Signed Sequential Rank Shiryaev-Roberts Schemes

C. van Zyl
Centre for Business Mathematics and Informatics
North-West University, Potchefstroom 2520
F. Lombard
Department of Statistics
University of Johannesburg, Johannesburg 2000

Abstract

We develop Shiryaev-Roberts schemes based on signed sequential ranks to detect a persistent change in location of a continuous symmetric distribution with known median. The in-control properties of these schemes are distribution free, hence they do not require a parametric specification of an underlying density function or the existence of any moments. Tables of control limits are provided. The out-of-control average run length properties of the schemes are gauged via theory-based calculations and Monte Carlo simulation. Comparisons are made with two existing distribution-free schemes. We conclude that the newly proposed scheme has much to recommend its use in practice. Implementation of the methodology is illustrated in an application to a data set from an industrial environment.

Keywords: Distribution free, CUSUM, Sequential ranks, Shiryaev-Roberts, Gordon-Pollak

1 Introduction

Statistical process control schemes, including CUSUMs and Shiryaev-Roberts schemes, are designed to signal a persistent change in some characteristic of a process as soon as possible after its onset. The change manifests itself as a sustained change in the distribution along a sequence of independent and identically distributed observations $X_1, X_2, \ldots$ which occurs at an index τ, known as the the change point. Thus $X_1, X_2, \ldots, X_\tau$, the "in-control" observations, have a common density function f while $X_{\tau+1}, X_{\tau+2}, \ldots$, the "out-of-control" observations, have common density $g \neq f$. Perhaps the most widely known CUSUM is the one introduced by Page (1954) which aims to detect an increase in the mean from $\delta = 0$ to a value $\delta > 0$ in a normal distribution with known variance σ^2. The CUSUM sequence C_i, $i \geq 0$, is defined by setting $C_0 = 0$ and applying the recursion

$$C_i = \max\left\{0, C_{i-1} + \frac{X_i}{\sigma} - \zeta\right\}, \ i \geq 1. \tag{1}$$

An alarm is raised as soon as C_i exceeds a control limit $h > 0$, indicating that the mean has possibly increased and that an out-of-control situation prevails. Here, ζ is a positive reference value or drift allowance. This is often set equal to $\delta_1/2$ where δ_1, the target shift, is the smallest change (drift) in the mean of X/σ that is judged to be of practical import. Alternatively, ζ is sometimes chosen with a view to producing a specified form of out-of-control behaviour. The CUSUM behaves like a random walk with a reflecting barrier at 0 and is sure to produce an alarm somewhere along the sequence even if no change ever occurs. In the latter case, we have a false alarm. Because false alarms are inevitable, the control limit h is chosen so that a predetermined expected run length, known as the in-control average run length (ICARL), is guaranteed if no change ever occurs.

Girshick and Rubin (1952) introduced a control scheme based on a recursion that can be expressed as

$$D_i = log(1 + \exp(D_{i-1})) + 2\zeta\left(\frac{X_i}{\sigma} - \zeta\right), \ i \geq 1 \tag{2}$$

with $D_0 = -\infty$. A continuous time version of this scheme was developed by Shiryaev (1963) with the objective of detecting the onset of a drift in a Brownian motion. Roberts (1966) compared the performance of the CUSUM and some other schemes with the scheme (2) and found the latter to have merit. Schemes based on the recursion (2) have subsequently become known in the literature as Shiryaev-Roberts (S-R) schemes.

The assumption of normality and of a known variance are enduring problems in the application of these schemes to observed data. Versions that are distribution free under a broad class of in-control distributions would be extremely useful in statistical practice. Lombard and Van Zyl (2018) developed distribution-free CUSUMs that can detect deviations from a specified median in a symmetric continuous distribution. The scale parameter does not figure directly in the construction of these CUSUMs and no moment conditions are required to assure their validity. This development frees one from the restriction to an underlying normal distribution and the difficulties in designing CUSUMs that operate efficiently when the variance is estimated or the distribution is non–normal; see, for instance, Bagshaw and Johnson (1974), Jones et Al. (2004) and Diko et Al. (2020).

Gordon and Pollak (1994) constructed a distribution free S-R type scheme which they dubbed the NPSR (non parametric Shiryaev-Roberts) scheme, for detecting shifts away from a known median in a symmetric distribution. The NPSR is based on the signs of the data and the ranks of their absolute values and requires specification of three adjustable parameters. These permit it being "tuned" to any given symmetric distribution. However, the NPSR does not seem to have enjoyed widespread adoption among practitioners. This is possibly due to the complexity of the scheme which does not allow a simple recursion such as (2) and to computational difficulties surrounding the generation of control limits for a wide range of ICARLs. These difficulties seem to have thus far inhibited a fuller evaluation of the NPSR's properties.

Following ideas from Lombard and Van Zyl (2018), we propose in this paper an alternative distribution-free S-R scheme that is based on the signs s_i of the observations and on the *sequential* ranks r_i^+ of their absolute values. The sequential rank r_i^+ of $|X_i|$ in the sequence $|X_1|, \ldots, |X_i|$ is the number of observations that are less than or equal to $|X_i|$,

$$r_i^+ = 1 + \sum_{j=1}^{i-1} \mathbf{1}\left(|X_j| < |X_i|\right), \ i \geq 2$$

where $\mathbf{1}\left(\cdot\right)$ denotes the indicator function and $r_1^+ = 1$. Under any continuous in-control distribution, successive sequential ranks are mutually independent and uniformly distributed on the integers $1, 2, \ldots i$, hence are distribution free - see Barndorff-Nielsen (1963, Theorem 1.1). Furthermore, the r_i^+ sequence is then also independent of the sequence of signs

$$s_i = \mathbf{1}\,(X_i > 0) - \mathbf{1}\,(X_i < 0)$$

see Reynolds (1975, Theorem 2.1). Thus, upon replacing X_i in (2) by a function $J(R_i^s)$ of the signed sequential ranks

$$R_i^s = \frac{s_i r_i^+}{i+1}, \tag{3}$$

a new family of distribution-free S-R type change detection schemes appears. In view of their distribution free property, tables of control limits are easily generated by Monte Carlo simulation. Furthermore, by an appropriate choice of the function J, the scheme can be tuned to have good properties in any given distribution. For instance, if the in-control distribution is expected to be near normal, then upon replacing the X_i in (2) by

$$X_i^* = \Phi^{-1}((1 + R_i^s)/2), \tag{4}$$

where Φ^{-1} denotes the inverse of the standard normal cdf, there results a distribution-free S-R scheme that can be expected to be competitive with the parametric version when the data indeed come from a normal distribution. A key fact is that the in-control distribution of $(1+R_i^s)/2$ approximates a uniform distribution as i increases, whence the distribution of X_i^* approximates that of a normally distributed X_i. Furthermore, the out-of-control distribution of X_i^* approximates a shifted normal distribution. This construction is reminiscent of the manner in which the Van der Waerden two sample signed-rank statistic is obtained - see Hájek, Šidák and Sen (1999, page 118).

