The Autoregressive T^2 Chart for Monitoring Univariate Autocorrelated Processes

DANIEL W. APLEY

Texas A&M University, College Station, TX 77843-3131

FUGEE TSUNG

Hong Kong University of Science and Technology, Kowloon, Hong Kong

In this paper we investigate the autoregressive T^2 control chart for statistical process control of autocorrelated processes. The method involves the monitoring, using Hotelling's T^2 statistic, of a vector formed from a moving window of observations of the univariate autocorrelated process. It is shown that the T^2 statistic can be decomposed into the sum of the squares of the residual errors for various order autoregressive time series models fit to the process data. Guidelines for designing the autoregressive T^2 chart are presented, and its performance is compared to that of residual-based CUSUM and Shewhart individual control charts. The autoregressive T^2 chart has a number of characteristics, including some level of robustness with respect to modeling errors, that make it an attractive alternative to residual-based control charts for autocorrelated processes.

Introduction

In recent years, statistical process control (SPC) for autocorrelated processes has received a great deal of attention, due in part to the increasing prevalence of autocorrelation in process inspection data. With improvements in measurement and data collection technology, processes can be sampled at higher rates, which often leads to data autocorrelation. It is well known that the run length properties of common SPC methods like CUSUM and $\overline{X}$ charts are strongly affected by data autocorrelation, and the in-control average run length (ARL) can be much shorter than intended if the autocorrelation is positive (Johnson and Bagshaw (1974) and Vasilopoulos and Stamboulis (1978)).

The most widely researched methods of SPC for autocorrelated processes are residual-based control charts, which involve fitting some form of autoregressive moving average (ARMA) model to the data and monitoring the model residuals (i.e., the one-step-ahead prediction errors). If the model is exact, then the model residuals are independent. Consequently, standard SPC control charts can be applied to the residuals with well understood in-control run length properties (see e.g., Alwan and Roberts (1988); Apley and Shi (1999); Berthouex, Hunter, and Pallesen (1978); Chow, Wu, and Ermer (1979); English, Krishnamurthi, and Sastri (1991); Lin and Adams (1996); Montgomery and Mastrangelo (1991); Runger, Willemain, and Prabhu (1995); Superville and Adams (1994); Vander Wiel (1996); and Wardell, Moskowitz, and Plante (1994)).

Krieger, Champ, and Alwan (1992) and Alwan and Alwan (1994) proposed a different approach to monitoring autocorrelated processes. The basic idea was to form a multivariate vector of a moving window of observations from a univariate autocorrelated process, and then apply multivariate control charts. Krieger et al. (1992) applied a multivariate CUSUM, and Alwan and Alwan (1994) applied a T^2 chart to the constructed vectors. In this paper, we analyze a slightly modified version of the T^2 approach of Alwan and Alwan (1994). One difference between the approaches is that Alwan and Alwan (1994) recommended time delays between samples so that the constructed vectors have less statistical dependency,

---

Dr. Apley is an Assistant Professor in the Department of Industrial Engineering. His email address is apley@tamu.edu.

Dr. Tsung is an Assistant Professor in the Department of Industrial Engineering and Engineering Management. His email address is season@uxmail.ust.hk.

whereas we do not. In addition, Alwan and Alwan (1994) focused on the case of a first-order autoregressive processes, whereas we take a more general approach. The method, which will be referred to as the autoregressive T^2 control chart, has a close connection to residual-based control charts. It is shown in this paper that the autoregressive T^2 statistic can be decomposed into the sum of the squares of the residual errors for various order autoregressive models, similar to well-known decompositions of multivariate T^2 statistics (see, for example, Hawkins (1993) and Mason, Tracy, and Young (1997)). In spite of the relationship to residual-based control charts, the autoregressive T^2 chart does not explicitly require an ARMA model of the process. Only the process autocovariance function up to a pre-specified lag is required.

