© 1990 American Statistical Association and
the American Society for Quality Control

Editor's Note: This article and the first two accompanying discussions were presented orally at the *Technometrics* session of the 33rd Annual Fall Technical Conference in Houston, Texas, October 25–27, 1989. The conference was cosponsored by the Chemical and Process Industries and the Statistics Divisions of the American Society for Quality Control and the Section on Physical and Engineering Sciences of the American Statistical Association.

Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements

James M. Lucas

Engineering Department
E. I. du Pont de Nemours and Company
Newark, DE 19714-6090

Michael S. Saccucci

Department of Quantitative Methods
Drexel University
Philadelphia, PA 19104

Roberts (1959) first introduced the exponentially weighted moving average (EWMA) control scheme. Using simulation to evaluate its properties, he showed that the EWMA is useful for detecting small shifts in the mean of a process. The recognition that an EWMA control scheme can be represented as a Markov chain allows its properties to be evaluated more easily and completely than has previously been done. In this article, we evaluate the properties of an EWMA control scheme used to monitor the mean of a normally distributed process that may experience shifts away from the target value. A design procedure for EWMA control schemes is given. Parameter values not commonly used in the literature are shown to be useful for detecting small shifts in a process. In addition, several enhancements to EWMA control schemes are considered. These include a fast initial response feature that makes the EWMA control scheme more sensitive to start-up problems, a combined Shewhart EWMA that provides protection against both large and small shifts in a process, and a robust EWMA that provides protection against occasional outliers in the data that might otherwise cause an out-of-control signal. An extensive comparison reveals that EWMA control schemes have average run length properties similar to those for cumulative sum control schemes.

KEY WORDS: Average run length; CUSUM; Fast initial response; Geometric moving average; Robust EWMA; Shewhart EWMA.

1. INTRODUCTION

Since Walter Shewhart introduced the control-chart technique in 1924, control schemes have found widespread application in improving the quality of manufacturing processes. Although the exponentially weighted moving average (EWMA) is known to have optimal properties in some forecasting and control applications (Box, Jenkins, and MacGregor 1974; Muth 1960), it has been largely neglected as a tool by quality-control analysts. Only recently has the EWMA control scheme been exploited and its properties evaluated analytically (Crowder 1987; Hunter 1986; Lucas and Saccucci 1987; Montgomery, Gardiner, and Pizzano 1987; Robinson and Ho 1978; Waldmann 1986).

Like Shewhart and cumulative sum (CUSUM) control schemes, an EWMA control scheme is easy to implement and interpret. It is based on the statistic

$$Z_i = \lambda Y_i + (1 - \lambda)Z_{i-1}, \qquad 0 < \lambda \leq 1, \quad (1.1)$$

together with upper control limits (UCL's) and lower control limits (LCL's). The starting value Z_0, which we shall discuss in more detail later, is often taken to be the target value. The sequentially recorded observations, Y_i, can be individually observed values from the process, although they are often sample averages obtained from a designated sampling plan. The process is considered out of control and action should be taken whenever Z_i falls outside the range of the control limits.

An EWMA has alternatively been referred to as a geometric moving average because Z_i can be equivalently written as a moving average of the current and past observations:

$$Z_i = \lambda \sum_{j=0}^{i-1} (1 - \lambda)^j Y_{i-j} + (1 - \lambda)^i Z_0, \quad (1.2)$$

where the weights of the past observations fall off exponentially as in a geometric series. In addition, many of the properties of EWMA's can be obtained from the formula for the sum of a geometric series.

When the Y_i are iid with common variance, σ_Y^2, the variance of the control statistic is given by

$$\sigma^2(Z_i) = [\{1 - (1 - \lambda)^{2i}\}\lambda/(2 - \lambda)]\sigma_Y^2.$$

Unless λ is small, the effect of the starting value soon dissipates and the variance quickly converges to its asymptotic value, $\sigma_Z^2 = \{\lambda/(2 - \lambda)\}\sigma_Y^2$. The control limits are usually based on the asymptotic standard deviation of the control statistic as LCL = Target $- L_1\sigma_Z$ and UCL = Target $+ L_2\sigma_Z$, respectively. In general, the control limits are chosen symmetrically about the process target value so that $L = L_1 = L_2$. For a discussion of one-sided EWMA control schemes, see Robinson and Ho (1978).

Roberts (1959) first described the use of EWMA control schemes. Using simulation, he developed nomograms of average run lengths (ARL's) for the case of normally distributed observations. In a subsequent article, Roberts (1966) compared their performance to other procedures, including CUSUM and Shewhart control schemes. More recently, Robinson and Ho (1978) numerically evaluated the ARL's of EWMA control schemes using an Edgeworth series expansion. Although they considered a wider range of parameter values than did Roberts, their results are inaccurate for small values of λ. Crowder (1987) evaluated the properties of EWMA's by formulating and solving a system of integral equations. Tables of the first and second moments of the run-length distribution are given in his article.

In this article, we evaluate the run-length properties of EWMA control schemes by representing the EWMA statistic as a continuous-state Markov chain. Its properties can be approximated by a finite-state Markov chain following a procedure similar to that of Brook and Evans (1972). This allows the properties of EWMA's to be evaluated more easily and completely than has previously been done (Lucas and Saccucci 1986; Yashchin 1987). A detailed discussion of the Markov-chain approach is given in Appendix A.

In Section 2, we give an example of an EWMA control scheme. In Section 3, we discuss the ARL tables that were obtained using the Markov-chain approach, and in Section 4, we give a design procedure based on these tables. Although ARL's can be evaluated for iid observations from any distribution, the ARL's provided here are evaluated for normally distributed observations. In Section 5, we evaluate and discuss the properties of several enhancements for EWMA's. These include a fast initial response (FIR) feature, a combined Shewhart EWMA, and a robust EWMA. Lucas and Crosier (1982a,b) and Lucas (1982) evaluated these enhancements for CUSUM control schemes. In Section 6, we compare EWMA and CUSUM control schemes. EWMA control schemes are shown to have ARL properties similar to those of CUSUM control schemes.

2. EXAMPLE

To illustrate an EWMA control scheme, we use a set of simulated observations taken from Lucas and Crosier (1982a). The data, together with the corresponding EWMA values, are shown in Table 1. The target value is taken to be 0, so the process is in control for the first 10 observations. The mean level was shifted upward by approximately one standard deviation for the last nine observations.

The parameters of the EWMA are chosen to be $\lambda = .25$ and $L = 3.0$, giving control limits of ± 1.134 $(\pm L[\lambda/(2 - \lambda)]^{1/2}\sigma_Y)$ when $\sigma_Y = 1$. For normally distributed observations, the in-control ARL is equal to 500.

