Department of Applied Mathematics
Faculty of EEMCS

University of Twente
The Netherlands

P.O. Box 217
7500 AE Enschede
The Netherlands
Phone: +31-53-4893400
Fax: +31-53-4893114
Email: memo@math.utwente.nl
www.math.utwente.nl/publications

Memorandum No. 1713

Self adapting control charts

W. Albers and W.C.M. Kallenberg

February, 2004

ISSN 0169-2690

Self Adapting Control Charts

Willem Albers and Wilbert C.M. Kallenberg
Department of Applied Mathematics
Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente
P.O. Box 217, 7500 AE Enschede
The Netherlands

Abstract When the distributional form of the observations differs from normality, standard control charts are often seriously in error. Such model errors can be avoided with (modified) nonparametric control charts. Unfortunately, these control charts suffer from large stochastic errors due to estimation. In between are so called parametric control charts. All three of them are discussed in this paper as well as a combined chart, which chooses one of the three control charts according to the appropriate model assumption on the underlying distribution. The data themselves tell us which of the three control charts to select. Ready-made formulas for the several control charts are presented accompanied by an application on real data. Apart from bias removal, criteria based on exceedance probability and semi-variance are investigated.

Keyword and phrases: statistical process control, Phase II control limits, unbiasedness, exceedance probability, semi-variance, normal power family, model error, nonparametric, model selection.

2000 Mathematics Subject Classification: 62 G 32, 62 P 30, 65 C 05.

[0]This research was supported by the Technology Foundation STW, applied science division of NWO and the technology programme of the Ministry of Economic Affairs.

1 Introduction

Consider the traditional (individual) Shewhart $\overline{X}$-chart for monitoring the mean of a production process. An upper limit UL and a lower limit LL are calculated and an out-of-control signal is produced as soon as an observation falls outside these limits. Normal control charts, corrected for the effects of applying estimators of the parameters, are investigated in Albers and Kallenberg (2003, 2004a, b). When normality does not hold, normal control charts are often seriously in error, see e.g. Chan et al. (1988), Pappanastos and Adams (1996), Albers et al. (2002, 2004). To solve this problem parametric control charts are developed in Albers et al. (2002, 2004). The idea is that if we have heavier tails than in the case of a normal distribution, UL should be taken larger and LL smaller to get the same false alarm rate and vice versa for lighter tailed distributions. The extra shape parameter in the parametric model does the job. If the parametric model also fails, a nonparametric approach can be considered, see Albers et al. (2004c). The obvious question is when to choose which chart. The data can be used to make the decisions, see Albers et al. (2003). In this way a chart is chosen which adapts itself to the situation at hand. A non-technical review of normal control charts is presented in Albers and Kallenberg (2003). Here we proceed by reviewing the parametric and nonparametric chart and the final step of the data driven choice between the three types of charts.

For all three types of charts, normal, parametric and nonparametric, two forms of corrected control charts are designed, associated with two performance criteria. The corrections are needed because quantiles of the distributions should be estimated and the estimation process causes (so called stochastic) errors. The first form concerns bias removal for the control chart (as a whole). Bias deals with the long run behavior of charts. Even with a small bias a single chart may be in error, since there still may be large variability due to the estimation step. The second form is devoted to correct the influence of the estimators in a more stringent way aiming at sensible behavior not in the long run, but for a single chart. We call this the exceedance probability approach. A more detailed description is given in Section 2.

In contrast to the papers mentioned above, here we concentrate on the two-sided versions of the charts as they are more important in practice. We present explicit formulas for the corresponding control limits. The one-sided versions are easily obtained from the two-sided ones. Moreover, we introduce some other criteria like the so called semi-variance. It turns out that both the exceedance criterion and the semi-variance can be regarded as special cases of a more general set of criteria. This generalization in itself is an interesting new aspect in this paper, but more importantly for practice, it shows that protection in terms of exceedance probabilities gives at the same time protection against a lot of other criteria as well with among them the semi-variance.

In Section 2 the general approach is presented, including precise definitions of the bias and exceedance probability criterion. The three control charts are discussed in Section 3. Moreover, ready-made formulas of the corrected control limits are given and they are exemplified by application to real data. Here also a modified version of the nonparametric control chart is presented. Section 4 is devoted to the choice of the model. The three control charts and the decision rule telling us which one to choose among them, in fact together form a new control chart, the combined control chart. Again, the required formulas are written out and the calculations are shown by application to the real data example. In Section 5 the exceedance criterion is put into a more general framework, containing the semi-variance as another special case. The paper is closed with a discussion on how to handle with the available methods when using this type of control charts in practice.

2 General approach

Let $X_1, \dots, X_n$ be a sample of so-called Phase I observations. It is assumed that these observations are in-control and we use them for estimating unknown quantities in the basic control limits and in addition for making necessary corrections after plugging in the estimators. For instance, the well-known '3σ'-limits in the traditional normal control chart, corresponding to $UL = \mu + 3\sigma$ and $LL = \mu - 3\sigma$, are firstly replaced by $\widehat{\mu} + 3\widehat{\sigma}$ and $\widehat{\mu} - 3\widehat{\sigma}$, respectively, where $\widehat{\mu} = \overline{X}, \widehat{\sigma} = S = \sqrt{S^2}$ with $S^2 = (n-1)^{-1} \sum \left(X_i - \overline{X}\right)^2$. When dealing with the false alarm rate (FAR) itself in the bias case they are subsequently corrected to $\widehat{UL} = \widehat{\mu} + 3\widehat{\sigma}\left(1 + 3n^{-1}\right)$ and $\widehat{LL} = \widehat{\mu} - 3\widehat{\sigma}\left(1 + 3n^{-1}\right)$, see (5.3) in Albers and Kallenberg (2003).

The prescribed false alarm rate will be denoted by p. For example, $UL = \mu + 3\sigma$ and $LL = \mu - 3\sigma$ (with known parameters μ and σ!) correspond to $p = 0.0027$. Due to estimation the probability of a false alarm is no longer a number, but a random variable. Let X_{n+1} be a new observation from Phase II. The upper and lower limit depend on the observations $X_1, \dots, X_n$. They will be denoted by $\widehat{UL}$ and $\widehat{LL}$, respectively. Let P_n be the observed alarm rate, i.e. the probability of an alarm, given $X_1, \dots, X_n$. In formula

$$P_n = \Pr\left(X_{n+1} > \widehat{UL} \text{ or } X_{n+1} < \widehat{LL}\,|X_1, \dots, X_n\right). \tag{1}$$

In the in-control situation, where X_{n+1} has the same distribution as $X_1, \dots, X_n$, we want P_n to be close to the prescribed false alarm rate p. (If not stated otherwise we consider in this section the in-control situation.) As P_n is random, there is no unique way to express the discrepancy. Two popular approaches are as follows. The first is the mildest one: look at the bias $EP_n - p$; if it is sufficiently small, the chart is considered to be O.K. To the same category belongs the more general comparison of $Eg(P_n)$ with $g(p)$ with for instance the function $g(p) = 1/p$ corresponding to the average run length (ARL). When considering the bias we will not take the ARL into account, e.g. because, contrary to simple intuition, $1/P_n$ has a positive bias (due to the occurrence of extremely long runs, which although having small probability nevertheless strongly determine the expectation). As a consequence, the correction for estimating brings the control limits closer to each other, while one has the feeling that one should "pay" something for estimating the parameters and not "gain" something. Several authors, for example Roes (1995, page 34), therefore remark that $E\left(1/P_n\right)$ does not adequately summarize the run length properties of the chart, cf. also Quesenberry (1993, page 242) and Jones et al. (2004, page 100).