Guidelines for designing the autoregressive T^2 chart are presented, and its performance is compared to residual-based CUSUM, recently studied by Lu and Reynolds (2001), and residual-based Shewhart individual charts. In many situations, the autoregressive T^2 chart compares favorably. In addition, it possesses other advantages over residual-based CUSUMs. For example, the optimal CUSUM design (i.e., selection of reference value and decision threshold) that minimizes the out-of-control ARL for a specified in-control ARL depends strongly on the mean shift size of interest (Montgomery (2001)). A single autoregressive T^2 chart design, in contrast, is often suitable for a wide range of mean shift sizes. This has practical significance, in that it is generally difficult to select one mean shift size that is of primary interest. It is, therefore, desirable that a control chart performs well over a range of mean shift sizes.

Another disadvantage of residual-based CUSUMs is that an accurate ARMA model of the process is required. If the model is inaccurate, the residuals will not be uncorrelated, and the in-control ARL of a residual-based CUSUM may be substantially shorter than what is intended. This results in too many nuisance alarms. Since ARMA models are always estimated from data, some level of model uncertainty will be present; if the data are limited, then the uncertainty may be large. The autoregressive T^2 chart incorporates a mechanism for taking into account model uncertainty (or, analogously, uncertainty in the autocovariance structure of the process) due to limited data from which the process characteristics must be estimated.

In the remainder of this paper, we present the autoregressive T^2 chart and a performance comparison with residual-based control charts, provide an interpretation in terms of an autoregressive decomposition of the T^2 statistic, give guidelines for designing the chart, and investigate issues related to model uncertainty and limited data.

The Autoregressive T^2 Control Chart

Let x_t represent a measurement from an autocorrelated process at sampling instant t. Suppose the process is Gaussian with mean μ and autocovariance function $\gamma_k = E[(x_t - \mu)(x_{t+k} - \mu)]$, where E[·] denotes the expectation operator. Let μ_0 denote the mean when the process is in-control, and assume μ_0 is known. Consider the sequence of p-dimensional vectors $\mathbf{X}_t = [x_{t-p+1}\ x_{t-p+2}\ \ldots\ x_t]'$ formed from observations of the univariate process. Assume the process x_t is such that $\mathbf{X}_t$ is multivariate normal. This will be the case if, for example, x_t follows a stable ARMA model. The covariance matrix of $\mathbf{X}_t$ is

$$\boldsymbol{\Sigma} = \begin{pmatrix} \gamma_0 & \gamma_1 & \cdots & \gamma_{p-1} \\ \gamma_1 & \gamma_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \gamma_1 \\ \gamma_{p-1} & \cdots & \gamma_1 & \gamma_0 \end{pmatrix}, \quad (1)$$

and when the process is in-control the mean vector is $\boldsymbol{\mu}_0 = [\mu_0\ \mu_0\ \ldots\ \mu_0]'$. If $\boldsymbol{\Sigma}$ is known, the T^2 statistic

$$T_t^2 = [\mathbf{X}_t - \boldsymbol{\mu}_0]' \boldsymbol{\Sigma}^{-1} [\mathbf{X}_t - \boldsymbol{\mu}_0], \quad (2)$$

which follows a chi-square distribution with p degrees-of-freedom when the process is in-control, could be used to monitor for departures from the in-control state. This method will be referred to as the autoregressive T^2 chart. One can specify a false alarm probability α and compare T_t^2 to the $1-\alpha$ percentile of the chi-square distribution with p degrees-of-freedom, denoted $\chi^2(1-\alpha, p)$. We discuss how to specify α to ensure a desired in-control ARL later in this paper.

In Alwan and Alwan (1994), the intent was that $\mathbf{X}_t$ be formed from a sample of p successive observations of the process. The next vector would be $\mathbf{X}_{t+p}$, formed from a sample of observations that do not overlap those used to form the previous vector. It was also suggested that there could be time delays between samples, so that the T^2 statistics are closer to being independent. The intent in this paper is slightly different, in that T_t^2 is to be formed for $t = p, p+1, p+2, p+3, \ldots$. In other words, T_t^2 is calculated and charted for *every* value of t, as opposed to $t = p, 2p, 3p, \ldots$.