The third column of Table 1 contains the values

Table 1. Example of an EWMA Control Scheme Using Data From a Process Initially in Control

			FIR EWMA[a] (50% HS)	
i	Observed value	EWMA[a]	Z_0^-	Z_0^+
0	—	.0	.567	.567
1	1.0	.250	−.175	.675
2	−.5	.063	−.256	.381
3	.0	.047	−.192	.286
4	−.8	−.165	−.344	.015
5	−.8	−.324	−.458	−.189
6	−1.2	−.543	−.644	−.442
7	1.5	−.032	−.108	.044
8	−.6	−.174	−.231	−.117
9	1.0	.119	.077	.162
10	−.9	−.135	−.167	−.103
11	1.2	.198	.175	.222
12	.5	.274	.256	.292
13	2.6	.855	.842	.869
14	.7	.817	.806	.827
15	1.1	.887	.880	.895
16	2.0	1.166[b]	1.160	1.171[b]
17	1.4	1.224[b]	1.220	1.228[b]
18	1.9	1.393[b]	1.390	1.396[b]
19	.8	1.245[b]	1.242	1.247[b]

[a] $\lambda = .25$; $L = 3.00$; CL = ±1.134.
[b] Out-of-control signal.

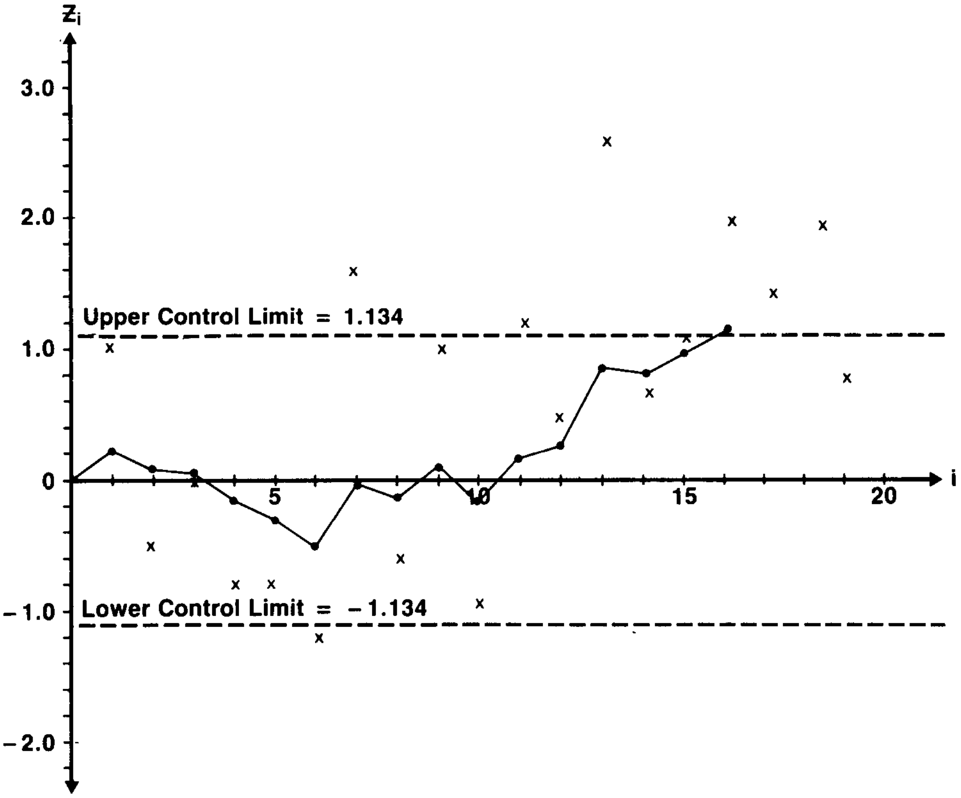

Figure 1. EWMA Control Scheme.

of the EWMA control statistic, $Z_i = .25Y_i + .75Z_{i-1}$. This statistic remains close to the target value for the first 10 observations but then grows in size after the shift in the process occurs. Like the CUSUM control scheme that was originally used with this data, the EWMA gives an out-of-control signal at the 16th observation. A Shewhart control scheme with the same in-control ARL has control limits of ± 3.09. The Shewhart control scheme does not give an out-of-control signal within the first 19 observations. In fact, for a one-standard deviation shift in the process mean, the EWMA control scheme will give an out-of-control signal in an average of 10.9 observations, whereas the Shewhart control scheme will give an out-of-control signal in an average of 54.6 observations.

Table 2. Example of an EWMA Control Scheme Using Data From a Process Initially Out of Control

i	*Observed value*	*EWMA*[a]	*FIR EWMA*[a] *(50% HS)* Z_0^-	Z_0^+
0	—	.0	−.567	.567
1	1.2	.300	−.125	.725
2	.5	.350	.031	.669
3	2.6	.913	.673	1.152[b]
4	.7	.859	.680	1.039[c]
5	1.1	.920	.785	1.054[c]
6	2.0	1.190[b]	1.089	1.291[b]
7	1.4	1.242[b]	1.167	1.318[b]
8	1.9	1.407[b]	1.350	1.463[b]
9	.8	1.255[b]	1.212	1.298[b]

[a] $\lambda = .25$; $L = 3.00$; CL = ± 1.134.
[b] Out-of-control signal.
[c] Observations after an out-of-control signal can drop the EWMA statistic below the control limits.

In quality-control applications, it is often useful to graphically display control schemes. Figure 1 is a plot of the EWMA control statistic, together with the original data given in Table 1. The solid line connects the EWMA values, and the individual observations are represented by X's. Roberts (1959) provided a manual plotting procedure that does not require evaluation of the EWMA values. Hunter (1986) suggested writing the current EWMA as the previous EWMA plus a fraction of the difference between the current observation and the previous EWMA, $Z_i = Z_{i-1} + \lambda(Y_i - Z_{i-1})$. In this form, the EWMA can be considered a one-step-ahead forecast for the process so that Z_i should be plotted at the $(i + 1)$st sequential position. Table 2 illustrates the use of an EWMA control scheme for a process with a mean level that is initially off target. It will be discussed in more detail in Section 4.

3. THE RUN-LENGTH DISTRIBUTION

Figure 2 illustrates the run-length distribution for an EWMA control scheme with $L = 2.414$ and $\lambda =$

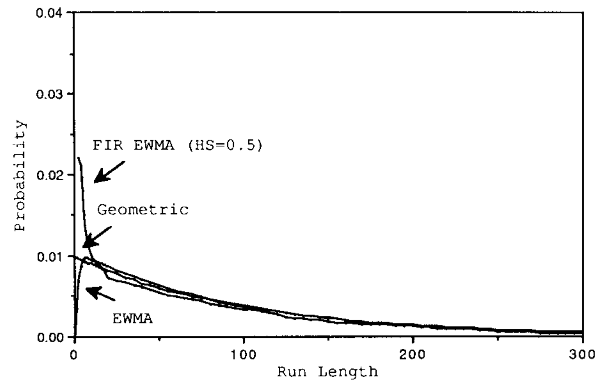

Figure 2. Run-Length Distribution for an In-Control EWMA.