The bias criterion is satisfactory if we are interested in the behavior of a long series of applications of charts. However, if the bias of P_n is indeed small, but its variability is still large, very long runs are mixed up with unwelcome very short ones. Therefore, for an application to a single chart the second criterion is more appropriate. It concerns the probability that P_n exceeds p by more than a given percentage (for example 10%). If this exceedance probability is sufficiently small (e.g. 10%), the chart is O.K. in this more strict way. So, with high probability P_n is with at most 10%, say, in error. Since the starting point of a two-sided control chart with prescribed false alarm rate p is a combination of two one-sided control charts each with prescribed false alarm rate $p/2$, we will treat the upper and lower limit here also separately. Therefore, we present corrected control limits such that the probability that P_{nL} exceeds $p/2$ by more than a given percentage only with some given small probability, where

$$P_{nL} = \Pr\left(X_{n+1} < \widehat{LL}\,|X_1, \dots, X_n\right)$$

and similarly for P_{nU} with

$$P_{nU} = \Pr\left(X_{n+1} > \widehat{UL}\,|X_1, \dots, X_n\right).$$

To be more precise, we require that the intended FAR $p/2$ should not be exceeded by the outcome P_{nL} and P_{nU} by more than a fraction ε in more than $100\alpha\%$ of the applications. In formula,

$$\Pr\left(P_{nL} > \frac{p}{2}(1+\varepsilon)\right) \le \alpha \text{ and } \Pr\left(P_{nU} > \frac{p}{2}(1+\varepsilon)\right) \le \alpha. \tag{2}$$

We consider also the more general situation of $g(P_{nL})$ exceeding $g(p/2)$ and similarly for $g(P_{nU})$. The problems with $g(p) = 1/p$ in the bias case do not occur here, since we deal with probabilities and hence extremely long runs are not dominating. We require that unpleasant values of $1/P_{nL}$ and $1/P_{nU}$ should be sufficiently rare, leading to

$$\Pr\left(\frac{1}{P_{nL}} < \frac{1}{p/2}(1-\varepsilon)\right) \le \alpha \text{ and } \Pr\left(\frac{1}{P_{nU}} < \frac{1}{p/2}(1-\varepsilon)\right) \le \alpha. \tag{3}$$

It is easily seen that (3) is the same as (2) with ε replaced by $\varepsilon/(1-\varepsilon)$.

The exceedance criterion given in (3) can also be interpreted as follows. The conditional ARL's due to the lower limit, $1/P_{nL}$, and due to the upper limit, $1/P_{nU}$, depend on the Phase I distributions $X_1, \ldots, X_n$. Consider these ARL's as random variables. The $100\alpha^{\text{th}}$ percentile of those random variables are at least $(1-\varepsilon)/(p/2)$. For instance with $p/2 = 0.001, \varepsilon = 0.1, \alpha = 0.1$ this gives that the 10^{th} percentile of the ARL's $1/P_{nL}$ and $1/P_{nU}$ are at least 900. It may be better to settle this kind of percentiles of the ARL than to base oneself on the bias (that is the expectation) of the ARL, see also Jones et al. (2004, page 102).

3 Three types of control charts

If the distribution function of the observations were known to be F, we could simply find the lower control limit LL by solving $F(LL) = p/2$ and the upper limit UL by solving $F(UL) = 1 - p/2$ For instance, when normality holds with known μ and σ, this leads to $LL = \mu - u_{p/2}\sigma$, $UL = \mu + u_{p/2}\sigma$, where $u_{p/2}$ is the point that is exceeded by a standard normal variable with probability $p/2$, in formula $1 - \Phi(u_{p/2}) = p/2$ with Φ the standard normal distribution function. When the parameters μ and σ are unknown, they should be estimated and this causes an error, which we call the stochastic error. Corrections can be made to repair this stochastic error according to the bias or exceedance criterion. When the distributional form of the observations differs from normality huge errors may occur. We call this the model error. Basically, the problem is that normality may be reasonable for the central part of the distribution, but we deal with the far tail of the distribution, due to the typically very small values of p, and large relative errors are produced.

To avoid the model error, a nonparametric approach can be proposed. However, estimating the 0.999-quantile with 100 observations is not very successful in a nonparametric setting. A parametric family based on the family of normal distributions with an extra shape parameter serves as an intermediate step between the classical normal chart with its possibly large model error and the nonparametric chart with its large stochastic error. The idea is to enlarge the $p/2$-quantile for heavy tailed distributions and to reduce it for light tailed distributions. Essentially, we replace $u_{p/2}$ by $u_{p/2}^{1+\gamma}$ which has the required effect: when $\gamma > 0$, we get larger values (heavy tail!), for $\gamma < 0$ we get smaller values (light tail!), while, when the normal fit is still good, $\gamma = 0$ returns the normal quantile.

Having in general not enough data to estimate the very far tail, we estimate in the parametric family the ordinary tail (where data are available) to make inference on the very far tail. This is done separately at the lower tail and the upper tail. Consequently, not only the thickness of the tail is taken into account, but also the skewness: the value of γ at the lower tail may be different from that at the upper tail.

The three types of control charts with the appropriate corrections are presented in the next subsections. They are exemplified by application on a real life example concerning the production of electric shavers by Philips. In an electrochemical process razor heads are formed. The measurements concern the thickness of the razor heads. The sample consists of 835 measurements. A histogram of the sample is given in Figure 1.

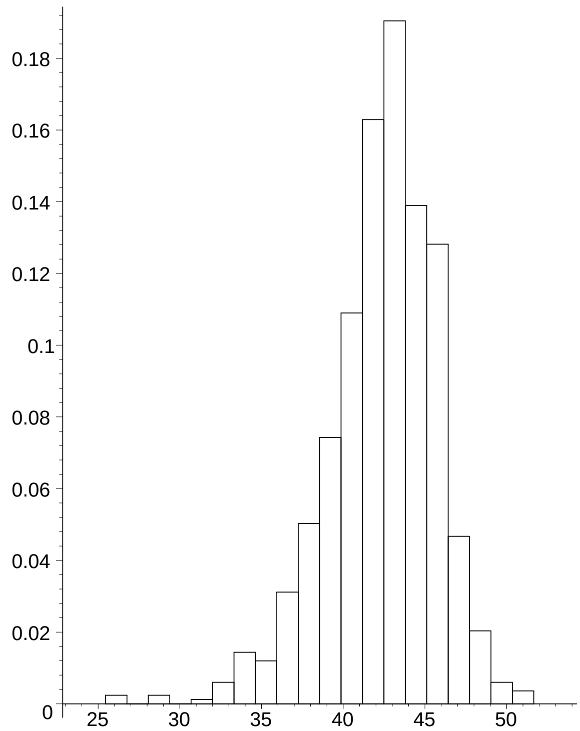

Figure 1. Histogram of the thickness of razor heads for a sample of 835 measurements.