The matrix $\boldsymbol{\Sigma}$ could be considered known if a sufficiently large data set were available to estimate the autocovariance function. The situation where the data is insufficient to assume that the T^2 statistic in Equation (2) follows a chi-square distribution is treated later in this paper and is closely related to the issue of model uncertainty in residual-based control charts. An alternative to directly estimating the autocovariance function is to fit an ARMA model to the data, and calculate the autocovariance function theoretically from the model parameters. Define $y_t = x_t - \mu$, and suppose y_t follows the ARMA(n, m) model (Box et al., 1994)

$$\Phi(B) y_t = \Theta(B) a_t, \tag{3}$$

where a_t is an independent Gaussian process with mean 0 and variance σ_a^2, and B is the backshift operator. Here $\Phi(B)$ and $\Theta(B)$ are polynomials in B of degree n and m, respectively, and are parameterized as $\Phi(B) = 1 - \phi_1 B - \phi_2 B^2 - \ldots - \phi_n B^n$ and $\Theta(B) = 1 - \theta_1 B - \theta_2 B^2 - \ldots - \theta_m B^m$.

For the special case of a first order model, when $n = m = 1$, a closed-form expression for the autocovariance function is (Pandit and Wu (1990))

$$\gamma_k = \begin{cases} \frac{\sigma_a^2(\phi_1-\theta_1)(1-\theta_1\phi_1)}{(1-\phi_1^2)}(\phi_1)^{k-1} & k \geq 1 \\ \frac{\sigma_a^2(1-2\theta_1\phi_1+\theta_1^2)}{(1-\phi_1^2)} & k = 0. \end{cases} \tag{4}$$

Although no closed-form expression exists for the more general situation, γ_k can easily be calculated recursively from the ARMA parameters, as described in Appendix A. The matrix $\boldsymbol{\Sigma}$ would then be formed directly from γ_k via Equation (1). In either method, σ_a^2 and the ARMA parameters must be known or estimated using standard time series modeling procedures (e.g., as described in Box et al., 1994). If the process is not well represented by an ARMA model, then γ_k must be estimated directly from the data in order to form the expressions in Equations (1) and (2). Unless otherwise noted, it will be assumed that the sample of data from which the ARMA parameters or γ_k are estimated is large enough that the effects of modeling errors can be ignored. In a later section of the paper we address the issue of model uncertainty.

Performance Comparison

Monte Carlo simulation results comparing the autoregressive T^2 control chart with residual-based CUSUM and Shewhart individual charts for a variety of ARMA(1,1) processes are shown in Tables 1 through 3. The residual-based charts involve applying standard (two-sided) CUSUM and Shewhart charts to the ARMA model residuals (see e.g., Alwan and Roberts (1988); Runger, Willemain, and Prabhu (1995); and Superville and Adams (1994)). We used 10,000 Monte Carlo trials in all simulations. The data were generated as an ARMA(1,1) process, according to Equation (3), with $\sigma_a^2 = 1$ and $\mu_0 = 0$. The true model was assumed known. A range of values for ϕ_1 and θ_1 were used. Only positive values for ϕ_1 were considered, however, since negative values of ϕ_1 generally result in processes that have negative autocorrelation, a situation that is not commonly encountered in industry. A range of out-of-control mean shift magnitudes, denoted μ_1 through μ_4 in Tables 1 through 3, were also considered. The mean shifts are expressed in units of σ_a (=1) for each model. The particular values of μ_1 through μ_4 vary with the model parameters, and were chosen to span a range from what can be considered difficult to detect to what can be considered easy to detect. An alternative would have been to use a single set of mean shift magnitudes (scaled in terms of the process standard deviation σ_x) consistent for each model. It was felt that the comparisons would have been less informative had this been done. For example, a mean shift of $2.5\sigma_x$ (=2.88) for the AR(1) process with $\phi_1 = 0.5$ yields an interesting comparison between the various tests. In contrast, a mean shift of $2.5\sigma_x$ (=12.55) for the AR(1) process with $\phi_1 = 0.98$ would have been detected by all tests on the initial observation with a probability of almost one. Table 4 expresses the mean shifts considered for each model in units of σ_x.