.25. This control scheme has an in-control zero-state ARL equal to 100. Zero-state run lengths refer to the run lengths of control schemes initialized at the target value. Figure 2 also illustrates the run-length distribution for a geometric random variable with an ARL equal to 100 and an EWMA control scheme with a FIR feature. The FIR feature is discussed in detail in Section 5.

The run-length distribution for the EWMA control scheme was obtained using the Markov-chain approach given in Appendix A. The run-length distribution for the FIR EWMA was obtained by simulation. Examination of Figure 2 shows that, except for small run lengths, the in-control run-length distribution is nearly geometric. Hence, for an in-control process, the run-length distribution of an EWMA can be adequately characterized by its ARL.

Figure 3 illustrates the run-length distributions for a process that is off aim by two σ_Y. The EWMA's in Figure 3 have the same parameter values [L, λ, and HS (head start)] as those in Figure 2, and the geometric distribution was chosen to match the ARL of the zero-state EWMA. Although the run-length distributions for the EWMA's are clearly not geometric, they are closely centered about their average values of 2.81 and 1.93, respectively. Hence the ARL also

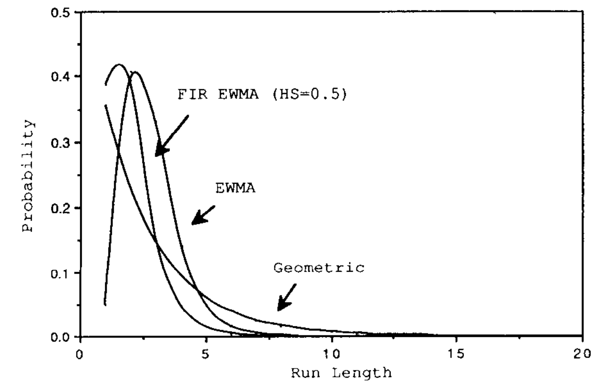

Figure 3. Run-Length Distribution for an EWMA When the Process Has Shifted Two Standard Deviations.

characterizes the run-length distribution for an out-of-control process.

Using the Markov-chain approach described in Appendix A, we have calculated zero-state and cyclical steady-state ARL's for a wide range of parameters. Steady-state run lengths refer to the run lengths of control schemes evaluated after the control statistic has reached steady state. A control statistic is in steady state if the process has been in control long enough for the effect of the starting value to become negligible. See Appendix A for a detailed discussion of steady-state run lengths.

With the exception of those of Ho (1978), ARL tables are usually given for fixed values of λ and L. We feel that it is convenient to have tables with fixed values of λ and in-control ARL's. Table 3 gives ARL values for EWMA control schemes with λ equal to 1.00, .75, .50, .40, .30, .25, .20, .10, .05, and .03. More detail is given for values of λ less than .5 because small values of λ are useful for detecting small shifts in a process. The corresponding values of L were obtained so that the in-control zero-state ARL's are equal to 500. The same values of λ and L were used for both zero-state and steady-state ARL's so that a direct comparison is possible.

Lucas and Saccucci (1987) also provided tables for in-control ARL's of 100, 300, 1,000, 2,000, and 5,000. The complete set of tables indicates that the differences between zero-state and steady-state ARL's depend on both λ and the in-control ARL. For large values of λ, there is essentially no difference in ARL's. For small values of λ, however, the percentage of difference increases as the in-control ARL decreases. Within the range of parameters considered, the difference between zero-state and steady-state ARL's ranges from less than 2% for an EWMA control scheme with an in-control ARL of 5,000 to approximately 10% for an EWMA control scheme with an in-control ARL of 100. For most practical purposes, the difference between zero-state and steady-state ARL's is unimportant and either one suffices. For subtle comparisons, such as those between EWMA and CUSUM control schemes given later, however, it is helpful to have both sets of ARL values.

4. DESIGN PROCEDURE

In addition to characterizing the run-length distribution of an EWMA control scheme, the ARL is also proportional to the amount of production from a process. Hence design procedures are usually based on the ARL properties of control schemes. The ARL should be long when the process is operating near its target value and short when the process shifts to an unacceptable level. Depending on the particular

Table 3. Average Run Lengths for an EWMA Control Scheme

Shift	Type	L =	3.090	3.087	3.071	3.054	3.023	2.998	2.962	2.814	2.615	2.437
		λ =	1.00	.75	.50	.40	.30	.25	.20	.10	.05	.03
.00	Zero state		500	500	500	500	500	500	500	500	500	500
	Steady state		500	500	499	498	497	496	496	492	487	480
.25	Zero state		374	321	255	224	189	170	150	106	84.1	76.7
	Steady state		374	321	254	223	188	169	149	104	81.7	74.1
.50	Zero state		201	140	88.8	71.2	55.4	48.2	41.8	31.3	28.8	29.3
	Steady state		201	140	88.4	70.7	54.9	47.7	41.2	30.6	28.0	28.6
.75	Zero state		103	62.5	35.9	28.4	22.5	20.1	18.2	15.9	16.4	17.6
	Steady state		103	62.4	35.7	28.1	22.2	19.8	17.8	15.5	16.0	17.3
1.00	Zero state		54.6	30.6	17.5	14.3	12.0	11.1	10.5	10.3	11.4	12.6
	Steady state		54.6	30.5	17.3	14.1	11.8	10.9	10.3	10.1	11.2	12.5
1.50	Zero state		17.9	9.90	6.53	5.88	5.53	5.46	5.50	6.09	7.12	8.08
	Steady state		17.9	9.86	6.44	5.79	5.43	5.37	5.40	5.99	7.03	8.00
2.00	Zero state		7.26	4.54	3.63	3.52	3.54	3.61	3.74	4.36	5.23	5.99
	Steady state		7.26	4.52	3.58	3.47	3.49	3.56	3.69	4.31	5.18	5.95
2.50	Zero state		3.60	2.69	2.50	2.54	2.65	2.74	2.88	3.44	4.17	4.80
	Steady state		3.60	2.67	2.47	2.50	2.61	2.71	2.84	3.41	4.14	4.78
3.00	Zero state		2.15	1.88	1.93	2.02	2.16	2.26	2.38	2.87	3.50	4.03
	Steady state		2.15	1.87	1.91	1.99	2.12	2.22	2.35	2.85	3.48	4.02
3.50	Zero state		1.52	1.46	1.58	1.69	1.85	1.95	2.07	2.47	3.04	3.49
	Steady state		1.52	1.46	1.58	1.68	1.82	1.91	2.03	2.47	3.02	3.49
4.00	Zero state		1.22	1.22	1.34	1.44	1.61	1.73	1.86	2.19	2.69	3.11
	Steady state		1.22	1.23	1.36	1.46	1.60	1.69	1.80	2.20	2.68	3.09
5.00	Zero state		1.03	1.04	1.07	1.12	1.22	1.32	1.48	1.94	2.16	2.55
	Steady state		1.03	1.04	1.10	1.17	1.29	1.38	1.49	1.83	2.22	2.55

NOTE: L values are based on zero-state in-control ARL = 500.

needs of the analyst, control schemes can be evaluated using zero-state and/or steady-state ARL's.