3.1 Normal control charts

An overview of the normal control chart is given in Albers and Kallenberg (2003). Therefore, here we simply present the results. The two-sided control limits for the (corrected) normal control chart are given in Table 1. They are denoted by $\widehat{LL}_N$ and $\widehat{UL}_N$, where the subscript N refers to "normality". In case of the lower limit $\widehat{LL}_N$ the $-$ sign should be read and for $\widehat{UL}_N$ the $+$ sign is used. The one-sided versions are simply obtained by replacing $p/2$ in the control limits by p.

Table 1 *Two-sided (corrected) normal control limits.*

Aim	$\widehat{LL}_N, \widehat{UL}_N$
$EP_n = p$	$\overline{X} \pm u_{p/2} S \left\{ 1 + \frac{u_{p/2}^2 + 3}{4n} \right\}$
$\Pr\left(P_{nL} > \frac{p}{2}(1+\varepsilon)\right) \leq \alpha$, $\Pr\left(P_{nU} > \frac{p}{2}(1+\varepsilon)\right) \leq \alpha$	$\overline{X} \pm u_{p/2} S \left\{ 1 + \frac{u_\alpha \left(\frac{1}{2} + u_{p/2}^{-2}\right)^{1/2}}{\sqrt{n}} - \frac{\varepsilon}{u_{p/2}^2} \right\}$
$\Pr\left(\frac{1}{P_{nL}} < \frac{1}{p/2}(1-\varepsilon)\right) \leq \alpha$, $\Pr\left(\frac{1}{P_{nU}} < \frac{1}{p/2}(1-\varepsilon)\right) \leq \alpha$	$\overline{X} \pm u_{p/2} S \left\{ 1 + \frac{u_\alpha \left(\frac{1}{2} + u_{p/2}^{-2}\right)^{1/2}}{\sqrt{n}} - \frac{\varepsilon}{u_{p/2}^2 (1-\varepsilon)} \right\}$

Application to the example mentioned before gives for $p/2 = 0.001, \varepsilon = 0.1, \alpha = 0.1$ the following two-sided control limits. The sample mean $\overline{X}$ equals 42.366 and the sample standard deviation $S = 3.311$, while $u_{p/2} = 3.090, u_\alpha = 1.282$ and $n = 835$. We get

$$\begin{aligned} \text{bias} \quad & \widehat{LL}_N = 32.096, \widehat{UL}_N = 52.635 \\ \text{exceedance } FAR \quad & \widehat{LL}_N = 31.889, \widehat{UL}_N = 52.842 \\ \text{exceedance } ARL \quad & \widehat{LL}_N = 31.901, \widehat{UL}_N = 52.830. \end{aligned} \tag{4}$$

For the derivation of the normal control charts and their properties we refer to Albers and Kallenberg (2003, 2004a,b).

3.2 Parametric control charts

To embed the normal distributions in a larger family with heavier or lighter tails we consider essentially powers of the standard normal quantiles as the new quantiles. The family of distributions obtained in this way is called the normal power family. More precisely, replace u_p (for $0 < p < 1/2$) by

$$c(\gamma)u_p^{1+\gamma}, \tag{5}$$

where $\gamma > -1$ and where $c(\gamma)$ is a normalizing constant (to make the variance equal to one) given by

$$c(\gamma) = \pi^{1/4}2^{-(1+\gamma)/2}\Gamma(\gamma + \tfrac{3}{2})^{-1/2}$$

with Γ the Gamma-function. It is immediately seen that $\gamma = 0$ leads to the standard normal quantile u_p.

For estimation of the extra shape parameter γ we use sample quantiles in the ordinary tail. Although the normal power family for a given γ is symmetric, we deliberately avoid the symmetry by estimating not the whole distribution at once, but the lower and upper tail separately. On the one hand, this has been done with the idea that the ordinary tail tells us more about the far tail than the middle part of the distribution; on the other hand, we want to deal also with skew distributions as it occurs very often that the behavior in the lower tail differs from that in the upper tail, see also Figure 1. In this way the extension w.r.t. normality goes into two directions: skewness is allowed, as well as different tail behavior.

More precisely, let $X_{(1)} \leq \ldots \leq X_{(n)}$ be the order statistics of $X_1, \ldots, X_n$. The estimator of γ at the upper tail, $\widehat{\gamma}_U$, is given by ($ent(x)$ is the integer part of x, for instance $ent(0.95 \times 835 + 1) = 794$)

$$\widehat{\gamma}_U = 1.1218 \log \left(\frac{X_{(ent(0.95n+1))} - \overline{X}}{X_{(ent(0.75n+1))} - \overline{X}} \right) - 1.$$

Note that $\left(X_{(ent(0.95n+1))} - \overline{X}\right) / \left(X_{(ent(0.75n+1))} - \overline{X}\right)$ estimates $\left(c(\gamma)u_{0.05}^{1+\gamma}\right) / \left(c(\gamma)u_{0.25}^{1+\gamma}\right) = \left(u_{0.05}/u_{0.25}\right)^{1+\gamma}$ and that $1/\log\left(u_{0.05}/u_{0.25}\right) = 1.1218$. Similarly, the estimator of γ at the lower tail, $\widehat{\gamma}_L$, is given by

$$\widehat{\gamma}_L = 1.1218 \log \left(\frac{\overline{X} - X_{(n-ent(0.95n))}}{\overline{X} - X_{(n-ent(0.75n))}} \right) - 1.$$

As an example of a skew distribution we consider the Normal Inverse Gaussian (2, 1.5, 0, 1) distribution, cf. Barndorff-Nielsen (1996). Figure 2 clearly shows that the member of the normal power family chosen by the limiting value γ_U of our estimator $\widehat{\gamma}_U$ fits the distribution at the right tail and not globally, just as wanted. Similarly, the left-hand side is fitted by another member of the normal power family (corresponding to $\widehat{\gamma}_L$) leading to a serious reduction of the model error on the left-hand side.