The paper is structured as follows. In Section 2 the original Shiryaev-Roberts scheme and the Gordon and Pollak (1994) NPSR scheme are discussed. In Section 3, we introduce the signed sequential rank S-R schemes, hereafter referred to as SSR S-R schemes. These schemes are constructed specifically with a view to detecting a change away from a *known* median in an unspecified *symmetric* distribution, thus generalizing the normal distribution based S-R scheme. Tables of control limits guaranteeing a nominal in-control ARL are provided. In Section 4 we discuss the specification of an appropriate reference value ζ in the sequential rank schemes and in Section 5 we deal in some detail with the out-of-control run length properties of the Wilcoxon SSR S-R scheme, that is, the SSR S-R scheme based on the Wilcoxon score $J(u) = u$. In particular we indicate how its out-of-control behaviour may be assessed on a theoretical basis by using both formal and informal calculations. It is seen that the Wilcoxon SSR S-R scheme is quite efficient in normal distributions and can usefully serve as an omnibus distribution-free scheme. The results of some pertinent Monte Carlo simulations are reported in Section 6. Specifically, the performance of the Wilcoxon SSR S-R scheme is compared to that of the Wilcoxon SSR CUSUM and the NPSR scheme. Finally, in Section 7, implementation of the schemes is illustrated by application to a new data set that has not been considered in the literature before. In Section 8 we provide a summary of our main results and conclusions and indicate some areas for further research.

2 The Shiryaev-Roberts and Gordon-Pollak schemes

An important aspect of any sequential change detection scheme is the extent to which the realized ICARL agrees with the nominal value, which we denote generically by ARL_0. Successful implementation of the normal S-R scheme (2) requires that σ be known and that the underlying distribution be normal. We now examine the extent to which deviations from these assumptions affect its in-control behaviour. Suppose first that σ $(= 1)$ is unknown to the analyst and that a Phase I estimated standard deviation $\hat{\sigma}_m$, computed from m observations, is used as a proxy for σ. Then the analyst will rescale the Phase II data, replacing X_i by $X_i/\hat{\sigma}_m$, and run the S-R scheme on the rescaled data. Consequently, since $X_i/\hat{\sigma}_m$ does not have unit variance, the ICARL will differ from the nominal ARL_0. To illustrate, suppose $ARL_0 = 500$ is desired and that an estimate $\hat{\sigma}_{50} = 1.1$ has been found from some Phase I data. Application of the S-R scheme to the rescaled data will then produce an ICARL of 1081 (estimated from 10^5 Monte Carlo trials). To provide some context to this result we note that when $\sigma = 1$,

$$Pr\left[\hat{\sigma}_{50} > 1.1\right] = Pr\left[\hat{\sigma}_{50} < 0.9\right] = 0.125,$$

which means that about one in every four estimates will be larger that 1.1 or smaller than 0.9. Table 1a shows Monte Carlo estimated true ICARL values for $m = 50$, two estimates $\hat{\sigma}_{50}$ of σ, two reference values ζ and three ARL_0 values.

Table 1a Estimated ICARL when σ is estimated from 50 observations.

		ARL_0		
$\hat{\sigma}_{50}$	ζ	100	500	1,000
1.1	0.25	**130**	**875**	**1,983**
1.1	0.5	**158**	**1,081**	**2,495**
0.9	0.25	**75**	**305**	**544**
0.9	0.5	**64**	**245**	**424**

The sizes of the discrepancies between nominal and true ICARLs are clearly a cause for concern. To ameliorate the situation one can increase the size of the sample on which the estimate is based. However, doubling the number of Phase I observations does not improve the situation much - see Table 1b, in which

$$Pr\left[\hat{\sigma}_{100} > 1.08\right] = Pr\left[\hat{\sigma}_{50} < 0.92\right] = 0.125.$$

Table 1b Estimated ICARL when σ is estimated from 100 observations.

		ARL_0		
$\hat{\sigma}_{100}$	ζ	100	500	1,000
1.08	0.25	**125**	**764**	**1,723**
1.08	0.5	**147**	**923**	**2,066**
0.92	0.25	**80**	**334**	**609**
0.92	0.5	**70**	**282**	**507**

Suppose next that the variance is known to equal 1 but that the in-control distribution is non-normal. The last three columns in Table 2 show Monte Carlo estimates of the ICARL values of the normal distribution S-R scheme when the data actually come from logistic, Laplace and Student t_3 distributions, all standardised to unit variance. Each of the estimates was made on 10^5 Monte Carlo trials. In the table, the first three columns show the reference values $\zeta = 0.1$, 0.15, 0.5 and 0.75, the nominal ICARL values $ARL_0 = 500$ and $ARL_0 = 1,000$ and the corresponding control limits h that guarantee an ICARL equal to ARL_0 under a normal distribution.

The differences between the estimated true in-control ARLs and the nominal ARL_0 values could be considered acceptable "for practical purposes" only in the logistic and Laplace cases at $\zeta = 0.1$, which is a reference value that would not be used frequently. The results shown in Tables 1a and 1b and in Table 2 clearly indicate that distribution-free and scale invariant S-R schemes would be valuable additions to the toolbox of a practitioner who contemplates using an S-R type scheme.

Table 2 Comparison of estimated true and nominal ICARL of Shiryaev-Roberts scheme in non-normal distributions.

			Estimated ICARL		
ζ	ARL_0	h	Logistic	Laplace	t_3
0.1	500	6.10	**512**	**513**	**581**
	1,000	6.79	**1007**	**1020**	**1156**
0.25	500	5.92	**484**	**483**	**560**
	1,000	6.62	**961**	**946**	**970**
0.5	500	5.63	**418**	**353**	**362**
	1,000	6.32	**804**	**622**	**527**
0.75	500	5.35	**330**	**244**	**242**
	1,000	6.06	**598**	**392**	**334**

To the best of our knowledge the first, and to date only, distribution-free S-R type scheme applicable to symmetric distributions with known median is the NPSR of Gordon and Pollak (1994). The NPSR is based on a double sequence Λ_k^n, $1 \leq k \leq n$, $n \geq 1$, of nonlinear two-sample rank and sign statistics, not on signed ranks alone. An expression for Λ_k^n, which has a somewhat complicated structure, can be found in eqn. (6) of Gordon and Pollak (1994) who also provide efficient Matlab code for the calculation. The run length is the first index n at which the sum $R_n = \sum_{k=1}^{n} \Lambda_k^n$ exceeds a control limit, A, which makes the ICARL equal to a specified number ARL_0. Given a large nominal ARL_0, Gordon and Pollak (1994) provide an approximation to A in terms of ARL_0 and a further three adjustable parameters. However, the approximation needs to be supplemented by Monte Carlo simulations to find a value of A that produces an ICARL sufficiently close to ARL_0 for use in a practical application. Here one encounters two kinds of difficulties. The first difficulty is that the complicated structure of R_n results in excessively time consuming simulations at large ARL_0 values. The second, and perhaps more important, difficulty is that a small fraction of simulated run lengths are so large that the only way in which a practicable scheme results is if the run length is truncated. Gordon and Pollak truncated all run lengths at $2,500$. In our simulations we truncated all NPSR run lengths at $5,000$ and encountered a negligible proportion of NPSR run lengths that had to be truncated.