For each model and each mean shift magnitude, four different autoregressive T^2 charts (with $p = 2, 5, 10$, and 20), four different CUSUM charts (each with different reference value K and decision threshold H), and a Shewhart chart were compared. All charts were designed so that the in-control ARL was 500. Results for other in-control ARLs exhibited the same general trends and are not shown. A subsequent section of this paper discusses guidelines for selecting p and specifying the α that is needed to achieve a desired in-control ARL for the autoregressive T^2 chart.

The Shewhart chart signaled when an individual residual fell outside the control limits, which coincided with the upper and lower 0.001 percentiles

TABLE 1. Comparison of Out-of-Control ARLs for Various Magnitude Mean Shifts and Various ARMA(1,1) Model Parameters. (All charts have an in-control ARL of 500.)

		Mean Shift				T^2 ARL				CUSUM ARL		Shewhart ARL	
ϕ_1	θ_1	μ_1	μ_2	p	α	μ_1	μ_2	K	H	μ_1	μ_2	μ_1	μ_2
0.98	0	3	4	2	0.0037	217.9	67.9	0.30	7.56	396.6	325.6	264.0	88.8
0.98	0	3	4	5	0.0049	253.8	94.2	0.50	5.07	412.3	327.3		
0.98	0	3	4	10	0.0073	278.6	122.6	0.75	3.54	388.6	233.0		
0.98	0	3	4	20	0.0122	302.7	150.5	1.00	2.67	344.7	161.0		
0.9	−0.9	2	3	2	0.0034	351.1	226.4	0.10	14.80	233.1	144.5	300.5	76.3
0.9	−0.9	2	3	5	0.0048	342.8	148.7	0.20	10.00	267.0	165.9		
0.9	−0.9	2	3	10	0.0073	271.2	56.0	0.30	7.56	298.9	194.3		
0.9	−0.9	2	3	20	0.012	201.2	19.4	0.50	5.07	354.1	242.3		
0.9	−0.5	3	4	2	0.0033	144.2	50.3	0.10	14.80	106.6	71.4	213.3	60.5
0.9	−0.5	3	4	5	0.0046	139.8	33.2	0.20	10.00	118.1	72.5		
0.9	−0.5	3	4	10	0.007	160.0	42.2	0.30	7.56	141.8	83.4		
0.9	−0.5	3	4	20	0.0119	173.1	54.0	0.50	5.07	192.0	113.6		
0.9	0	2	3	2	0.003	181.9	65.1	0.10	14.80	106.9	60.6	356.5	180.7
0.9	0	2	3	5	0.0046	222.7	88.3	0.20	10.00	118.5	58.7		
0.9	0	2	3	10	0.0071	253.8	108.2	0.30	7.56	142.2	65.3		
0.9	0	2	3	20	0.0119	262.4	121.8	0.50	5.07	193.2	91.8		
0.9	0.5	2	3	2	0.003	52.7	13.5	0.20	10.00	34.6	14.8	218.8	73.2
0.9	0.5	2	3	5	0.0046	67.1	14.2	0.30	7.56	35.4	11.7		
0.9	0.5	2	3	10	0.0071	80.6	17.2	0.50	5.07	44.7	9.7		
0.9	0.5	2	3	20	0.012	90.7	20.9	0.75	3.54	66.8	11.1		
0.5	−0.9	2	3	2	0.0031	97.5	31.6	0.20	10.00	28.3	15.9	131.5	22.6
0.5	−0.9	2	3	5	0.0049	82.7	20.5	0.30	7.56	27.4	13.7		
0.5	−0.9	2	3	10	0.0073	74.1	12.9	0.50	5.07	32.5	12.6		
0.5	−0.9	2	3	20	0.0123	57.0	7.7	0.75	3.54	46.4	13.3		
0.5	0	1	2	2	0.003	112.0	18.0	0.20	10.00	31.1	12.2	198.9	48.1
0.5	0	1	2	5	0.0046	117.2	17.4	0.30	7.56	30.9	10.4		
0.5	0	1	2	10	0.0071	110.0	16.5	0.50	5.07	38.6	9.2		
0.5	0	1	2	20	0.0121	99.3	16.3	0.75	3.54	56.2	10.0		
0.5	−0.5	2	3	2	0.003	54.7	13.8	0.20	10.00	20.5	11.8	111.6	29.7
0.5	−0.5	2	3	5	0.0047	50.0	11.8	0.30	7.56	18.9	9.9		
0.5	−0.5	2	3	10	0.0072	46.5	11.4	0.50	5.07	20.2	8.6		
0.5	−0.5	2	3	20	0.012	43.7	11.8	0.75	3.54	27.2	8.9		