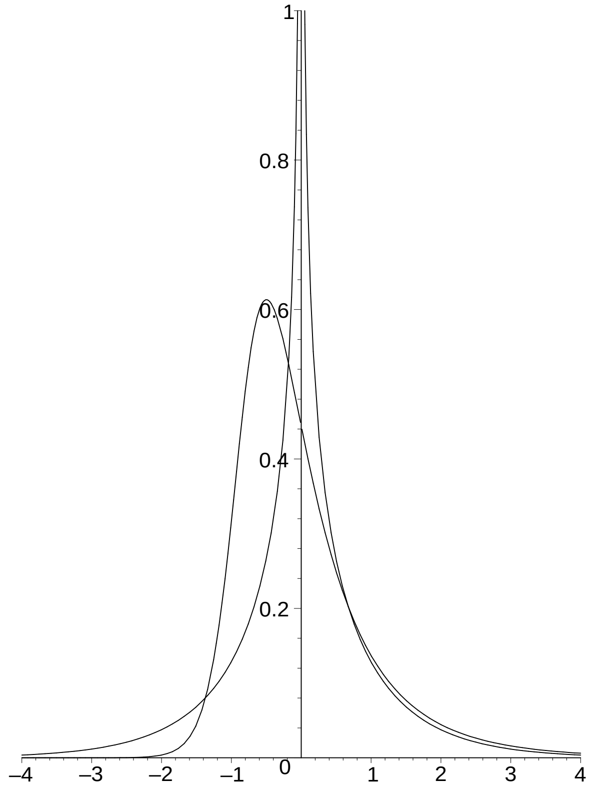

Figure 2. Density of the Normal Inverse Gaussian ($2, 1.5, 0, 1$*) and the corresponding density of the normal power family with* $\gamma_U = 0.77$.

Just adding an extra shape parameter actually makes it far more complicated to derive control charts. Fortunately, the control charts in this family can be applied quite straightforwardly. Before presenting the control limits we introduce some notation

$$\begin{aligned}
C1(\gamma, u_p) &= -1.23 - 0.63\gamma + 0.73\gamma^2 + 0.74u_p - 0.08\gamma u_p - 0.14\gamma^2 u_p, \\
C2(\gamma) &= \left(\frac{u_{a_n}}{u_{b_n}}\right)^{1+\gamma} - 2.4387^{1+\gamma} \\
\text{with } a_n &= 1 - \frac{ent\,(0.95n+1)}{n+1}, b_n = 1 - \frac{ent\,(0.75n+1)}{n+1}, \\
C3(\gamma, u_p) &= -76.37 - 120.12\gamma - 81.93\gamma^2 + 35.53u_p + 53.71\gamma u_p + 37.18\gamma^2 u_p, \\
A(\gamma, u_p) &= -4.00 - 12.54\gamma - 10.02\gamma^2 + 2.91u_p + 6.47\gamma u_p + 4.42\gamma^2 u_p.
\end{aligned}$$

The two-sided control limits for the (corrected) parametric control chart are given in Table 2. They are denoted by $\widehat{LL}_P$ and $\widehat{UL}_P$, where the subscript P refers to "parametric". In case of the lower limit $\widehat{LL}_P$ the $-$ sign should be read and $\widehat{\gamma} = \widehat{\gamma}_L$ should be inserted, while for $\widehat{UL}_P$ the $+$ sign is used and $\widehat{\gamma} = \widehat{\gamma}_U$. The one-sided versions are simply obtained by replacing $p/2$ in the control limits by p.

Table 2 *Two-sided (corrected) parametric control limits based on the normal power family.*

Aim	$\widehat{LL}_P, \widehat{UL}_P$
$EP_n = p$	$\overline{X} \pm S\left\{c(\widehat{\gamma})u_{p/2}^{1+\widehat{\gamma}} - C1(\widehat{\gamma}, u_{p/2})C2(\widehat{\gamma}) + \frac{C3(\widehat{\gamma}, u_{p/2})}{n}\right\}$
$\Pr\left(P_{nL} > \frac{p}{2}(1+\varepsilon)\right) \leq \alpha$, $\Pr\left(P_{nU} > \frac{p}{2}(1+\varepsilon)\right) \leq \alpha$	$\overline{X} \pm S\left\{c(\widehat{\gamma})u_{p(1+\varepsilon)/2}^{1+\widehat{\gamma}} + \frac{A(\widehat{\gamma}, u_{p/2})u_\alpha}{\sqrt{n}}\right\}$
$\Pr\left(\frac{1}{P_{nL}} < \frac{1}{p/2}(1-\varepsilon)\right) \leq \alpha$, $\Pr\left(\frac{1}{P_{nU}} < \frac{1}{p/2}(1-\varepsilon)\right) \leq \alpha$	$\overline{X} \pm S\left\{c(\widehat{\gamma})u_{p/\{2(1-\varepsilon)\}}^{1+\widehat{\gamma}} + \frac{A(\widehat{\gamma}, u_{p/2})u_\alpha}{\sqrt{n}}\right\}$

Application to the example mentioned before gives for $p/2 = 0.001, \varepsilon = 0.1, \alpha = 0.1$ the following two-sided control limits. We have $\overline{X} = 42.366, S = 3.311, n = 835, X_{(n-ent(0.95n))} = 36.54, X_{(n-ent(0.75n))} = 40.62, \widehat{\gamma}_L = 0.352, X_{(ent(0.95n+1))} = 47.03, X_{(ent(0.75n+1))} = 44.54$, $\widehat{\gamma}_U = -0.144$ while $u_{p/2} = 3.090, u_{p(1+\varepsilon)/2} = 3.062, u_{p/\{2(1-\varepsilon)\}} = 3.059$ and $u_\alpha = 1.282$. We get

$$\begin{aligned} &\text{bias} && \widehat{LL}_P = 29.100, \widehat{UL}_P = 51.606 \\ &\text{exceedance } FAR && \widehat{LL}_P = 28.306, \widehat{UL}_P = 52.001 \\ &\text{exceedance } ARL && \widehat{LL}_P = 28.324, \widehat{UL}_P = 51.994. \end{aligned} \quad (6)$$

As is seen in the histogram (Figure 1), we have a heavy lower tail and indeed $\widehat{\gamma}_L > 0$. As a consequence $\widehat{LL}_P$ is much lower than $\widehat{LL}_N$. In the upper tail the situation is reversed. We have a thin upper tail (see Figure 1) giving $\widehat{\gamma}_U < 0$ and hence the upper control limits are somewhat smaller in the parametric model than those in the normal control charts.

For the derivation of these control chart and more details about their properties we refer to Albers et al. (2004,2002).

3.3 Nonparametric control charts

Here we do not assume any model for the distribution function of our observations (except continuity). Therefore, the control limits are based on the highest and lowest order statistics. For instance, the unconditional probability $P\left(X_{n+1} > X_{(n)}\right) = 1/(n+1)$ and hence taking $X_{(n)}$ as upper control limit in a one-sided control chart gives $P_n = 1 - F\left(X_{(n)}\right)$ and $EP_n = 1/(n+1)$. This makes clear that an expected FAR equal to 0.001 can only be achieved with at least 999 observations. For smaller n, even the largest observation is too small to serve as upper limit. A possible way out is to take a randomized procedure with $X_{(n)}$ as upper control limit, but only with such probability that $EP_n = p$. For example, when $p = 0.001$ and $n = 300$, we throw a coin with probability of head equal to 0.301. When head turns up the upper control limit is $X_{(n)}$ and a signal is produced when this upper control limit is exceeded. When tail turns up, never an out-of-control signal is given. Although this delivers the right expected FAR, such a procedure is hardly acceptable, because with probability 0.699 we will never get an out-of-control signal!