To facilitate the design of EWMA control schemes, Table 4 contains a list of "optimal" parameters. For a specified in-control ARL and shift in the process, these parameters will give an EWMA control scheme having the minimum ARL at the specified shift. Whenever the minimum ARL's were nearly equivalent over a range of λ values, the entire range of λ was given. An EWMA control scheme designed with the smallest value of λ in this range will provide more protection against small shifts in the process, but an EWMA control scheme designed with the largest value of λ will provide more protection against large shifts in the process.

Examination of Table 4 illustrates that, for a fixed in-control ARL, the optimal value of λ increases as the shift in the process increases. Lucas (1973) proved that a Shewhart control scheme, an EWMA with λ equal to 1, is optimal for detecting large shifts. Consequently, if $\sigma_Y^2 = \sigma^2/n$ and there is no restriction on choosing n, there is little need to consider control schemes other than standard Shewhart control schemes. Shewhart control schemes can be optimized for detecting a given size shift by increasing the number of samples included in each observation. On the other hand, it is often appropriate to consider more complicated error structures. For many chemical processes, the error structure can be approximated as $\sigma_Y^2 = \sigma_{\text{between}}^2 + \sigma_{\text{within}}^2/n$. In this situation, only short-term variation is reduced when multiple samples are taken, and the EWMA control scheme is useful for detecting small shifts in the process. For a further discussion of this point, see Goel (1968) or Lucas (1976).

To design an EWMA control scheme, we recommend the following procedure. First, specify the desired in-control ARL and the shift in the process that is to be detected quickly. Using Table 4, obtain the EWMA parameters that will result in the minimum ARL for the specified shift in the process. Finally, the entire ARL profile for this EWMA should be evaluated to determine whether it provides sufficient protection against other shifts. The ARL profile can be obtained from the ARL tables given by Lucas and Saccucci (1987).

Note that the values of λ recommended for detecting one-σ_Y shifts in the process are small relative to the values between .25 and .50 that have often been suggested in the literature. Table 4 indicates

Table 4. Optimal EWMA Control Schemes

	In-control average run length					
Shift	*100*	*300*	*500*	*1,000*	*2,000*	*5,000*
.5						
λ	.07–.06	.06–.05	.05	.04	.04–.03	.03
L	2.015–1.954	2.462–2.399	2.616	2.817	3.069–2.989	3.299
ARL_{min}	17.3	24.9	28.7	34.3	40.1	47.7
1.0						
λ	.19–.16	.15–.14	.15–.12	.13–.10	.12–.10	.09
L	2.346–2.298	2.723–2.707	2.907–2.858	3.113–3.059	3.317–3.283	3.538
ARL_{min}	6.97	9.14	10.2	11.7	13.2	15.2
2.0						
λ	.52–.47	.42–.38	.37–.36	.35–.31	.32–.28	.29–.26
L	2.538–2.526	2.895–2.885	3.047–3.044	3.253–3.241	3.445–3.433	3.686–3.677
ARL_{min}	2.62	3.23	3.51	3.90	4.29	4.81
3.0						
λ	.81–.77	.74–.71	.70–.66	.66–.59	.61–.53	.53–.47
L	2.572–2.569	2.931–2.930	3.086–3.084	3.286–3.283	3.477–3.473	3.714–3.711
ARL_{min}	1.45	1.72	1.86	2.06	2.26	2.51
4.0						
λ	1.00–.85	.97–.84	.95–.82	.91–.80	.91–.75	.84–.72
L	2.576–2.573	2.935–2.934	3.090–3.089	3.290–3.289	3.480–3.480	3.719–3.718
ARL_{min}	1.08	1.16	1.21	1.29	1.39	1.53

NOTE: The in-control average run length is based on zero-state average run lengths.

that this range is optimal for detecting two-σ_Y shifts. Furthermore, our design suggestions differ from those of Hunter (1986). He suggested choosing λ to minimize the one-step-ahead forecast error by using the past history of the data. Although past history provides guidance, we feel that it is important to design the EWMA to guard against possible future shifts in the process mean.

5. ENHANCEMENTS

5.1 FIR Feature

Lucas and Crosier (1982a) showed that a FIR feature is useful for CUSUM control schemes because processes are more likely to be away from the target value when a control scheme is initiated due to start-up problems or because of ineffective control action after the previous out-of-control signal. A FIR feature is especially useful for EWMA control schemes designed with small values of λ. When λ is small, the variance of the control statistic converges slowly to its asymptotic value so that control schemes based on the asymptotic standard deviation tend to be insensitive at start-up.

A FIR feature for an EWMA control scheme can be obtained by simultaneously implementing two one-sided EWMA's, each with a head start (HS). One EWMA has an HS, or starting value, above the target value and the other EWMA has an HS below the target value. If the process is off aim at start-up, the EWMA with the appropriate HS will tend to give an out-of-control signal more quickly. On the other hand, if the process is initially in control, the two EWMA's will tend to converge. In practice, one of the control schemes can be discontinued when they are sufficiently close—for example, whenever they differ by less than .1 σ_Z.

Table 1 gives an example of the FIR feature for a process that is initially in control. The two simultaneous EWMA control statistics are represented by Z^+ and Z^-. For purposes of illustration, the starting values are taken halfway between the process target value and the control limits (i.e., 50% HS). Since the process is initially in control, both EWMA's remain well within the control limits as they converge toward each other. By the 12th observation, both schemes are within .1 σ_Z and Z^- could be discontinued. Although the HS value never completely disappears, the performances of the EWMA and the FIR EWMA are similar after the first seven observations.

To illustrate the advantage of the FIR feature when the process is initially out of control, Table 2 contains the last nine observations from Table 1. Although both EWMA's immediately start to increase, the FIR EWMA with a 50% HS gives an out-of-control signal at the 3rd observation, whereas the EWMA does not signal until the 6th observation.