3 Signed sequential rank schemes

The independence, distribution freeness and naturally sequential nature of signed sequential ranks R_i^s from (3) makes them ideally suited to the construction of CUSUM and S-R schemes for independently distributed time ordered data. A class of signed sequential rank analogues of (1) and (2) is obtained upon replacing X_i there by

$$X_i^* = \frac{J\left(R_i^s\right)}{v_i} \tag{5}$$

where $J(u),\ -1 < u < 1$ is an odd, square-integrable, function on the interval $(-1, 1)$ and where

$$v_i = \sqrt{\frac{1}{i}\sum\nolimits_{j=1}^{i} J^2\left(\frac{j}{i+1}\right)}. \tag{6}$$

It is customary in the rank statistic literature to refer to $J(u)$ as a score function. In particular, the Wilcoxon score

$$J_W(u) = u,\ -1 \leq u \leq 1$$

leads to a particularly useful omnibus scheme. In this case

$$\nu_i = \sqrt{6(i+1)/(2i+1)}$$

and the corresponding SSR S-R and CUSUM schemes are defined by

$$D_i = \log(1 + \exp\left(D_{i-1}\right)) + 2\zeta\left(\frac{R_i^s}{\nu_i} - \zeta\right) \tag{7}$$

and

$$C_i = \max\left(0, C_{i-1} + \frac{R_i^s}{\nu_i} - \zeta\right). \tag{8}$$

For underlying distributions close to the normal, one could use an SSR scheme based on the Van der Waerden score already mentioned in the introduction, namely

$$J_V(u) = \Phi^{-1}((1+u)/2).$$

However, the Pearson correlation coefficient between $J_W(R_i^s)$ and $J_V(R_i^s)$ tends as $i \to \infty$ to $\sqrt{\pi/3} = 0.98$, which implies that not much will be lost if the computationally convenient J_W is used in place of the somewhat more complicated J_V. Furthermore the large sample correlations between the Wilcoxon score and the efficient scores in some heavy-tailed Student t-distributions are quite high while their correlations with the Van der Waerden score are somewhat lower. Table 3 shows these correlations in four t-distributions.

Table 3 Correlations of efficient scores in Student t_ν distributions with Wilcoxon and Van der Waerden scores.

	Distribution				
Score	*normal*	t_4	t_3	t_2	t_1
Wilcoxon	0.98	0.99	0.98	0.94	0.79
Van der Waerden	1.00	0.95	0.92	0.86	0.67

Consequently, the SSR S-R and SSR CUSUM schemes that are based on the Wilcoxon score J_W can be expected to be quite efficient in a broad range of symmetric underlying distributions, obviating to a large extent the necessity of tuning the scheme to any specific distribution. We will refer to them as the Wilcoxon SSR S-R and the Wilcoxon SSR CUSUM respectively. The latter one of these was dealt with in detail by Lombard and Van Zyl (2018). The term VdW SSR S-R will be used to denote the SSR S-R scheme that is based on the Van der Waerden score J_V. Construction of these schemes does not require knowledge of the numerical value of any scale parameter σ because signed sequential ranks are scale invariant.

Tables 4 and 5 give control limits for a matrix of (ζ, ARL_0) pairs for the Wilcoxon SSR S-R and CUSUM schemes. The tables were generated by Monte Carlo simulation using the method detailed in Lombard and Van Zyl (2018).

Table 4 Control limits for the Wilcoxon SSR S-R .

	ARL_0						
ζ	100	200	300	400	500	1,000	2,000
0.05	4.55	5.24	5.65	5.93	6.16	6.85	7.55
0.10	4.49	5.18	5.60	5.86	6.10	6.80	7.48
0.15	4.43	5.14	5.53	5.83	6.05	6.73	7.24
0.20	4.37	5.07	5.47	5.76	5.98	6.68	7.38
0.25	4.31	5.01	5.41	5.70	5.92	6.59	7.27
0.30	4.29	4.95	5.34	5.61	5.83	6.51	7.18
0.35	4.21	4.86	5.28	5.55	5.74	6.41	7.08
0.40	4.13	4.78	5.17	5.43	5.66	6.31	6.96
0.45	4.04	4.68	5.07	5.33	5.54	6.18	6.83
0.50	3.95	4.58	4.95	5.24	5.43	6.03	6.69
0.75	3.38	3.95	4.26	4.50	4.67	5.22	5.77
1.00	2.65	3.09	3.36	3.55	3.69	4.12	4.54

Table 5 Control limits for the Wilcoxon SSR CUSUM.

	ARL_0						
ζ	100	200	300	400	500	1,000	2,000
0.00	8.92	13.07	16.24	18.9	21.3	30.24	43.95
0.05	7.61	10.51	12.49	14.01	15.33	19.89	25.03
0.10	6.45	8.62	10.05	11.12	12.01	14.79	17.93
0.15	5.65	7.34	8.42	9.21	9.86	11.88	14.06
0.20	5.00	6.37	7.24	7.87	8.37	9.96	11.57
0.25	4.46	5.61	6.33	6.85	7.25	8.52	9.84
0.30	4.01	5.00	5.60	6.03	6.37	7.45	8.53
0.35	3.62	4.48	5.00	5.37	5.66	6.58	7.51
0.40	3.29	4.04	4.49	4.81	5.06	5.87	6.66
0.45	2.99	3.66	4.05	4.34	4.56	5.25	5.96
0.50	2.73	3.31	3.68	3.93	4.13	4.74	5.34
0.75	1.72	2.06	2.2	2.42	2.53	2.89	3.25
1.00	1.02	1.24	1.34	1.42	1.49	1.71	1.92

4 Specification of the reference value

Consider a situation in which the median shifts after a considerable time τ from 0 to $\delta \neq 0$. Then, with

$$\xi_i = \frac{R_i^s}{\nu_i} = \frac{s_i r_i^+}{v_i(i+1)},$$

a calculation which is detailed in the Appendix shows that for τ large and δ small,

$$E[\xi_{\tau+1}] \approx \theta_0 \delta, \tag{9}$$

with

$$\theta_0 = \sqrt{12} \int_{-\infty}^{\infty} f^2(x) dx, \tag{10}$$

f denoting the pdf of the in-control distribution. In analogy with the parametric scheme (2), the relation (9) suggests $\zeta = \delta_1 \theta_0 / 2$ as an appropriate choice for targeting a shift of size δ_1. Some values of θ_0 that are likely to be encountered in practice are given in Table 6. In the case of the t_2 and t_1 (Cauchy) distributions the interquartile range served as the scale parameter, that is, the shift δ is expressed in units of the interquartile range. Even though the density function underlying the data is unknown, we can still use these values of θ_0 to make a priori an informed choice of θ_0, hence of ζ, based on the expected heaviness of the tails. On the other hand, if some Phase I data are available, θ_0 can be estimated non-parametrically as indicated in Lombard and Van Zyl (2018, Section 3.1).

Table 6 Values of θ_0 for a range of symmetric distributions.