In fact, the preceding procedure takes $UL = X_{(n)}$ with probability 0.301 and $UL = \infty$ with probability 0.699. We modify this procedure somewhat by taking a "very large value" instead of ∞. A convenient choice for this very large value is $X_{(n)} + S$. Formally, this is no longer a nonparametric chart, but from a practical point of view it turns out to work very well.

When $1/(n+1) < p < 2/(n+1)$, we should randomize between $X_{(n-1)}$ and $X_{(n)}$. To make things a little bit more simple (without much loss of accuracy) we take the deterministic mixture $\delta X_{(n-1)} + (1-\delta)X_{(n)}$ with $\delta = p(n+1) - 1$. For larger p we proceed in a similar way.

Determination of the control limits using the exceedance probability as criterion is done in a similar way. For details and further discussion of the nonparametric control charts we refer to Albers and Kallenberg (2004c) and Albers et al.(2003).

The two-sided control limits for the nonparametric control chart are given in Table 3. In this table we use the following notation

$$\begin{aligned}
&r = ent[(n+1)p/2], \delta = (n+1)p/2 - r,\\
&k(\varepsilon) : \text{smallest integer such that } Po\,(n(1+\varepsilon)p/2, r-1-k) \le \alpha,\\
&\text{where } Po\,(\mu, x) = \Pr(Y \le x) \text{ with } Y \sim Poisson(\mu),\\
&\lambda(\varepsilon) = \frac{\alpha - Po\,(n(1+\varepsilon)p/2, r-1-k(\varepsilon))}{Po\,(n(1+\varepsilon)p/2, r-k(\varepsilon)) - Po\,(n(1+\varepsilon)p/2, r-1-k(\varepsilon))},\\
&X_{(0)} = X_{(1)} - S, X_{(n+1)} = X_{(n)} + S.
\end{aligned} \tag{7}$$

The control limits are denoted by $\widehat{LL}_{NP}$ and $\widehat{UL}_{NP}$, where the subscript NP refers to "non-parametric". The one-sided versions are simply obtained by replacing $p/2$ in the control limits by p.

Table 3 *Two-sided nonparametric control limits; notation see (7)*

Aim	$\widehat{LL}_{NP}, \widehat{UL}_{NP}$
$EP_n = p$	$\widehat{LL}_{NP}: \; r=0: \begin{cases} X_{(1)} - S & \text{with probability } 1-\delta \\ X_{(1)} & \text{with probability } \delta \end{cases}$ $r \ge 1: \; (1-\delta)X_{(r)} + \delta X_{(r+1)}$ $\widehat{UL}_{NP}: \; r=0: \begin{cases} X_{(n)} & \text{with probability } \delta \\ X_{(n)} + S & \text{with probability } 1-\delta \end{cases}$ $r \ge 1: \; \delta X_{(n-r)} + (1-\delta)X_{(n-r+1)}$
$\Pr\left(P_{nL} > \frac{p}{2}(1+\varepsilon)\right) \le \alpha$, $\Pr\left(P_{nU} > \frac{p}{2}(1+\varepsilon)\right) \le \alpha$	$\widehat{LL}_{NP}: \begin{cases} X_{(r-k(\varepsilon))} & \text{with probability } 1-\lambda(\varepsilon) \\ X_{(r+1-k(\varepsilon))} & \text{with probability } \lambda(\varepsilon) \end{cases}$ $\widehat{UL}_{NP}: \begin{cases} X_{(n+k(\varepsilon)-r)} & \text{with probability } \lambda(\varepsilon) \\ X_{(n+k(\varepsilon)-r+1)} & \text{with probability } 1-\lambda(\varepsilon) \end{cases}$
$\Pr\left(\frac{1}{P_{nL}} < \frac{1}{p/2}(1-\varepsilon)\right) \le \alpha$, $\Pr\left(\frac{1}{P_{nU}} < \frac{1}{p/2}(1-\varepsilon)\right) \le \alpha$	$\widehat{LL}_{NP}: \begin{cases} X_{(r-k(\varepsilon/(1-\varepsilon)))} & \text{with prob. } 1-\lambda(\varepsilon/(1-\varepsilon)) \\ X_{(r+1-k(\varepsilon/(1-\varepsilon)))} & \text{with prob. } \lambda(\varepsilon/(1-\varepsilon)) \end{cases}$ $\widehat{UL}_{NP}: \begin{cases} X_{(n+k(\varepsilon/(1-\varepsilon))-r)} & \text{with prob. } 1-\lambda(\varepsilon/(1-\varepsilon)) \\ X_{(n+k(\varepsilon/(1-\varepsilon))-r+1)} & \text{with prob. } \lambda(\varepsilon/(1-\varepsilon)) \end{cases}$

Application to the example mentioned before gives for $p/2 = 0.001, \varepsilon = 0.1, \alpha = 0.1$ the following two-sided control limits. We have $r = 0, \delta = 0.836, k(0.1) = 0, k(0.1/(1-0.1)) = 0, \lambda(0.1) = 0.251, \lambda(0.1/(1-0.1)) = 0.253, X_{(1)} = 25.45, X_{(835)} = 51.66$ and $S = 3.311$. We get

$$\begin{array}{lll}
\text{bias} & \widehat{LL}_{NP} = \begin{cases} 22.139 & \text{with prob. } 0.164 \\ 25.45 & \text{with prob. } 0.836 \end{cases} & \widehat{UL}_{NP} = \begin{cases} 51.66 & \text{with prob. } 0.836 \\ 54.971 & \text{with prob. } 0.164 \end{cases} \\
\begin{array}{c}\text{exceedance}\\ FAR\end{array} & \widehat{LL}_{NP} = \begin{cases} 22.139 & \text{with prob. } 0.749 \\ 25.45 & \text{with prob. } 0.251 \end{cases} & \widehat{UL}_{NP} = \begin{cases} 51.66 & \text{with prob. } 0.251 \\ 54.971 & \text{with prob. } 0.749 \end{cases} \\
\begin{array}{c}\text{exceedance}\\ ARL\end{array} & \widehat{LL}_{NP} = \begin{cases} 22.139 & \text{with prob. } 0.747 \\ 25.45 & \text{with prob. } 0.253 \end{cases} & \widehat{UL}_{NP} = \begin{cases} 51.66 & \text{with prob. } 0.253 \\ 54.971 & \text{with prob. } 0.747 \end{cases}.
\end{array} \tag{8}$$

Comparing the numerical results with the previous control limits for the normal and normal power families, it is clearly seen that in particular the acceptance regions here are much wider than those for the normal and normal power families. Moreover, it is seen that in all situations the exceedance criterion is a more restrictive criterion than the bias criterion, leading to smaller lower and larger upper limits. The difference between FAR and ARL in the exceedance case is fortunately very small, implying that more or less protection against an unwanted large FAR gives at the same time protection against an unpleasantly long ARL.