Appendix B describes the method used to evaluate the properties of the FIR feature. Table 5 gives FIR

Table 5. FIR Average Run Lengths for an EWMA Control Scheme

Shift	%HS		3.090	3.087	3.071	3.054	3.023	2.998	2.962	2.814	2.615	2.437
		L =	3.090	3.087	3.071	3.054	3.023	2.998	2.962	2.814	2.615	2.437
Shift	%HS	λ =	1.00	.75	.50	.40	.30	.25	.20	.10	.05	.03
.00	0		500	500	500	500	500	500	500	500	500	500
	25		500	498	497	497	495	491	491	487	470	465
	50		500	496	487	487	485	483	475	468	434	406
	75		500	495	478	471	456	444	429	382	312	258
.50	0		201	140	88.8	71.2	55.4	48.2	41.8	31.3	28.8	29.3
	25		201	140	87.8	70.0	53.9	46.5	39.7	28.3	24.7	24.3
	50		201	139	86.1	67.8	51.2	43.6	36.6	24.2	19.5	18.4
	75		201	138	82.7	63.5	46.2	38.2	30.8	17.9	12.9	11.4
1.00	0		54.6	30.6	17.5	14.3	12.0	11.1	10.5	10.3	11.4	12.6
	25		54.6	30.2	16.9	13.5	11.1	10.1	9.40	8.75	9.30	10.1
	50		54.6	29.7	15.9	12.4	9.82	8.79	7.93	6.87	6.93	7.36
	75		54.6	29.1	14.5	10.8	7.99	6.86	5.91	4.56	4.28	4.37
2.00	0		7.26	4.54	3.63	3.52	3.54	3.61	3.74	4.36	5.23	5.99
	25		7.26	4.33	3.29	3.13	3.08	3.11	3.17	3.57	4.19	4.74
	50		7.26	4.09	2.87	2.64	2.52	2.50	2.51	2.72	3.08	3.43
	75		7.26	3.81	2.41	2.11	1.90	1.82	1.76	1.76	1.90	2.07
3.00	0		2.15	1.88	1.93	2.02	2.16	2.26	2.38	2.87	3.50	4.03
	25		2.15	1.75	1.69	1.73	1.83	1.90	2.01	2.35	2.80	3.20
	50		2.15	1.63	1.45	1.44	1.46	1.49	1.54	1.80	2.11	2.34
	75		2.15	1.51	1.26	1.21	1.18	1.17	1.16	1.19	1.29	1.41
5.00	0		1.03	1.04	1.07	1.12	1.22	1.32	1.48	1.94	2.16	2.55
	25		1.03	1.02	1.03	1.04	1.07	1.09	1.15	1.50	1.93	2.05
	50		1.03	1.01	1.01	1.01	1.01	1.02	1.02	1.07	1.27	1.57
	75		1.03	1.01	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.01

NOTE: L values are based on zero-state in-control ARL = 500.

ARL's for EWMA control schemes with an in-control ARL equal to 500 and HS's of 0%, 25%, 50%, and 75%. ARL's for other HS values can be obtained by interpolation.

In general, we find that this FIR feature is most useful for EWMA control schemes with λ less than or equal to .25. Unlike CUSUM control schemes, there is no theory that suggests the appropriate HS value, so the HS value is obtained by examining the trade-off between the percentage of decrease in the in-control ARL and the percentage of decrease in the out-of-control ARL's. An HS value of approximately 50% works well in the situations that we have examined.

A disadvantage of this approach is that it requires two separate EWMA's for each process being monitored. This can be cumbersome when the number of processes being monitored is large. Moreover, this procedure is not easily generalized to multivariate EWMA's. An alternative FIR feature can be achieved by using tighter control limits for the initial few observations. A detailed discussion of this approach will be given in a future article.

5.2 Combined Shewhart EWMA

Although EWMA control schemes can be designed to quickly detect small shifts in level, Lucas (1973) showed that Shewhart control schemes are superior for detecting large shifts. A combined Shewhart EWMA often gives improved properties when both large and small shifts are to be detected. This is achieved by adding Shewhart limits to an EWMA control scheme so that an out-of-control signal is given if the EWMA statistic is outside the control limits or if the current observation is outside the Shewhart limits.

Tables of ARL's for the combined Shewhart EWMA were given by Lucas and Saccucci (1987). Examination of these tables indicates that the combined Shewhart EWMA shows improved ARL values similar to those obtained by a combined Shewhart CUSUM (Lucas 1982). In general, we suggest choosing Shewhart limits larger than would be used for a standard Shewhart control scheme to prevent the Shewhart limits from causing a large reduction in the in-control ARL. Although the recommended Shewhart control limits depend on the in-control ARL and the value of λ, Shewhart control limits between 4.0 and 4.5 σ_Y work well for an EWMA with an in-control ARL equal to 500.

5.3 Robust EWMA

It is useful to consider a robust EWMA whenever the process produces occasional outliers. The combined Shewhart EWMA should not be used in these situations because the addition of Shewhart limits causes the combined Shewhart EWMA to be sensitive to the occurrence of outliers.

In this article, we examined the two-in-a-row rule, which worked well for CUSUM control schemes (Lucas and Crosier 1982b). This procedure requires the specification of outlier limits for the observations entering the EWMA (i.e., Y_i). In practice, these limits are often set at $\pm 4\,\sigma_Y$ about the target value. A single observation outside of the outlier limits does not enter the EWMA, but two outliers in a row are considered to be an out-of-control signal. More complicated rules are generally not as useful; they decrease the sensitivity of the scheme when large shifts occur.

The performance of the two-in-a-row outlier rule was evaluated using a contaminated normal distribution, formed by combining two normal random variables with common mean. The standard deviation of the first normal random variable was chosen to be $(1 + 8\alpha)^{-1/2}\,\sigma_Y$, where α is the proportion of contamination of the second normal random variable. The standard deviation of the second normal random variable was taken to be three times that of the first normal random variable. Although the overall standard deviation of the contaminated distribution remains constant, the tails of the distribution become heavier as the percentage of contamination increases.

The ARL properties of the robust EWMA were examined for distributions with 0%, 1%, 3%, and 10% levels of contamination. This is identical to the approach used by Lucas and Crosier (1982b) in evaluating robust CUSUM control schemes, so a comparison can be made between the two schemes.

Tables of ARL's for the two-in-a-row robust EWMA were given by Lucas and Saccucci (1987). Examination of the robust ARL tables indicates that the ARL's for the robust EWMA are always larger than the ARL's for the EWMA at the same level of contamination. Furthermore, the effect of outliers depends on both the level of contamination and the parameters of the EWMA control scheme. For large values of λ, the ARL's substantially decrease as the level of contamination increases, indicating that Shewhart control schemes are extremely sensitive to outliers. For small values of λ, the ARL's tend to increase as the level of contamination increases.

We consider robust control procedures applicable for situations in which outliers are known to exist, such as those involving tough analytical procedures. We do not recommend the routine implementation of robust control procedures, however. The presence of outliers is itself an indication of quality problems, and work should be done to reduce these problems.

6. COMPARISONS BETWEEN EWMA AND CUSUM CONTROL SCHEMES

Lucas and Saccucci (1987) compared the ARL's of EWMA and CUSUM control schemes over a wide range of parameter values. These comparisons indicate that there is little practical difference between the ARL properties of the two control schemes. Although one scheme may be superior in terms of other properties, it seems likely that nonstatistical criteria could be used to decide which particular procedure should be used in a given situation. If one scheme is currently being used in most applications, we would recommend using the same scheme throughout.