Distribution					
normal	Laplace	t_4	t_3	t_2	t_1
0.98	1.2	1.18	1.37	1.18	1.10

5 Run length properties

An important practical requirement is to assess a priori the properties of the run length

$$N = min\{i \geq 1;\ D_i \geq h\} \tag{11}$$

with D_i from (7). Given any specific out-of-control density or range of densities, the behaviour of N can be assessed quite simply by Monte Carlo simulation. The question nevertheless arises whether it is possible to obtain useful conclusions based on less information. To see that this is indeed possible, we observe that D_i can be expressed as a functional of the partial sums $S_i = \xi_1 + \cdots + \xi_i$, namely

$$exp\,(D_i) = \sum_{j=1}^{i-1} exp\left\{2\zeta\,(S_i - S_j) - \zeta^2(i-j)\right\}$$

- see, for instance, Pollak and Siegmund (1991, page 396). Because the partial sums S_i are for large i approximately normally distributed, we can expect the behaviour of a Wilcoxon SSR S-R at a small out-of-control shift δ to be close to that of a normal S-R scheme (2) with the same ζ and h at a shift $\theta_0\delta$, provided ARL_0 and the change point τ are large - see (9) and (10). Large values of ARL_0, hence h, and τ allows enough time for a normal approximation to manifest itself. Indeed, Lombard and Van Zyl (2018, Appendix) showed via a continuous time approximation involving Brownian motion that under these conditions the ARL of any SSR CUSUM behaves in the indicated manner. They also provided supporting numerical evidence - see their Tables 4.1 and 4.2 and Tables S3 and S4 in the Supplementary Material. The same method substantiates the conclusion in respect of general SSR S-R schemes and the following numerical evidence provides support specifically in respect of the Wilcoxon SSR S-R scheme.

The out-of-control performance criterion we use here is the conditional average delay time (CADT) $E_\tau[N-\tau|N>\tau]$, where the subscript τ in E_τ denotes that the expected value is computed under the assumption that the change occurs at time $t=\tau+1$. The ARL_0 value was set at 500, the target out of control shifts were $\delta_1 = 0.2$ and $\delta_1 = 0.5$. The Laplace distribution and the much heavier tailed t_3 distributions both served as in-control distributions. The design parameters θ_0 (the tuning constant) and $(\zeta, h) = (\delta_1\theta_0/2, h)$ =(reference constant,control limit) are shown in the second row of Tables 7a and 7b. The entries in the tables are Monte Carlo estimates (10^5 trials) of the conditional average delays, $\mathcal{W}(\delta)$ and $\mathcal{N}(\delta\theta_0)$, of the Wilcoxon SSR S-R and normal S-R schemes respectively at a range of out-of-control means δ. The results are shown for changes of size δ at change points $\tau = 0$ and $\tau = 100$. If the out-of-control behaviour of an SSR S-R scheme is indeed similar to that of a normal S-R scheme, using normal data only, with a θ_0-adjusted out-of-control shift, we would expect to see

$$\mathcal{W}(\delta) \approx \mathcal{N}(\delta\theta_0) \qquad (12)$$

to good approximation if δ is "small" and τ is "large". Inspection of the results in Table 7a indicates that the approximation is quite satisfactory at $\tau = 100$. Also, though rather unexpectedly, the approximation is also quite good at $\tau = 0$ (Table 7b) except at $\delta \geq 0.5$, that is, when the underlying process is substantially out-of-control from the outset.

Table 7a Wilcoxon SSR S-R CADT approximations in Laplace and t_3 distributions. $ARL_0 = 500$; change point $\tau = 100$.

	Laplace: ($\theta_0 = 1.2$)				t_3: ($\theta_0 = 1.37$)			
(ζ, h)	$(0.12, 6.09)$		$(0.3, 5.83)$		$(0.15, 6.05)$		$(0.35, 5.74)$	
δ	$W(\delta)$	$N(\theta_0\delta)$	$W(\delta)$	$N(\theta_0\delta)$	$W(\delta)$	$N(\theta_0\delta)$	$W(\delta)$	$N(\theta_0\delta)$
0.125	91	95	124	123	85	85	116	118
0.25	44	45	49	50	38	37	43	42
0.5	22	21	19	18	18	17	16	15
0.75	15	14	12	11	12	11	10	8
1.0	12	11	10	8	10	8	8	6
1.5	10	7	7	7	8	6	6	5

Table 7b Wilcoxon SSR S-R CADT approximations in Laplace and t_3 distributions. $ARL_0 = 500$; change point $\tau = 0$.

	Laplace: ($\theta_0 = 1.2$)				t_3: ($\theta_0 = 1.37$)			
(ζ, h)	$(0.12, 6.09)$		$(0.3, 5.83)$		$(0.15, 6.05)$		$(0.35, 5.74)$	
δ	$W(\delta)$	$N(\theta_0\delta)$	$W(\delta)$	$N(\theta_0\delta)$	$W(\delta)$	$N(\theta_0\delta)$	$W(\delta)$	$N(\theta_0\delta)$
0.125	124	125	135	132	107	107	126	118
0.25	67	65	60	56	55	53	**54**	**47**
0.5	**40**	**34**	**27**	**22**	**32**	**27**	**24**	**18**
0.75	**31**	**24**	**19**	**14**	**26**	**18**	**17**	**11**
1.0	**28**	**19**	**16**	**10**	**23**	**14**	**14**	**8**
1.5	**25**	**13**	**14**	**7**	**20**	**10**	**12**	**6**

The preceding results indicate that in the absence of a specified out-of-control density, useful estimates of out-of-control ARLs can be had if an estimate of θ_0 is available. The behaviour of the ARLs described above and seen in Tables 7a and 7b can be deduced from formal limit theorems involving contiguous and fixed alternatives. The approximation (12) is an informal interpretation of Theorem 1 in Lombard (1981) which deals with the convergence in distribution of signed sequential rank statistics under contiguous alternatives to a drifted Brownian motion. Here, "contiguous" is interpreted informally as indicating that δ and the reference value ζ are "small" when h is "large" - see also the Appendix in Lombard and Van Zyl (2018). The failure of (12) when τ $(= 0)$ is "small" and δ (≥ 0.5) is "large" is to a large extent explained by Theorem 1.1 in Müller-Funk (1983) and Theorem 2 in Lombard and Mason (1985) which deal with the convergence in distribution of sequential rank statistics under *fixed* alternatives. In this case there is distributional convergence to Gaussian processes, which are not simply drifted Brownian motions, and for which run length distributions are not available. As far as we are aware, similar approximation results are not available for the NPSR.

To summarize, there are strong indications (except when the underlying process is substantially out of control from the outset) that the out of control behaviour of the two Wilcoxon SSR schemes can usefully be gauged from the behaviour of their normal distribution counterparts by the simple device of replacing in the latter any shift δ by $\hat{\theta}_0\delta$ where $\hat{\theta}_0$ is an estimator of θ_0 made from some Phase I data. Thus, for instance, if an estimate of the CADT at a given shift δ and change point τ is required, then this can be had by simulating observations $X_1, \ldots, X_\tau$ from a normal$(0, 1)$ distribution and observations $X_{\tau+1}, X_{\tau+2}, \ldots$ from a normal$(\hat{\theta}_0\delta, 1)$ distribution.