4 Choosing the model

When the observations are close to normality, we want to select the normal control chart. If the departure from normality is too large, we apply the parametric control chart, unless the parametric family also does not fit. In the latter case the nonparametric control chart comes in.

A first idea might be to execute a (standard) goodness of fit test to investigate normality. If normality is not rejected, use the normal control chart. If we do reject, apply a goodness of fit test for the normal power family. Again, when not rejecting, apply the parametric control chart and otherwise use the nonparametric control chart.

Although this way of thinking looks attractive, it has a serious drawback. Standard goodness of fit tests are looking at the majority of the data, and as such concentrate on the middle of the distribution, while here we are not interested in this middle part, but in the (extreme) tail. Therefore, standard goodness of fit tests are not appropriate for the situation at hand. For the same reason, less formal methods like "a good look at the data" or "an inspection of a histogram" are completely insufficient to judge the possible normality in the far tail.

The choice between the three control charts can be seen as a model selection problem. Again, we cannot use standard selection rules, since the common selection rules are intended for the bulk of the data and not for the extreme tail. The motivation to switch from the normal control chart to the parametric control chart or even to the nonparametric control chart is provided by the model error. We keep this in mind when developing our selection rule. Next we will describe how the selection is done for the upper control limits.

Let us start with considering the question whether to use the normal control chart or not. The supposed distribution function is that of $\mu + \sigma Z$ (with Z having a standard normal distribution). Therefore, the model error (w.r.t. the upper control limit) concerns the behavior of $(X - \mu)/\sigma$ in the far upper tail (p is small!) and we want to check that behavior (based on our observations $X_1, \ldots, X_n$). The most obvious quantity to look at is the standardized maximum of our observations: $(X_{(n)} - \overline{X})/S$.

The next point is the determination of the cut-off points for staying at the normal chart. Distributions with heavier tails than the normal one give problems with the in-control behavior, leading for common distributions to EP_n being 4 or even 12 times as large as it should be, see Table 1 in Albers et al. (2004) and hence the control chart is invalid. Distributions with thinner tails are conservative in the in-control case with as consequence a loss in the out-of-control. Because errors in the in-control are more serious than those in the out-of-control case and moreover, since a positive model error as large as p or larger can easily occur, whereas the negative model error is at most $-p$, we take the selection rule *unbalanced*. In particular, we will prefer the normal control chart when

$$u_{(-0.7+0.5\log n)/n} \leq \frac{X_{(n)} - \overline{X}}{S} \leq u_{5/(n\sqrt{n})}.$$

To see what this implies, let $X_1, \ldots, X_n$ be i.i.d. random variables with a standard normal distribution. Then

$$P\left(X_{(n)} < u_{(-0.7+0.5\log n)/n}\right) = \left(1 - \frac{-0.7 + 0.5\log(n)}{n}\right)^n \approx \exp\left(0.7 - 0.5\log n\right) \approx \frac{2}{\sqrt{n}}$$

and

$$P\left(X_{(n)} > u_{5/(n\sqrt{n})}\right) = 1 - \left(1 - \frac{5}{n\sqrt{n}}\right)^n \approx 1 - \exp\left(-\frac{5}{\sqrt{n}}\right) \approx \frac{5}{\sqrt{n}}.$$

Hence, when normality holds we stay with high probability at the normal chart. When the standardized maximum $(X_{(n)} - \overline{X})/S$ is very large, this indicates that the tail may be heavier than the normal one and we change to the parametric or nonparametric chart. Similarly we change for relatively small values of $(X_{(n)} - \overline{X})/S$ pointing to thinner tail behavior. That the latter is less severe, is expressed in the probabilities $2/\sqrt{n}$ and $5/\sqrt{n}$, respectively.

The cut-off points for staying at the parametric chart are chosen in an analogous way and for the lower control limits the same procedure holds. Therefore, the total two-sided combined control chart is given by the following scheme. Let

$$
\begin{aligned}
&d_{1N} = u_{(-0.7+0.5\log n)/n}, d_{2N} = u_{5/\left(n\sqrt{n}\right)},\\
&d_{1P}\left(\widehat{\gamma}\right) = c\left(\widehat{\gamma}\right) u^{1+\widehat{\gamma}}_{(-0.2+0.5\log n)/n}, d_{2P}\left(\widehat{\gamma}\right) = c\left(\widehat{\gamma}\right) u^{1+\widehat{\gamma}}_{3/\left(n\sqrt{n}\right)}
\end{aligned}
$$

and denote the lower and upper limit by $\widehat{LL}_C$ and $\widehat{UL}_C$, respectively. The upper limit is obtained as follows

$$
\begin{array}{ccc}
\boxed{d_{1N} \le \frac{X_{(n)}-\overline{X}}{S} \le d_{2N}} & \overset{\text{yes}}{\rightarrow} & \boxed{\widehat{UL}_C = \widehat{UL}_N} \\
\downarrow \text{no} & & \\
\boxed{d_{1P}\left(\widehat{\gamma}_U\right) \le \frac{X_{(n)}-\overline{X}}{S} \le d_{2P}\left(\widehat{\gamma}_U\right)} & \overset{\text{yes}}{\rightarrow} & \boxed{\widehat{UL}_C = \widehat{UL}_P} \\
\downarrow \text{no} & & \\
\boxed{\widehat{UL}_C = \widehat{UL}_{NP}} & &
\end{array}
$$

Here $\widehat{UL}_N \widehat{UL}_P$ and $\widehat{UL}_{NP}$ are taken from Table $1-3$ according to the required aim: bias, exceedance with FAR or exceedance with ARL.

For the lower limit $\widehat{LL}_C$ we get

$$
\begin{array}{ccc}
\boxed{d_{1N} \le \frac{\overline{X}-X_{(1)}}{S} \le d_{2N}} & \overset{\text{yes}}{\rightarrow} & \boxed{\widehat{LL}_C = \widehat{LL}_N} \\
\downarrow \text{no} & & \\
\boxed{d_{1P}\left(\widehat{\gamma}_L\right) \le \frac{\overline{X}-X_{(1)}}{S} \le d_{2P}\left(\widehat{\gamma}_L\right)} & \overset{\text{yes}}{\rightarrow} & \boxed{\widehat{LL}_C = \widehat{LL}_P} \\
\downarrow \text{no} & & \\
\boxed{\widehat{LL}_C = \widehat{LL}_{NP}} & &
\end{array}
$$

Here $\widehat{LL}_N, \widehat{LL}_P$ and $\widehat{LL}_{NP}$ are taken from Table $1-3$ according to the required aim: bias, exceedance with FAR or exceedance with ARL.

In Albers et al. (2003) the properties of the combined control chart are extensively discussed, both from a theoretical point of view as by simulations to see the finite sample behavior. We shortly present the conclusions.

1. Theoretically it is shown that the combined control chart in each of the three situations (normality, normal power family, outside the normal power family) asymptotically behaves as the specific corresponding control chart.

