications Statistics - Theory and Methods*, 31(8), 1423-1440.

17. Montgomery, D.C. (2001). *Introduction to Statistical Quality Control*. John Wiley and Sons, New York.

18. Perez, J.M. and Palacin, A.F. (1987). Estimating the quantile function by Bernstein polynomials. *Computational Statistics & Data Analysis*, 5, 391-397.

19. Quesenberry, C.P. (1993). The effect of sample size on estimated limits for $\bar{X}$ and X control charts. *Journal of Quality Technology*, 25, 237-247.

20. Quesenberry, C.P. (1997). *SPC Methods for Quality Improvement*. John Wiley and Sons, New York.

21. Rodriguez, R.N. (1999). Recent issues in statistical process control SAS solutions using statistical modeling procedures. *Unpublished manuscript available form authors*, SAS Institute.

22. Roes, C.B. (1995). *Shewhart-Type Charts in Statistical Process Control*, Ph.D. Thesis, University of Amsterdam, The Netherlands.

23. Roes, K.C.B., Does, R.J.M.M. and Schurink, Y. (1993). Shewhart-type control charts for individual observations. *Journal of Quality Technology*, 25, 188-198.

24. Schilling, E.G. and Nelson, P.R. (1976). The effect of non-normality on the control limits of $\bar{X}$ charts. *Journal of Quality Technology*, 8(4), 183-188.

25. Shewhart, W.A. (1931). *Economic Control of Quality of Manufactured Product*, Van Nostrand, Princeton, New York.

26. Shore, H. (2001). Process control for non-normal populations based on an inverse normalizing transformation. In *Frontiers in Statistical Quality Control 6*, Lenz, H.-J. and Wilrich, P.-Th. (eds.), Springer Verlag, New York, 195-206.

27. Stoumbos, Z.G. and Reynolds, M.R. Jr. (2000). Robustness to non-normality and autocorrelation of individuals control charts. *Computational Statistics & Data Analysis*, 66, 145-187.

28. Sturm, G.W., Feltz, C.J., and Yousry, M.A. (1991). An empirical Bayes strategy for analysing manufacuring data in real time. *Quality and Reliability Engineering International*, 7, 159-167.

29. Tagaras, G. and Nikolaidis, Y. (2002). Comparing the effectiveness of various Bayesian $\bar{X}$ control charts. *Operations Research*, 50(5), 878-888.

30. Vermaat, M.B., Ion, R.A., Does, R.J.M.M. and Klaassen, C.A.J. (2003). A comparison of Shewhart individuals control charts based on normal, non-parametric, and extreme-value theory. *Quality and Reliability Engineering International*, 19, 337-353.

31. Wheeler, D.J. (1991) Shewhart's charts: myths, facts, and competitors. *45th annual quality congress transactions ASQC*, 533-538.

32. Wheeler, D.J. (1995). *Advanced Topics in Statistical Process Control*, SPC Press, Knoxville, Tennessee.

33. Woodall, W.H. and Montgomery, D.C. (1999). Research issues and ideas in statistical process control. *Journal of Quality Technology*, 31(4), 376-386.

34. Woodward, P.W. and Naylor, J.C. (1993). An application of Bayesian methods in SPC. *The Statistician*, Vol. 42, 461-469.

35. Yourstone, S.A. and Zimmer, W.J. (1992). Non-normality and the design of control charts for averages. *Decision Sciences*, 23, 1099-1113.

Appendix A

The Bernstein Approximation

Suppose we have a sample X_1, …, X_k. Denote the sample order statistics by $X_{(1)} < X_{(2)} < \cdots < X_{(k)}$. Suppose ψ is an a priori density function for f and Ψ the corresponding cumulative a priori distribution function for the cumulative distribution F. Assume that f is continuous and strictly positive on the interval (a,b). Transform the order statistics $X_{(1)}, \cdots, X_{(k)}$ by the transformation $Y = \Psi(X)$, such that

$$Y_{(0)} = 0, \quad Y_{(i)} = \Psi(X_{(i)}), \quad i = 1,\dots,k, \quad Y_{(k+1)} = 1.$$