6 Simulation results

In a numerical comparison of the normal distribution S-R and CUSUM schemes, Moustakides, Polunchenko and Tartakovsky (2009) conclude that the only marked difference in out-of-control performances is seen at small shifts. The approximation argument formulated in Section 5 suggests that this may also be the major difference between the signed sequential rank analogues of the two schemes. To investigate this suggestion, we compared in a simulation study the performances of the Wilcoxon SSR S-R with the Wilcoxon SSR CUSUM of Lombard and Van Zyl (2018) and with the NPSR of Gordon and Pollak (1994). We report below some of the results that are indicative of the behaviours of the schemes. When there is no danger of confusion, we will often refer in what follows to the S-R, the CUSUM and the NPSR, dropping the "Wilcoxon" and "SSR" prefixes as well as the "scheme" suffix.

All three schemes were tuned to the detection of mean shifts in two distributions, namely, the normal distribution and the Laplace distribution. The use of the Laplace distribution is justified in view of the inclusion of the

NPSR in the comparisons. The NPSR is derived from a mixture of two exponential distributions, the null instance of which is the Laplace distribution. To begin with, we limit our reporting to target sizes $\delta_1 = 0.5$ and $\delta_1 = 0.25$ and $ARL_0 = 500$. In the simulations, each estimated CADT of the S-R and CUSUM is the result of $10,000$ independent Monte Carlo trials while the estimates for the NPSR resulted from $5,000$ independent trials. The smaller number of trials used in the case of the NPSR was necessary in order to keep its runtimes within reasonable bounds. The somewhat "wobbly" appearance of some of the NPSR plots is a consequence of the reduced number of trials, but we do not believe that this misrepresents the true behaviour of the NPSR. In all the figures that follow, the dotted line represents the S-R, the dashed line represents the NPSR and the solid line the CUSUM.

Beginning with the normal distribution, the tuning parameters used for a target shift $\delta_1 = 0.5$ were $\zeta = 0.25$ for the Wilcoxon SSR S-R and CUSUM and $(\alpha, \beta, p) = (0.735, 1.324, 0.691)$ for the NPSR, the latter found by numerical calculation from equations (10) and (15) in Gordon and Pollak (1994). The control limits were $h = 5.92$ for the S-R, $h = 7.25$ for the CUSUM and $A = 375$ for the NPSR. For a target shift $\delta_1 = 0.25$ the tuning parameters used were $\zeta = 0.125$ and $(\alpha, \beta, p) = (0.860, 1.155, 0.0.599)$ with control limits $h = 6.07$ for the S-R, $h = 10.94$ for the CUSUM and $A = 455$ for the NPSR. Suppose first that the data do come from a normal distribution. Figures 1 and 2 show plots of the CADTs against a series of change points $\tau = 0 : 50 : 500$ involving actual shifts of sizes $\delta = 0.125$ and $\delta = 0.5$. Clearly, the differences between the CADTs of the schemes are largest when the underlying process is out of control from the outset. Also, while the three schemes perform similarly when the actual shift is equal to the target ($\delta = \delta_1$, Figure 2), the CUSUM seems to fare rather poorly compared to the other two when the actual shift is substantially less than the target ($\delta = \delta_1/4$, Figure 1). A striking feature in both plots is that the CADTs seem to have reached stationary values at or near change point $\tau \approx 50$. This feature also appeared in numerous other configurations (not shown here).

In Figures 3 and 4 the CADTs at two target shifts $\delta_1 = 0.25$ and $\delta_1 = 0.5$ are plotted against a series of actual shift sizes occurring at a change point $\tau = 100$. We notice that, as was seen in Figures 1 and 2, the CADT of all three schemes increases as the actual shift decreases further away from the target. However, this is not necessarily an indicator of poor performance because a change of size much less than the target is often deemed to be undesirable or even as constituting a false alarm. Clearly, the S-R and NPSR again perform quite similarly at small shifts but with somewhat smaller CADTs there than the CUSUM.

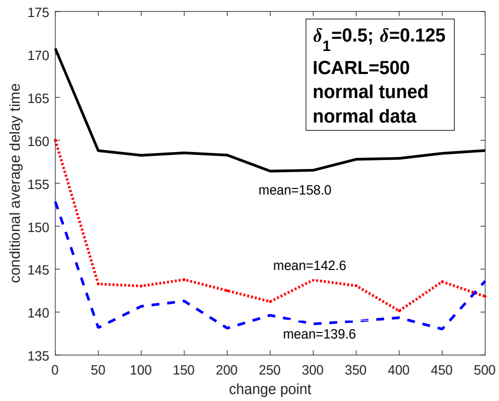

Figure 1

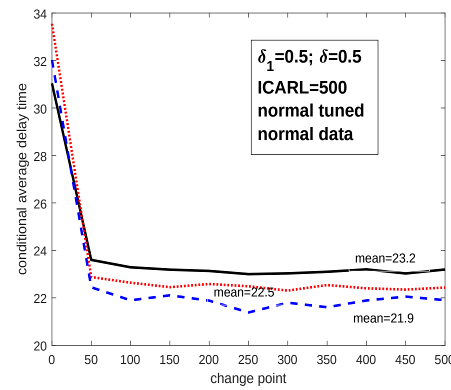

Figure 2

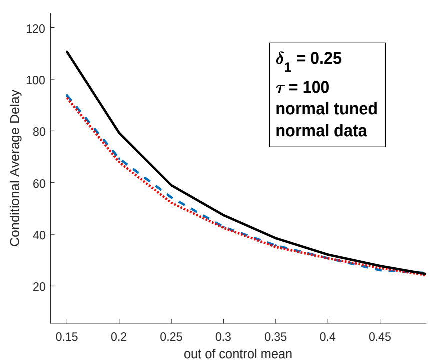

Figure 3

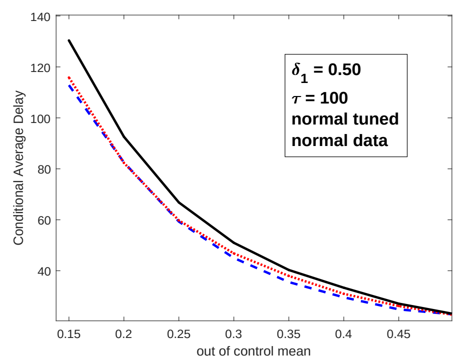

Figure 4

Figures 1 to 4 Conditional average delay simulation results.
Tuned for normal distribution.
Data from a normal distribution.

The preceding discussion indicates that the differences between the CADT of the schemes are relatively small when all three are tuned for data coming from a normal distribution. However, it is rather interesting to see what transpires when the data actually come from a non-normal distribution, that is, when the schemes have been tuned to the wrong distribution. Figures 5 and 6 are the counterparts of Figures 1 and 2 for data coming from a Laplace distribution. The most obvious difference between the two sets of figures is that the NPSR has, in a manner of speaking, gone from "best" to "worst". All three schemes detect the change in the mean of the data from the heavier tailed Laplace distribution sooner than before, but the improvement in the CADT of the S-R and CUSUM schemes is markedly better than that of the NPSR. This finding could perhaps be paraphrased by saying that the two SSR schemes adapt better to erroneous tuning than the NPSR.