2. If normality holds, the combined control chart has a substantial gain w.r.t. the non-parametric control chart and only a small loss w.r.t. the normal and parametric control chart.

3. If the true distribution belongs to the normal power family, the combined control chart has for not too large n and γ a pretty gain w.r.t. the nonparametric control chart and in general only a small loss w.r.t. the parametric control chart; the normal control chart cannot be applied in that case unless γ is very close to 0; for positive γ's this is due to its bad in-control behavior and for negative γ's it has a very low alarm rate under out-of-control.

4. Outside the normal power family, the combined control chart exhibits only a small loss w.r.t. the nonparametric control chart, while the normal control chart cannot be applied in this case due to its bad in-control behavior or its low alarm rate (unless the distribution is very close to normality); the same holds for the parametric control chart, albeit to a much smaller extent.

We exemplify the application of the procedure by doing the calculations for the above example. We have $d_{1N} = 2.728, d_{2N} = 3.531, X_{(835)} = 51.66, \overline{X} = 42.366, S = 3.311$ and hence

$$\frac{X_{(n)} - \overline{X}}{S} = 2.807 \in [2.728, 3.531].$$

Therefore, we apply for the upper limit that from the normal chart: $\widehat{UL}_C = \widehat{UL}_N$. Because $X_{(1)} = 25.45$, we get

$$\frac{\overline{X} - X_{(1)}}{S} = 5.109 \notin [2.728, 3.531]$$

and at the lower part we do not use the normal chart. Next we calculate the cut-off points for staying at the parametric chart. We have $\widehat{\gamma}_L = 0.352$ (see above) and hence $c(\widehat{\gamma}_L) = 0.857, u^{1+\widehat{\gamma}_L}_{(-0.2+0.5\log n)/n} = 3.773, d_{1P} = 3.232, u^{1+\widehat{\gamma}_L}_{3/(n\sqrt{n})} = 5.786, d_{2P} = 4.957$, implying that for the lower limit the nonparametric chart is chosen: $\widehat{LL}_C = \widehat{LL}_{NP}$. This means that we arrive at the following control limits

$$\begin{array}{lll}
\text{bias} & \widehat{LL}_C = \begin{cases} 22.139 & \text{with probability } 0.164 \\ 25.45 & \text{with probability } 0.836 \end{cases} & \widehat{UL}_C = 52.635 \\
\text{exceedance } FAR & \widehat{LL}_C = \begin{cases} 22.139 & \text{with probability } 0.749 \\ 25.45 & \text{with probability } 0.251 \end{cases} & \widehat{UL}_C = 52.842 \\
\text{exceedance } ARL & \widehat{LL}_C = \begin{cases} 22.139 & \text{with probability } 0.747 \\ 25.45 & \text{with probability } 0.253 \end{cases} & \widehat{UL}_C = 52.830.
\end{array}$$

5 Generalizing the exceedance criterion

Due to estimation or an inadequate model P_n, the (conditional) FAR, may be different from the prescribed one, p. The main problem arises when P_n is (substantially) larger than p and therefore the bias criterion may be replaced by

$$E\,(P_n - p)^+, \tag{9}$$

where $y^+ = \max(y, 0)$ and hence $E\,(P_n - p)^+$ looks only at those values of P_n that are larger than p. Define for $k = 0, 1, 2, ...$

$$E\left\{\left(Y^+\right)^k\right\} = E\left\{Y^k I\,(Y > 0)\right\},$$

that is we take the expectation of Y^k, but restrict attention to positive values of Y. (The function I is the indicator function, $I(A) = 1$ if A holds and 0 otherwise.) We consider as criterion for the performance of P_n

$$E\left[\{P_n - p\,(1+\varepsilon)\}^+\right]^k. \tag{10}$$

For $k = 0$, this gives

$$EI\,\{P_n - p\,(1+\varepsilon) > 0\} = \Pr\left(P_n > p\,(1+\varepsilon)\right)$$

and that is just the exceedance criterion (in a one-sided setting). For $k=1$ and $\varepsilon=0$, the expression (10) equals (9) and this is comparable to the so-called stop-loss criterion in insurance. For $k=2$ and $\varepsilon=0$ criterion (10) is known as the semi-variance.

When dealing with the criterion (10) we have to say what the requirement is. Of course, it should be small, but in which sense? To fix the idea, let Y be normally distributed and write $Y=bZ-a$ with $b>0$ and Z a random variable with a standard normal distribution Then

$$E\left\{\left(Y^{+}\right)^{k}\right\}=E\left\{(bZ-a)^{k} I\left(bZ-a>0\right)\right\}=b^{k} h_{k}\left(a/b\right)$$

with

$$h_{k}(x)=E\left\{(Z-x)^{k} I\left(Z>x\right)\right\}=\int_{x}^{\infty}(z-x)^{k}\varphi(z)dz,$$

where φ denotes the standard normal density. Note that $h_0(x)=1-\Phi(x), h_1(x)=\varphi(x)-x\left\{1-\Phi(x)\right\}, h_2(x)=\left(x^2+1\right)\left\{1-\Phi(x)\right\}-x\varphi(x)$. To set "universal "or "dimensionless" limits we say that $E\left\{\left(Y^{+}\right)^{k}\right\}$ is small if

$$E\left\{\left(Y^{+}\right)^{k}\right\}=b^{k} h_{k}\left(a/b\right)\leq \alpha b^{k}. \tag{11}$$

To see how control limits should be determined when dealing with (11) we assume normality and discuss the upper control limit. Writing the upper limit as $\overline{X}+u_p S\left\{1+c\right\}$ for some (small) c, still to be determined, we get, assuming that $X_1,\ldots,X_n,X_{n+1}$ are i.i.d. random variables with a $N(\mu,\sigma^2)$-distribution,

$$\begin{aligned} P_n &= 1-\Phi\left(\frac{\overline{X}-\mu}{\sigma}+u_p\frac{S}{\sigma}\left\{1+c\right\}\right) \\ &\approx p-\varphi\left(u_p\right)\left\{\frac{\overline{X}-\mu}{\sigma}+u_p\left(\frac{S}{\sigma}-1\right)+u_p c\right\}. \end{aligned}$$

Because $\left(\overline{X}-\mu\right)/\sigma+u_p\left(S/\sigma-1\right)$ is approximately $N(0,\tau^2/n)$ with $\tau^2=\left(u_p^2+2\right)/2$, we apply (11) with

$$\begin{aligned} Y &= P_n-p\left(1+\varepsilon\right) \\ &\approx -p\varepsilon-\varphi\left(u_p\right)\left\{\frac{\overline{X}-\mu}{\sigma}+u_p\left(\frac{S}{\sigma}-1\right)+u_p c\right\} \end{aligned}$$

and

$$b=\frac{\varphi\left(u_p\right)\tau}{\sqrt{n}}, a=p\varepsilon+u_p\varphi\left(u_p\right)c.$$

In order to satisfy (11) we take

$$c=\frac{h_k^{-1}(\alpha)b-p\varepsilon}{u_p\varphi\left(u_p\right)}.$$

For $k=0$ we arrive at

$$c=\frac{u_\alpha b-p\varepsilon}{u_p\varphi\left(u_p\right)}=\frac{u_\alpha\left(\frac{1}{2}+u_p^{-2}\right)^{1/2}}{\sqrt{n}}-\frac{p\varepsilon}{u_p\varphi\left(u_p\right)}\approx\frac{u_\alpha\left(\frac{1}{2}+u_p^{-2}\right)^{1/2}}{\sqrt{n}}-\frac{\varepsilon}{u_p^2}.$$

Indeed, replacing p by $p/2$ for the two-sided case, the latter expression is exactly the same as in Table 1.