Table 6 provides a specific comparison between an EWMA and a CUSUM control scheme. The parameters of the EWMA's were chosen so that the in-

Table 6. ARL Comparisons Between EWMA and CUSUM Control Schemes

	Zero state			Steady state			FIR[a]			Worst case	
Shift	$EWMA_1$[b]	$EWMA_2$[c]	CUSUM	$EWMA_1$	$EWMA_2$	CUSUM	$EWMA_1$	$EWMA_2$	CUSUM	$EWMA_1$	$EWMA_2$
.00	465	465	465	459	459	459	434	435	430	310	315
.25	116	118	139	114	116	137	104	106	122	97.6	100
.50	33.3	33.8	38.0	32.6	33.1	36.4	27.0	27.6	28.7	34.2	34.7
.75	16.0	16.1	17.0	15.6	15.7	16.0	11.8	12.0	11.2	19.1	19.1
1.00	10.1	10.0	10.4	9.84	9.84	9.62	6.99	7.04	6.35	13.3	13.2
1.50	5.71	5.67	5.75	5.62	5.57	5.28	3.74	3.73	3.37	8.43	8.32
2.00	4.04	3.99	4.01	3.98	3.94	3.68	2.59	2.57	2.36	6.25	6.16
2.50	3.16	3.12	3.11	3.13	3.09	2.86	2.01	2.00	1.86	5.01	4.93
3.00	2.62	2.59	2.57	2.61	2.57	2.38	1.66	1.65	1.54	4.21	4.14
4.00	2.05	2.03	2.01	2.01	1.98	1.86	1.23	1.22	1.16	3.23	3.19
5.00	1.77	1.74	1.69	1.68	1.66	1.53	1.04	1.04	1.02	2.71	2.66

NOTE: EWMA is compared to a CUSUM with $h = 5.00$ and $k = .50$.
[a]FIR feature with a 50% HS.
[b]$EWMA_1$ with $L = 2.856$ and $\lambda = .133$; in-control zero-state ARL equated to CUSUM.
[c]$EWMA_2$ with $L = 2.866$ and $\lambda = .139$; in-control steady-state ARL equated to CUSUM.

control ARL's would match those of a CUSUM control scheme with $h = 5.0$ and $k = .5$. The matching was done for both zero-state and steady-state ARL's, so two EWMA's were evaluated. Both EWMA's ended up with similar properties. The first EWMA was obtained using parameters $\lambda = .133$ and $L = 2.856$, and the second EWMA was obtained using parameters $\lambda = .139$ and $L = 2.864$. Each of these schemes was designed to detect a 1-σ_Y shift in the process.

Our comparisons showed that the ARL's for the EWMA are usually smaller than the ARL's of the CUSUM up to a value of the shift near the one that the scheme was designed to detect. Beyond this shift, the ARL's of the EWMA are larger than the ARL's of the corresponding CUSUM. The steady-state comparisons favor the CUSUM more strongly than the zero-state comparisons because random observations on one side of the target value can delay detection of a shift to the other side for an EWMA.

Following a referee's suggestion, we evaluated worst-case ARL's for EWMA control schemes. In this situation, the EWMA control statistic is at or near one of its control limits when a shift occurs in the opposite direction. An example of worst-case ARL's is shown in Table 6. For small values of λ and moderate to large shifts, the worst-case ARL's are slightly larger than the steady-state ARL's. In this worst-case situation, the EWMA requires a few observations to overcome its initial inertia. Although this situation is highly unlikely, it can be guarded against by using the combined Shewhart EWMA.

7. SUMMARY AND CONCLUSIONS

We have described the properties of EWMA control schemes and have compared them with CUSUM control schemes. The results show that the properties of EWMA's are very close to those of CUSUM schemes.

Both schemes include the one-parameter Shewhart control scheme as a special case. The two parameters in the EWMA and CUSUM control schemes are used to average observations over time. This makes them less sensitive to outliers and enables them to detect small shifts more quickly than the standard Shewhart control scheme.

Several enhancements to EWMA control schemes were evaluated. These include a FIR feature that makes the scheme more sensitive at start-up, a combined Shewhart EWMA that provides protection against both large and small shifts in the process, and a robust EWMA that provides extra protection against outliers. These enhancements work as well for EWMA control schemes as they do for CUSUM control schemes.

ACKNOWLEDGMENTS

We thank the editor, the associate editor, and the referees for their helpful suggestions. We also express our appreciation to Ron Crosier for many helpful comments and for providing the CUSUM steady-state ARL values and to the computer-center staff at Drexel University for providing valuable assistance in the preparation of this manuscript. This research was partially supported by the Office of Sponsored Projects at Drexel University.

APPENDIX A: MARKOV-CHAIN APPROACH

The properties of an EWMA control scheme can be approximated using a procedure similar to that described by Brook and Evans (1972). Although they suggested discretizing the control statistic and then evaluating the exact properties of the discretized statistic, we evaluate the properties of the continuous-state Markov chain by discretizing the infinite-state transition probability matrix. This procedure involves dividing the interval between the upper and lower control limits into $t = 2m + 1$ subintervals of width 2δ. The control statistic, Z_i, is said to be in transient state (j) at time (i) if $S_j - \delta < Z_i \le S_j + \delta$ for $j = -m, -m + 1, \ldots, m$, where S_j represents the midpoint of the jth interval. The control statistic is in the absorbing state (a) if Z_i falls outside the control limits. The process is assumed to be in control whenever Z_i is in a transient state and is assumed to be out of control whenever Z_i is in the absorbing state. For obvious reasons, the transient states are often referred to as in-control states and the absorbing state is often referred to as the out-of-control state.

The run-length distribution of an EWMA is completely determined by its initial probability vector and transition probability matrix. The initial probability vector can be represented by

$$\mathbf{p}_{\text{int}}^T = (p_{-m}, \ldots, p_{-1}, p_0, p_1, \ldots, p_m \mid 0)$$
$$= (\mathbf{p}^T \mid 0),$$

where p_j represents the probability that Z starts in state (j). Note that p_a is equal to 0 because the control statistic is assumed to start in control. In practice, the initial probability vector will usually either contain a single element equal to 1, representing the initial starting state, or it will be a vector of steady-state probabilities.

The transition probability matrix, represented in partitioned matrix form, is given by

$$\mathbf{P} = \begin{pmatrix} \mathbf{R} & (\mathbf{I} - \mathbf{R})\mathbf{1} \\ \mathbf{0}^T & 1 \end{pmatrix},$$

where the submatrix $\mathbf{R}$ contains the probabilities of going from one transient state to another, $\mathbf{I}$ is the identity matrix, and $\mathbf{1}$ is a column vector of ones. Hence p_{jk} represents the probability that the control statistic goes from state (j) to state (k) in one step. To approximate this probability, we assume that the control statistic is equal to S_j whenever it is in state (j). This yields

$$\begin{aligned} p_{jk} &= \Pr(\text{going to } S_k \mid \text{in } S_j) \\ &\approx \Pr[\lambda^{-1}\{(S_k - \delta) - (1 - \lambda)S_j\} \\ &\qquad < Y_i \le \lambda^{-1}\{(S_k + \delta) - (1 - \lambda)S_j\}], \\ &\qquad\qquad j = -m, -m + 1, \ldots, m. \end{aligned}$$

The rows of the transition probability matrix must add to 1 so that the probabilities of going from an in-control state to the out-of-control state are found by subtraction.