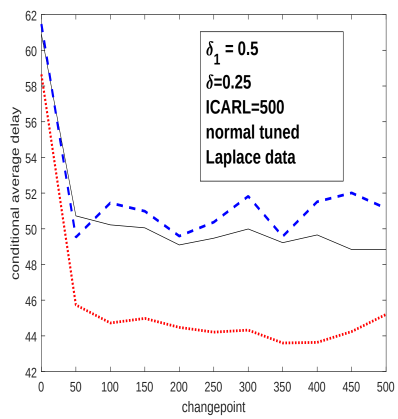

Figure 5

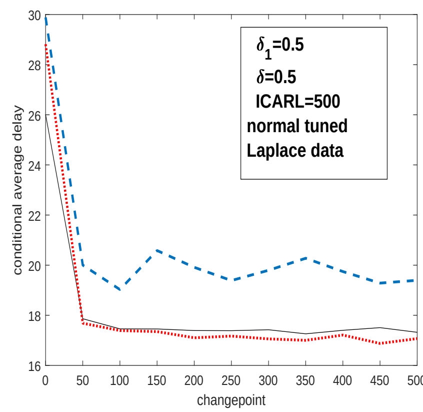

Figure 6

Figures 5 and 6 Conditional average delay simulation results.
Tuned for normal distribution.
Data from a Laplace distribution.

Next, consider a situation in which the procedures are tuned to a Laplace distribution, for which $\theta_0 = 1.2$ in the SSR schemes. As mentioned earlier, the use of the Laplace distribution is motivated by the NPSR scheme which is derived from a mixture of two exponential distributions, the null instance of which is the Laplace distribution. The tuning parameters at $\delta_1 = 1$ are $\zeta = 0.60$ for the Wilcoxon SSR S-R and CUSUM and $(\alpha, \beta, p) = (0.57, 1.00, 0.88)$ for the NPSR. At $\delta_1 = 0.5$ the parameters are $\zeta = 0.3$ and $(\alpha, \beta, p) = (0.79, 1.00, 0.75)$. The NPSR tuning parameters were obtained by exact calculation from equations (10) and (15) in Gordon and Pollak (1994). Figures 7 and 8 show the results when the schemes are tuned to the correct, i.e. Laplace, and incorrect, i.e. normal, underlying distributions respectively. In Figure 7 we see that the NPSR performs a little bit better than the S-R and a lot better than the CUSUM, especially at early changes. When the data come from the thinner tailed normal distribution then Figure 8 indicates, as expected, that the detection capability of each of the three schemes degenerates considerably. However, the S-R scheme seems to adapt itself better to the mistuning than the other two schemes.

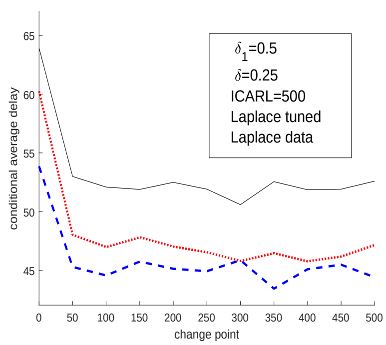

Figure 7

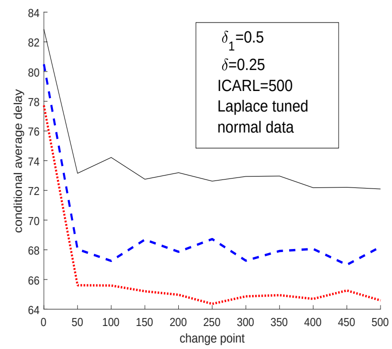

Figure 8

Figures 7 and 8 Conditional average delay simulation results. Tuned for Laplace distribution.

Finally, we compare some stationary average delay times (SADTs). The SADT is the average delay time measured from the time of the last false alarm, assuming that a stationary regime has established itself. Figures 1 and 2 already suggest strongly that stationarity sets in rather quickly and that the NPSR and SSR S-R schemes would perform similarly and exhibit smaller SADTs than the CUSUM. Here we take $ARL_0 = 500$, a target mean $\delta_1 = 0.5$ and change points $\tau = 500,\ 1,000$ and $1,500$ which ensures many restarts, thus a stationary situation, before the change takes place. For each of the schemes, the three SADT curves corresponding to the three change points are virtually indistinguishable, confirming what was anticipated earlier after looking at Figures 1 and 2, namely that a form of stationarity seems to become in force rather early. Therefore, only the results for $\tau = 1,000$ are plotted here.

These results are shown in Figure 9. The NPSR and S-R schemes are seen to behave very similarly and to have CADTs that are substantially smaller than those of the SSR CUSUM at shifts that are substantially less than the target. That the SADT performance of the S-R scheme is better than that of the CUSUM is a result that is in line with what is known about the behaviour of the corresponding parametric schemes - see Moustakides, Polunchenko and Tartakovsky (2009). Finally, Figure 10 shows that if the data actually come from a Laplace, rather than a normal, distribution then the NPSR again goes from "best" to "worst".

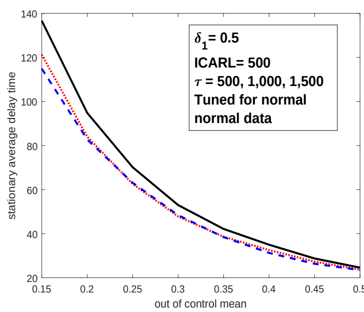

Figure 9

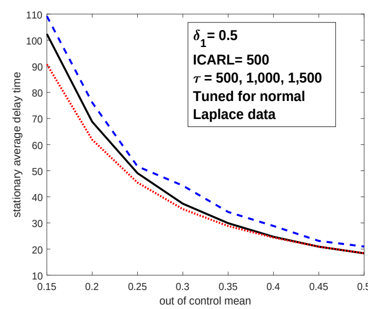

Figure 10

Figures 9 and 10 Stationary average delay time simulation results.

7 Application

The data and the particular application from which they arose are similar to those examined by Lombard and Van Zyl (2018) and consist of sequentially observed pairs of replicate coal ash values $(V_{1i}, V_{2i}),\ i \geq 1$. The measurements V_1 and V_2 come from two nominally identical coal samples analyzed by two independent laboratories. Denote the observations from the two laboratories by $V_k = T + \epsilon_k,\ k = 1,\ 2$ where T is the true ash content and the ϵ_k denote measurement error. These errors should be independent and identically distributed with zero means and common standard deviation σ. Then the difference $X = V_1 - V_2 = \epsilon_1 - \epsilon_2$ is independent of T and should be symmetrically distributed around zero. However, if the mean of X is nonzero then there exists a bias between the laboratory results which would call for an audit of their respective methodologies to identify the cause of the bias. Therefore, our interest is in monitoring the observed X_i sequence for sustained deviations away from a zero mean. The preceding is a typical matched pairs setup in which, were we dealing with a fixed sample of pairs, the Wilcoxon signed rank test would typically be used. Since the data are accruing one pair at a time, use of some sequential version of the test seems appropriate. We will compare the results produced by the Wilcoxon SSR S-R, SSR CUSUM and NPSR.

The three schemes considered thus far were designed to detect positive shifts. Denote by $S^{(k)}(X),\ k = 1, 2, 3$, any one of these schemes applied to the data $X = (X_1, X_2, \ldots)$. To also detect negative shifts, we run simultaneously the schemes $S^{(k)}(X)$ and $S^{(k)}(-X)$, that is, two schemes with the same ζ and h, one on the X data and the other on the sign-changed data $-X$. This is then the two-sided scheme with run length the smaller of the two constituent run lengths. The resulting ICARL is often close to one half that of each of the one-sided schemes. Thus, to find the control limits that produce a nominal ARL_0 in a two-sided scheme, the control limit applicable to a nominal $2 \times ARL_0$ in a one-sided scheme, adjusted after some Monte Carlo simulation, is used.