In order to see the influence of k on the correction term c, here are the values of $h_k^{-1}(0.1)$ for $k=0,1,2,3,4$, respectively: $1.28, 0.90, 0.87, 0.96, 1.11$. For illustration we give the (one-sided) lower control limits calculated for the data of the razor heads (see Section 3) with $p=0.001$

(corresponding to $p = 0.002$ in the two-sided case), $\varepsilon = 0.1, \alpha = 0.1$. For $k = 0$ we get 31.889, see also (4). For $k = 1, 2, 3, 4$ we obtain 31.993, 32.003, 31.977, 31.937, respectively. This clearly shows that the control limits do not change much, also compared to the change due to another modeling approach: the corresponding parametric upper control limit equals (see (6)) 28.306, while the nonparametric one gives 22.139 with probability 0.749 and 25.45 with probability 0.251. Since the control limits for different values of k do not vary much, protection in terms of one k gives also more or less protection for the other ones. For instance, taking the correction term of the exceedance case with $\alpha = 0.1$ results in protection in terms of the semi-variance according to $\alpha = 0.04$, because $h_2(1.28) = 0.04$. Similarly, we get $\alpha = 0.05, 0.04, 0.06$ for $k = 1, 3, 4$, respectively.

6 Further discussion

Essentially the main problem in determining two-sided control limits in the setting presented here is to estimate the $(p/2)^{\text{th}}$- and $(1-p/2)^{\text{th}}$-quantile of the distribution for very small p. In principle we deal with $p = 0.002$ or a value of that order of magnitude.

Suppose one believes firmly in normality for the data at hand. Then it "only" remains to estimate the normal parameters μ and σ and to deal with the stochastic error involved by the estimation. When using bias as a criterion, the normal chart without using a correction is already fine for $n \geq 300$. But after applying the suitable correction, see Table 1, even $n \geq 40$ will do. If instead of bias one selects exceedance probabilty as a criterion, considerably larger boundary values of n must be reckoned with. For instance, even for $n = 2000$ and $\alpha = 0.2$ still the correction is not yet superfluous. For more details we refer to Albers and Kallenberg (2004b). It should be remarked that even ostensibly small deviations from normality may cause large model errors, which will seriously spoil matters.

Suppose one's belief in normality is somewhat less firm. This may be due to possible skewness or tail behavior different from normality. Then the normal power family comes into consideration. Using bias as a criterion, the parametric chart gives reasonable results for $n \geq 100$. This holds after application of a suitable correction, see Table 2, which is substantially more complicated than the one for the normal case. However, the calculations are quite straightforward. Again, switching to exceedance probabilities increases the required values of n considerably; for $n = 250$ there are still some fluctuations, but from $n = 500$ on, the results are quite satisfactory for practical purposes. More details can be found in Albers et al. (2002).

Suppose one does not really believe in normality and neither does one want to assume closeness to the normal power family. Then a (modified) nonparametric chart has to be applied. For a fully nonparametric approach huge sample sizes are needed to get an acceptable solution, because the stochastic error is very large here. The modified nonparametric chart, see Table 3, works rather well for all kind of distributions occurring in practice. But some loss under out-of-control has to be accepted when the observations are normally distributed or close to that. More details can be found in Albers and Kallenberg (2004c) for the fully nonparametric chart and in Albers et al. (2003) for the modified version.

Since it is not easy to judge "by hand" whether normality or the parametric model of the normal power family can be assumed to be sufficiently adequate, it is more promising to let the data speak by themselves. That means application of the combined control chart. In the bias case this data driven procedure works very well for $n \geq 250$ in the sense that the selection rule is successful. If normality or the normal power family model holds after all, the loss w.r.t. the corresponding specific control charts is limited, while outside these models the occurrence of huge model errors is indeed effectively avoided. Therefore, we recommend to use this combined chart. For smaller values of n one should either be sufficiently convinced that normality or the normal power family model can be used, or one should switch to a less ambitious value of p, like $p = 0.01$ or that order of magnitude.

Acknowledgements The authors cordially thank dr. ir. J. Praagman (CQM Eindhoven) for providing us with the electric shaver data used in this paper.

References

Albers, W. and Kallenberg, W.C.M. (2003). New corrections for old control charts. Memorandum 1694, Faculty of Mathematical Sciences, University of Twente.

Albers, W. and Kallenberg, W.C.M. (2004a). Estimation in Shewhart control charts: effects and corrections. To appear in *Metrika*.

Albers, W. and Kallenberg, W.C.M. (2004b). Are estimated control charts in control? *Statistics* **38** 67–79.

Albers, W. and Kallenberg, W.C.M. (2004c). Empirical nonparametric control charts: estimation effects and corrections. Memorandum 1651, Faculty of Mathematical Sciences, University of Twente. To appear in *J. Appl. Statist.*

Albers, W., Kallenberg, W.C.M. and Nurdiati, S.(2002) Exceedance probabilities for parametric control charts. Memorandum 1650, Faculty of Mathematical Sciences, University of Twente.

Albers, W., Kallenberg, W.C.M. and Nurdiati, S.(2003) Normal, parametric and nonparametric control charts: a data driven choice. Memorandum 1674, Faculty of Mathematical Sciences, University of Twente.

Albers, W., Kallenberg, W.C.M. and Nurdiati, S.(2004) Parametric control charts. To appear in *J. Statist. Plann. Inference.*

Barndorff-Nielsen, O.E. (1996) Normal inverse Gaussian distributions and stochastic volatility modelling. *Scand. J. Statist.* **24** 1–13.

Chan, L.K., Hapuarachchi, K.P. and Macpherson, B.D. (1988). Robustness of $\overline{X}$ and R charts. *IEEE Trans. Reliability* **37** 117-123.

Jones, L.A., Champ, C.W. and Rigdon, S.E. (2004). The run length distribution of teh CUSUM with estimated parameters. *J. Quality Technol.* **36** 95-108.

Pappanastos, E.A. and Adams, B.M. (1996). Alternative designs of the Hodges-Lehmann control chart. *J. Quality Technol.* **28** 213-223.

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

Roes, C.B. (1995). *Shewhart-type Charts in Statistical Process Control.* Ph. D. thesis University of Amsterdam.