For the special case of iid normal observations with mean μ_Y and standard deviation σ_Y, the in-control transition probabilities are given by

$$\begin{aligned} p_{jk} &= \Phi[(\lambda\sigma_Y)^{-1}\{(S_k + \delta) - (1 - \lambda)S_j - \lambda\mu_Y\}] \\ &\quad - \Phi[(\lambda\sigma_Y)^{-1}\{(S_k - \delta) - (1 - \lambda)S_j - \lambda\mu_Y\}], \end{aligned}$$

where Φ represents the standard normal distribution function.

Crosier (1986) suggested methods for calculating both conditional and cyclical steady-state probability vectors. Note, however, that an exact steady-state probability vector does not exist because the transition probability matrix is not ergodic. The steady-state probability vector that we feel best models the way control schemes are used is a cyclical steady-state probability vector that is obtained by altering the transition probability matrix so that the control statistic is reset to state (0) whenever it goes into the out-of-control state; that is,

$$\mathbf{P}^* = \begin{pmatrix} \mathbf{R} & (\mathbf{I} - \mathbf{R})\mathbf{1} \\ 0 \cdots 1 \cdots 0 & 0 \end{pmatrix}.$$

This transition probability matrix is ergodic. The steady-state probability vector, $\mathbf{p}_{ss}$, is found by solving $\mathbf{p} = \mathbf{P}^{*T}\mathbf{p}$ subject to $\mathbf{1}^T\mathbf{p} = 1$. Then, $\mathbf{p}_{ss} = (\mathbf{1}^T\mathbf{q})^{-1}\mathbf{q}$, where $\mathbf{q}$ is a vector of length t obtained from $\mathbf{p}$ by deleting the entry corresponding to the absorbing state; that is, $\mathbf{p}_{ss}$ is the probability vector obtained from $\mathbf{p}$ by deleting the entry corresponding to absorbing state and normalizing so that the probabilities sum to 1.

The ith-stage transition probability matrix is useful for evaluating the run-length distribution because it contains the probabilities that the control statistic goes from one state to another state in i steps,

$$\mathbf{P}^i = \begin{pmatrix} \mathbf{R}^i & (\mathbf{I} - \mathbf{R}^i)\mathbf{1} \\ \mathbf{0}^T & 1 \end{pmatrix}.$$

Hence

$$\Pr(\text{RL} \le i) = \mathbf{p}^T(\mathbf{I} - \mathbf{R}^i)\mathbf{1}$$

and

$$\Pr(\text{RL} = i) = \mathbf{p}^T(\mathbf{R}^{i-1} - \mathbf{R}^i)\mathbf{1}. \qquad \text{(A.1)}$$

Using Equation (A.1), the ARL based on t in-control states is given by

$$\begin{aligned} \text{ARL}(t) &= \sum_{i=1}^{\infty} i \Pr(\text{RL} = i) \\ &= \sum_{i=1}^{\infty} i\, \mathbf{p}^T(\mathbf{R}^{i-1} - \mathbf{R}^i)\mathbf{1} \\ &= \sum_{i=1}^{\infty} \mathbf{p}^T\mathbf{R}^{i-1}\mathbf{1} = \mathbf{p}^T(\mathbf{I} - \mathbf{R})^{-1}\mathbf{1}. \qquad \text{(A.2)} \end{aligned}$$

In general, higher-order moments of the run-length distribution can easily be obtained using a recursive formula given by Brook and Evans (1972).

The discretized ARL's were evaluated for $t = 51$, 59, 67, 75, and 83. As t goes to infinity, ARL(t) approaches the continuous-state ARL for an EWMA control scheme. Following the procedure of Brook and Evans (1972), the continuous-state ARL's were approximated as the least squares intercept of the quadratic equation in the reciprocal of t, ARL(t) = asymptotic ARL + B/t + C/t^2. Except for small values of λ (e.g., $\lambda < .25$), the discretized ARL's quickly converged to the continuous-state ARL's. For small values of λ, we found the convergence to be slow. As a numerical accuracy check, we compared our zero-state ARL's with those given by Crowder (1987) and our steady-state ARL's with those given by Robinson and Ho (1978). Our results agree with Crowder's results to approximately three significant digits but only agree with Robinson and Ho's results for large values of λ. Examination of Robinson and Ho's procedure revealed that their algorithm did not converge for small values of λ.

APPENDIX B: ENHANCEMENTS

B.1 FIR Feature

The FIR feature requires the simultaneous implementation of two one-sided EWMA control schemes with different starting values. One EWMA is started above the target value, and the other is started below the target value. An out-of-control signal is given if the EWMA started on the high side falls outside the upper control limit or if the EWMA started on the low side falls outside the lower control limit.

The transition probability matrix requires t^2 in-control states, with $p_{jkj'k'}$ representing the probability that the EWMA with an HS on the high side moves from state j to state k and the EWMA with an HS on the low side moves from state j' to state k'. An upper bound for the FIR ARL's can be obtained using only t in-control states. This approach requires two initial probability vectors,

$$\mathbf{p}_L^T = (0, \ldots, 1, \ldots, 0, 0, 0, \ldots, 0, \ldots, 0)$$

and

$$\mathbf{p}_U^T = (0, \ldots, 0, \ldots, 0, 0, 0, \ldots, 1, \ldots, 0),$$

where $\mathbf{p}_L^T$ represents the initial probability vector for the EWMA with an HS on the low side and $\mathbf{p}_U^T$ represents the initial probability vector for the EWMA with an HS on the high side. When the desired HS does not correspond to the midpoint of any discrete state, the ARL can be approximated using quadratic interpolation of the three closest states.

Similarly, the transition probability matrix requires two out-of-control states,

$$\mathbf{P} = \begin{pmatrix} 1 & \mathbf{0}^T & 0 \\ \mathbf{1}_1 & \mathbf{R} & \mathbf{h}_1 \\ 0 & \mathbf{0}^T & 1 \end{pmatrix},$$

where the vector $\mathbf{1}_1$ contains the probabilities that the EWMA of the low-sided scheme goes out-of-control on the low side and $\mathbf{h}_1$ contains the probabilities that the EWMA of the high-sided scheme goes out-of-control on the high side. The upper bound for the FIR ARL's is given by