A practical limitation of the NPSR is that it is not computationally feasible to generate sufficiently accurate control limits guaranteeing ICARLs of $1,000$ or more at a wide range of reference values ζ. The data set shown in Figure 11, indicates *in retrospect* a change point relatively soon after initialization. Thus, we will implement the S-R, the CUSUM and the NPSR, using a two-sided ARL_0 of 400, making implementation of the NPSR feasible. In all three schemes the target change size is set at $\delta_1 = 0.25$ and each scheme is tuned to a t_4 distribution. For the S-R and CUSUM, this results in a reference value $\zeta = 0.15$ (rounded to two decimal places) and control limits $h = 6.52$ for the S-R and $h = 11.3$ for the CUSUM, found by interpolation from Tables 4 and 5. The tuning constants for the NPSR at $\delta_1 = 0.25$, namely $\alpha = 0.8912$, $\beta = 1.0774$ and $p = 0.6292$, were found by numerical computation from formulas (10) and (15) in Gordon and Pollak (1994). The approximation

$$\alpha \times ARL_0 = 0.8912 \times 800 = 712.96$$

to the ICARL of the (one-sided) NPSR (Gordon and Pollak, 1994, Theorem 2.2), supplemented with some Monte Carlo simulation, leads to a control limit $h = 725$ for an $ARL_0 = 400$ in a two-sided NPSR.

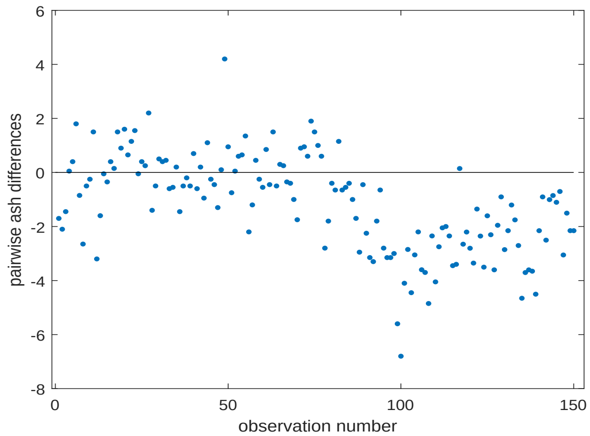

Figure 11 Plot of 150 successive pairwise ash differences.

Plots of the paths of the three schemes, applied to the data shown in Figure 11 are in Figures 12, 13 and 14. For visual presentation we plot $S^{(k)}(X)$ with control limit h (the upper path) and $-S^{(k)}(-X)$, $i \geq 1$ with control limit $-h$ (the lower path). The CUSUM scheme sounds an alarm at $i = 91$ while the S-R and NPSR both do so at $i = 93$. These conclusions seem to be in agreement with what is seen in Figure 11.

A distinguishing feature in a two-sided normal distribution CUSUM, also evident in Figure 12, is that the upper (lower) CUSUM is at zero whenever the lower (upper) CUSUM is non-zero. The usual CUSUM change point estimator is then the last index at which the hitting CUSUM sequence, upper or lower, was at zero. In the present instance, this estimator gives $\tau = 77$ as the change point. The same feature is not present in the S-R or the NPSR, so that an alternative estimator must be sought. A straightforward approach is to look upon the stopped sequence $X_1, X_2, \ldots, X_N$ as consisting of samples $\{X_1, \ldots, X_\tau\}$ and $\{X_{\tau+1}, \ldots, X_N\}$ from two distributions differing only in location and to estimate τ by least squares, conveniently ignoring the fact that the observed run length N is, in fact, a random variable. Then the least squares estimator of τ is

$$\widehat{\tau} = arg\ max_{1 \leq k \leq N-1}\ |T_k| \tag{13}$$

where

$$T_k = \sum_{i=k+1}^{N} X_i / \sqrt{N-k}. \tag{14}$$

This suggests using the estimator (13) after replacing X_i in (14) by ξ_i from (5). Denote this version of T_k by T_k^*. A plot of $|T_k^*|$ against k is shown in Figure 15. The maximum occurs at $\widehat{\tau} = 78$, which is almost the same estimate as that found from the CUSUM.

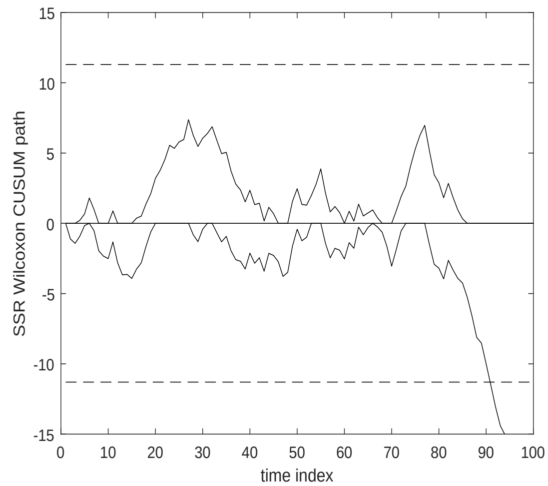

Figure 12 Wilcoxon SSR CUSUM paths.

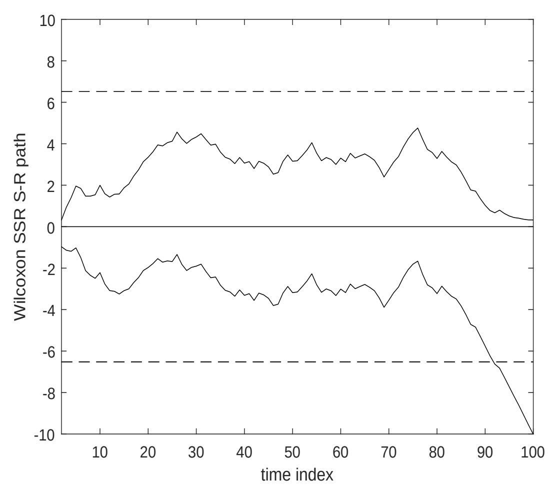

Figure 13 Wilcoxon SSR S-R paths.

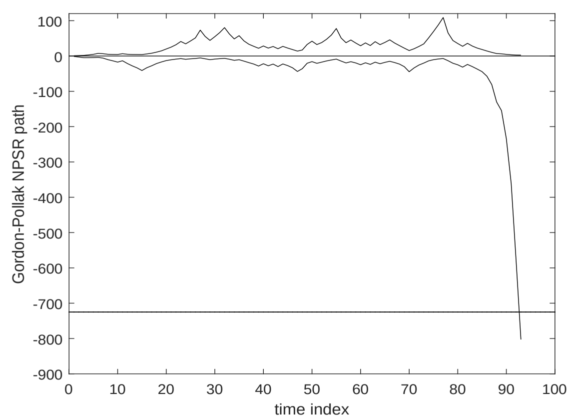

Figure 14 NPSR paths.