of the standard normal distribution. The CUSUM chart signaled when either the upper or lower one-sided CUSUM statistic exceeded H. The upper and lower CUSUM statistics are calculated recursively via

$$S_t^+ = \max\{0, S_{t-1}^+ + e_t - K\}$$

and

$$S_t^- = \max\{0, S_{t-1}^- - e_t - K\},$$

respectively, where e_t is the residual for the observation at time t.

In the Monte Carlo simulations, the process, the residuals, and the test statistics were first allowed to reach steady-state before the mean shift was introduced. The run length was taken to be the time from when the shift was introduced to when the chart first signaled. Consequently, the ARL can be viewed as a steady-state ARL. The steady-state ARL is an appropriate measure of performance (Lu and Reynolds (1999)) if one is interested in measuring signaling properties for mean shifts that occur after the chart has been running for some time. If one were more

TABLE 1. Continued

		Mean Shift				T^2 ARL				CUSUM ARL		Shewhart ARL	
ϕ_1	θ_1	μ_3	μ_4	p	α	μ_3	μ_4	K	H	μ_3	μ_4	μ_3	μ_4
0.98	0	5	6	2	0.0037	10.2	1.46	0.30	7.56	251.9	175.9	14.3	1.84
0.98	0	5	6	5	0.0049	19.6	2.12	0.50	5.07	195.5	74.73		
0.98	0	5	6	10	0.0073	31.3	4.32	0.75	3.54	81.9	15.53		
0.98	0	5	6	20	0.0122	47.1	9.11	1.00	2.67	37.5	4.64		
0.9	−0.9	4	5	2	0.0034	116.7	44.18	0.10	14.80	97.5	71.31	4.4	1.05
0.9	−0.9	4	5	5	0.0048	30.8	2.72	0.20	10.00	106.0	70.70		
0.9	−0.9	4	5	10	0.0073	3.6	1.20	0.30	7.56	123.1	76.92		
0.9	−0.9	4	5	20	0.012	1.9	1.29	0.50	5.07	144.5	61.59		
0.9	−0.5	5	6	2	0.0033	11.6	2.00	0.10	14.80	51.5	38.76	8.1	1.32
0.9	−0.5	5	6	5	0.0046	4.6	1.12	0.20	10.00	47.6	32.40		
0.9	−0.5	5	6	10	0.007	6.0	1.22	0.30	7.56	50.8	31.83		
0.9	−0.5	5	6	20	0.0119	9.6	1.56	0.50	5.07	53.6	17.44		
0.9	0	4	5	2	0.003	14.7	2.56	0.10	14.80	39.2	27.51	48.6	6.66
0.9	0	4	5	5	0.0046	23.3	3.88	0.20	10.00	33.0	20.35		
0.9	0	4	5	10	0.0071	31.1	6.00	0.30	7.56	33.2	17.30		
0.9	0	4	5	20	0.0119	39.5	8.62	0.50	5.07	38.5	13.56		
0.9	0.5	4	5	2	0.003	3.3	1.27	0.20	10.00	7.6	4.46	12.5	1.71
0.9	0.5	4	5	5	0.0046	2.7	1.16	0.30	7.56	5.0	2.89		
0.9	0.5	4	5	10	0.0071	3.4	1.28	0.50	5.07	3.1	1.84		
0.9	0.5	4	5	20	0.012	4.5	1.56	0.75	3.54	2.4	1.26		
0.5	−0.9	4	5	2	0.0031	9.6	3.22	0.20	10.00	10.6	7.67	2.2	1.06
0.5	−0.9	4	5	5	0.0049	4.5	1.46	0.30	7.56	8.5	5.76		
0.5	−0.9	4	5	10	0.0073	2.5	1.31	0.50	5.07	6.4	3.46		
0.5	−0.9	4	5	20	0.0123	2.4	1.45	0.75	3.54	4.5	1.75		
0.5	0	3	4	2	0.003	4.3	1.61	0.20	10.00	7.3	5.10	10.6	2.32
0.5	0	3	4	5	0.0046	4.4	1.76	0.30	7.56	5.8	3.93		
0.5	0	3	4	10	0.0071	4.9	2.12	0.50	5.07	4.4	2.74		
0.5	0	3	4	20	0.0121	5.9	2.66	0.75	3.54	3.7	1.99		
0.5	−0.5	4	5	2	0.003	3.8	1.41	0.20	10.00	8.1	6.02	5.4	1.32
0.5	−0.5	4	5	5	0.0047	3.1	1.28	0.30	7.56	6.5	4.61		
0.5	−0.5	4	5	10	0.0072	3.4	1.47	0.50	5.07	4.8	2.92		
0.5	−0.5	4	5	20	0.012	4.3	1.87	0.75	3.54	3.6	1.68		