$$\begin{aligned}
\mathrm{ARL_{FIR}}(t) &= \sum_{i=1}^{\infty} i \Pr(\mathrm{RL} = i) \\
&= \sum_{i=1}^{\infty} i \Pr(\mathrm{RL}_L = i \text{ or } \mathrm{RL}_H = i) \\
&\le \sum_{i=1}^{\infty} i\{\Pr(\mathrm{RL}_L = i) + \Pr(\mathrm{RL}_H = i)\} \\
&= \sum_{i=1}^{\infty} i\{\mathbf{p}_H^T(\mathbf{h}_i - \mathbf{h}_{i-1}) + \mathbf{p}_L^T(\mathbf{1}_i - \mathbf{1}_{i-1})\} \\
&= \sum_{i=1}^{\infty} i(\mathbf{p}_H^T\mathbf{R}^{i-1}\mathbf{h}_1 + \mathbf{p}_L^T\mathbf{R}^{i-1}\mathbf{1}_1) \\
&= \mathbf{p}_H^T\left(\sum_{i=1}^{\infty} i\,\mathbf{R}^{i-1}\right)\mathbf{h}_1 + \mathbf{p}_L^T\left(\sum_{i=1}^{\infty} i\,\mathbf{R}^{i-1}\right)\mathbf{1}_1 \\
&= \mathbf{p}_H^T(\mathbf{I} - \mathbf{R})^{-2}\mathbf{h}_1 + \mathbf{p}_L^T(\mathbf{I} - \mathbf{R})^{-2}\mathbf{1}_1,
\end{aligned}$$

where $\mathbf{1}_i = \mathbf{1}_{i-1} + \mathbf{R}^{i-1}\mathbf{1}_1$ and $\mathbf{h}_i = \mathbf{h}_{i-1} + \mathbf{R}^{i-1}\mathbf{h}_1$. Using simulation, we found that this upper bound closely approximates the FIR ARL's whenever the shift in the process target value is larger than .25 σ_Y. The FIR ARL's for shifts equal to 0 were obtained by simulation. We plan to implement the exact method for calculating FIR ARL's in our future work when we compare and optimize methods of selecting a FIR feature.

B.2 Combined Shewhart EWMA

The properties of the combined Shewhart EWMA can be obtained by modifying the transition probability matrix for an EWMA control scheme. The modified one-step transition probabilities are given by

$$p_{jk} \approx \Pr[\min\{\mathrm{SCL}_U, \max(\mathrm{SCL}_L, Y_L)\} < Y_i \le \max\{\mathrm{SCL}_L, \min(\mathrm{SCL}_U, Y_U)\}],$$

$$Y_L = \lambda^{-1}\{(S_k - \delta) - (1 - \lambda)S_j\},$$

and

$$Y_U = \lambda^{-1}\{(S_k + \delta) - (1 - \lambda)S_j\}, \qquad j = -m, -m + 1, \ldots, m,$$

where SCL_U and SCL_L represent the upper and lower Shewhart control limits, respectively.

B.3 Robust EWMA

The run-length properties of a robust EWMA using a two-in-a-row rule can be obtained by modifying the transition probability matrix of a combined Shewhart EWMA. The in-control transition probabilities will be identical to the in-control transition probabilities for the combined Shewhart EWMA, with the Shewhart limits representing outlier limits. When a single outlier is observed, the control statistic remains in the same state and a counter is set. If the next observation lies within the outlier limits, the counter is reset to 0; otherwise an out-of-control signal is given.

The exact form of the transition probability matrix was given by Lucas and Crosier (1982b). Although they also gave the ARL vector associated with this transition probability matrix, a simplified form of the ARL vector is given by

$$\mathrm{ARL}_R = (\mathbf{p}_1^T, \mathbf{p}_2^T)\begin{pmatrix} (1 + c)\{\mathbf{I} - (1 + c)\mathbf{R}_c\}^{-1}\mathbf{1} \\ \{\mathbf{I} - (1 + c)\mathbf{R}_c\}^{-1}\mathbf{1} \end{pmatrix}.$$

A similar result also applies to the modified cyclical steady-state probability vector discussed in Appendix A.

[Received September 1987. Revised April 1989.]

REFERENCES

Box, G. E. P., Jenkins, G. M., and MacGregor, J. F. (1974), "Some Recent Advances in Forecasting and Control," *Applied Statistics*, 23, 158–179.

Brook, D., and Evans, D. A. (1972), "An Approach to the Probability Distribution of CUSUM Run Length," *Biometrika*, 59, 539–549.

Crosier, R. B. (1986), "A New Two-Sided Cumulative Sum Quality Control Scheme," *Technometrics*, 28, 187–194.

Crowder, S. V. (1987), "A Simple Method for Studying Run Length Distributions of Exponentially Weighted Moving Average Control Charts," *Technometrics*, 29, 401–407.

Goel, A. L. (1968), "A Comparative and Economic Investigation of $\bar{X}$ and Cumulative Sum Control Charts," unpublished Ph.D. dissertation, University of Wisconsin, Dept. of Industrial Engineering.

Ho, T. Y. (1978), "Optimal Design of Geometric Moving Average Control Charts," unpublished Ph.D. dissertation, New York University, School of Engineering and Science.

Hunter, J. S. (1986), "The Exponentially Weighted Moving Average," *Journal of Quality Technology*, 18, 203–210.

Lucas, J. M. (1973), "A Modified *V*-Mask Control Scheme," *Technometrics*, 15, 833–847.

——— (1976), "The Design and Use of *V*-Mask Control Schemes," *Journal of Quality Technology*, 8, 1–12.

——— (1982), "Combined Shewhart-CUSUM Quality Control Schemes," *Journal of Quality Technology*, 14, 51–59.

Lucas, J. M., and Crosier, R. B. (1982a), "Fast Initial Response for CUSUM Quality-Control Schemes: Give Your CUSUM a Head Start," *Technometrics*, 24, 199–205.

——— (1982b), "Robust CUSUM: A Robustness Study for CUSUM Quality Control Schemes," *Communications in Statistics—Theory and Methods*, 11, 2669–2687.

Lucas, J. M., and Saccucci, M. S. (1986), "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements (Preliminary Results)," unpublished paper presented at the American Statistical Association spring regional meetings in Atlanta.

——— (1987), "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements," Faculty Working Series Paper 87-5, Drexel University, Dept. of Quantitative Methods.

Montgomery, D. C., Gardiner, J. S., and Pizzano, B. A. (1987), "Statistical Process Control Methods for Detecting Small Process Shifts," in *Frontiers in Statistical Quality Control*, eds. H.-J. Lenz, G. B. Wetherill, and P. T. Wilrich, Heidelberg, West Germany: Physica-Verlag, pp. 161–178.

Muth, J. F. (1960), "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55, 299–306.

Roberts, S. W. (1959), "Control Chart Tests Based on Geometric Moving Averages," *Technometrics*, 1, 239–250.

——— (1966), "A Comparison of Some Control Chart Procedures," *Technometrics*, 8, 411–430.

Robinson, P. B., and Ho, T. Y. (1978), "Average Run Lengths of Geometric Moving Average Charts by Numerical Methods," *Technometrics*, 20, 85–93.

Waldmann, K. H. (1986), "Bounds for the Distribution of the Run Length of Geometric Moving Average Charts," *Applied Statistics*, 35, 151–158.

Yashchin, E. (1987), "Some Aspects of the Theory of Statistical Control Schemes," *IBM Journal of Research and Development*, 31, 199–205.