Because $P(Y \le y) = P(\Psi(X) \le y) = P(X \le \Psi^{-1}(y)) = F(\Psi^{-1}(y))$ we have that $B = (F(\Psi^{-1}))^{-1}$ is the quantile function of the random variable $Y = \Psi(X)$. This quantile function is estimated by the so called Bernstein polynomial approximation of degree $(k+1)$ defined by

$$B_k(p) = \sum_{i=0}^{k+1} Y_{(i)} \binom{k+1}{i} p^i (1-p)^{k+1-i}$$

and $0 < p < 1$. Let $H_k = \Psi^{-1}(B_k)$ then $\hat{F}_{k,BA} = H_k^{-1} = B_k^{-1}(\Psi)$ is used to estimate F.

An improvement of the Bernstein approximation can be established by introducing a smoothing parameter m, cf. [2]. First, split the original sample into $k!/[m!(k-m)!]$ subsamples of size m each. Then, the approximation of B of the quantile function of Y is defined by

$$B_k^{(m)}(p) = \binom{k}{m}^{-1} \sum_{1 \le \alpha_1 \le \cdots \le \alpha_m \le k} B_m(p \mid Y_{\alpha_1}, \dots, Y_{\alpha_m}). \tag{9}$$

This can be rewritten as an *L*-statistic

$$B_k^{(m)}(p) = p^{m+1} + \sum_{j=1}^{m} \binom{m+1}{j} p^j (1-p)^{m+1-j} \sum_{i=j}^{k-m+j} \frac{\binom{i-1}{j-1}\binom{k-i}{m-j}}{\binom{k}{m}} Y_{(i)}. \tag{10}$$

From this we obtain $\hat{F}_{k,BA}^{(m)} = B_k^{(m)\,-1}(\Psi)$, as an estimator of F. Differentiating $\hat{F}_{k,BA}^{(m)}$ gives $\hat{f}_{k,BA}^{(m)}$ which is an estimator for f.

For the choice of the smoothing parameter m, we use the rule of thumb determined in [2]

$$m = 2.6k^{1/2}vw \tag{11}$$

and round this off to the nearest integer, where $v = \|\psi - f\|_1$ with $\|\cdot\|_1$ the L_1-distance, $w = 1$ if the number of sign changes of $\psi - f$ is equal to 1 and $w = 2$ if the number of sign changes is greater than 1. A sign change occurs if $\psi - f > 0$ changes in $\psi - f < 0$ or vice versa on the support of $\psi - f$. Note that $0 \le v \le 2$.

For practical purposes we suggest to take $w = 2$ (because most of the time we do not know the number of sign changes of $\psi - f$) and $v = 1$ (which puts a lot of weight on the data). This means that we choose $m = 5.2\sqrt{k}$ and round this off to the nearest integer (cf. formula (11)).

In [2] an example is given how the Bernstein approximation works in practice.

Authors' Biographies:

Thijs Vermaat received a Master's Degree in Econometrics and Operations Research at the University of Groningen in 2002 and a Master's Degree in Statistics at the same university in 2003. Currently he is a PhD student at the University of Amsterdam and a consultant in industrial statistics at the Institute for Business and Industrial Statistics. His research interests are control charts for nonstandard situations. He is a member of ENBIS.

Ronald J.M.M. Does received a MSc degree (cum laude) in Mathematics at the University of Leiden in 1976 and a PhD degree in Statistics at the same university in 1982. From 1981-1989, he worked at the University of Maastricht, where he became Head of the Department of Medical Informatics and Statistics. In 1989 he joined Philips Electronics as a senior consultant in Industrial Statistics. Since 1991 he is Professor of Industrial Statistics at the University of Amsterdam. In 1994 he founded the Institute for Business and Industrial Statistics, which operates as an independent consultancy firm within the University of Amsterdam. His current research activities are control charts for nonstandard situations and Lean Six Sigma. He is a member of the editorial boards of Quality Engineering and Quality Technology and Quantitative Management. He is director of ENBIS, vice president publications of ISBIS, Senior Member of ASQ, and elected member of the ISI.