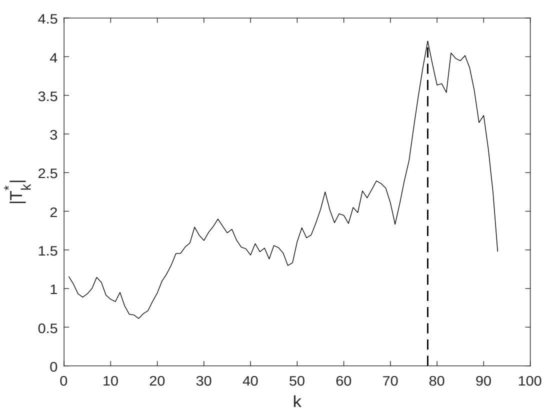

Figure 15 Least squares estimate of the change point.

8 Summary

We develop a Shiryaev-Roberts type scheme based on signed sequential ranks for detecting a change in the median from zero to a nonzero value in an unspecified symmetric distribution. The scheme is distribution free and scale invariant, meaning that a single set of control limits apply regardless of the functional form of the underlying distribution. Monte Carlo simulation results indicate that the scheme performs very well under a broad range of circumstances. In particular, it seems to be more adept at detecting small changes than a corresponding signed sequential rank CUSUM. Some

Monte Carlo simulations involving the signed sequential rank schemes and the distribution-free NPSR scheme developed by Gordon and Pollak (1994) suggest that the distribution-free Shiryaev-Roberts and NPSR schemes often show better performance in terms of out-of-control run length than the distribution-free CUSUM. The conceptual and computational simplicity of the distribution-free Shiryaev-Roberts scheme makes it an attractive alternative to both the distribution-free CUSUM and NPSR schemes. The focus in this paper has been on detecting a location change in a symmetric distribution. A matter for further research is the possibility of detecting scale changes via an appropriate construction of a Shiryaev- Roberts type sequential rank scheme. Furthermore, the possibility of constructing such schemes to deal with the detection of location and scale changes in asymmetric distributions, needs to be investigated.

References

[1] Bagshaw, M. and Johnson, R. A. (1974). The influence of reference value and estimated variance on the arl of cusum tests. Journal of the Royal Statistical Society - Series B, 37, 413-420.

[2] Barndorff-Nielsen, 0., (1963). On the limit behavior of extreme order statistics. The Annals of Mathematical Statistics, 34, 992-1002.

[3] Diko, M.D., Goedhart, R. and Does, R.J.M.M. (2020). A head-to-head comparison of the out-of-control performance of control charts adjusted for parameter estimation. Quality Engineering, 32 (4), 643-652

[4] Girshick, M. A. and Rubin, H. (1952). A Bayes approach to a quality control model, The Annals of Mathematical Statistics 23(1): 114–125.

[5] Gordon, L. and Pollak, M. (1994). An efficient sequential nonparametric scheme for detecting a change in distribution, The Annals of Statistics 22(2): 763–804.

[6] Hájek, J, Šidák, Z and Sen, P.K. (1999). Theory of Rank Tests. Second Edition. New York: Academic Press.

[7] Jones, L.A., Champ, C.W. and Rigdon, S.E. (2004). The Run Length Distribution of the CUSUM with Estimated Parameters. Journal of Quality Technology, 36, 95-108

[8] Lombard, F. (1981). An invariance principle for sequential nonparametric test statistics under contiguous alternatives. South African Statistical Journal, 15,129-152.

[9] Lombard, F. and Mason, D. M. (1985). Limit theorems for sequential rank statistics. Z. Wahrscheinlichkeitstheorie verw. Gebiete, 70, 395-410.

[10] Lombard, F. and Van Zyl, C. (2018). Signed sequential rank CUSUMs, Computational Statistics and Data Analysis 118(1): 30-39.

[11] Moustakides, G. V., Polunchenko, A. S. and Tartakovsky, A. G. (2009). Numerical comparison of CUSUM and Shiryaev-Roberts schemes for detecting changes in distribution, Communications in Statistics - Theory and Methods 38 (16-17): 3225-3239.

[12] Müller-Funk, U. (1983). Sequential signed rank statistics. Communications in Statistics-Sequential Analysis, 2 (2), 123:148

[13] Page, E. S. (1954). Continuous inspection schemes, Biometrika 41(1): 100-115.

[14] Pollak, M. and Siegmund, D. (1991). Sequential Detection of a Change in a Normal Mean when the Initial Value is Unknown. The Annals of Statistics, 19, 394-416.

[15] Reynolds, M.R., Jr. (1975). A Sequential Signed Rank Test for Symmetry. The Annals of Statistics, 3, 382-400.

[16] Roberts, S. W. (1966). A comparison of some control chart schemes. Technometrics 8, 411-430.

[17] Shiryaev, A. N. (1963). On optimum method in quickest detection problems, Theory of Probability and its Applications, 8, 22-46.

9 Appendix

Derivation of eqn. (9)

It is convenient to let $Y, Y_1, Y_2, \ldots$ be i.i.d. with a common symmetric around zero distribution with cdf F. Then, if a shift of size $\delta \neq 0$ occurs at index $\tau + 1$, the observations can be represented as $X_1 = Y_1, \ldots, X_\tau = Y_\tau$, $X_{\tau+1} = Y_{\tau+1} + \delta$. Furthermore, then

$$\begin{aligned}
&E\left[s(Y_{\tau+1}+\delta)\mathbf{1}\left(|Y_j| < |Y_{\tau+1}+\delta|\right)|Y_{\tau+1}=y\right] \\
&= s(y+\delta)\Pr\left(|Y| < |y+\delta|\right) \\
&= s(y+\delta)\left(2F(|y+\delta|)-1\right) \\
&= 2F(y+\delta)-1.
\end{aligned}$$

Consequently,

$$\begin{aligned}
E\left[\nu_{\tau+1}\xi_{\tau+1}\right] &= \frac{1}{\tau+2}E\left[s(X_{\tau+1})\textstyle\sum_{j=1}^{\tau}\mathbf{1}\left(|Y_j| \leq |Y_{\tau+1}+\delta|\right)\right] \\
&= \frac{\tau}{\tau+2}E\left[2F(Y+\delta)-1\right] + \frac{E\left[s(X_{\tau+1})\right]}{\tau+2} \\
&= E\left[2F(Y+\delta)-1\right] + 0\left(\frac{1}{\tau}\right) \qquad (15)
\end{aligned}$$

and by Taylor expansion

$$\begin{aligned}
E\left[2F(Y+\delta)-1\right] &= E\left[2F(Y)-1\right] + 2\delta E\left[f(Y)\right] + 0\left(\delta^2\right) \\
&= 2\delta\textstyle\int_{-\infty}^{+\infty} f^2(y)dy + 0\left(\delta^2\right). \qquad (16)
\end{aligned}$$

Also,

$$\nu_\tau = \sqrt{3} + 0\left(\frac{1}{\tau}\right). \qquad (17)$$

Putting (15), (16) and (17) together, we get

$$E\left[\xi_{\tau+1}\right] = \sqrt{12}\delta\textstyle\int_{-\infty}^{+\infty} f^2(y)dy + 0\left(\delta^2\right) + 0\left(\frac{1}{\tau}\right).$$