interested in mean shifts that are present at the time the control chart is started, zero-state ARLs would be more appropriate. In this situation, the T^2 chart ARL would be at least as large as p, since the T^2 statistic is not formed until at least p observations are available. One could use a modified procedure in which the T^2 statistic is formed beginning from the initial observation, however, if $\{x_0, x_{-1}, \ldots, x_{-p+2}\}$ are taken to be zero. This is similar to what is done in residual-based charts, if the residual is calculated on the initial observation with $\{x_0, x_{-1}, x_{-2}, \ldots\}$ taken to be zero. This is not recommended unless the control limits over the initial period are modified appropriately so as to avoid an excessively large probability of a false alarm before the test statistic reaches its steady-state distribution. All ARL results presented below are for the steady-state scenario.

In Table 1 we show the out-of-control ARLs for the different charts with the various size mean shifts. Generally speaking, the autoregressive T^2 chart (with appropriate p) outperforms the residual-based CUSUM and Shewhart charts when ϕ_1 is large (e.g. 0.9 or 0.98) and θ_1 is negative. For example, when $\phi_1 = 0.9$ and $\theta_1 = -0.9$, the autoregressive T^2 chart with $p = 20$ has substantially lower out-of-

TABLE 2. Comparison of P_1 Values for Various Magnitude Mean Shifts and Various ARMA(1,1) Model Parameters. (All charts have an in-control ARL of 500.)

		Mean Shift				T^2 P_1				CUSUM P_1		Shewhart P_1	
ϕ_1	θ_1	μ_1	μ_2	p	α	μ_1	μ_2	K	H	μ_1	μ_2	μ_1	μ_2
0.98	0	3	4	2	0.0037	0.43	0.79	0.30	7.56	0.00	0.00	0.46	0.82
0.98	0	3	4	5	0.0049	0.31	0.66	0.50	5.07	0.00	0.06		
0.98	0	3	4	10	0.0073	0.23	0.54	0.75	3.54	0.10	0.38		
0.98	0	3	4	20	0.0122	0.18	0.44	1.00	2.67	0.25	0.63		
0.9	−0.9	2	3	2	0.0034	0.05	0.18	0.10	14.80	0.00	0.00	0.14	0.46
0.9	−0.9	2	3	5	0.0048	0.07	0.24	0.20	10.00	0.00	0.00		
0.9	−0.9	2	3	10	0.0073	0.06	0.22	0.30	7.56	0.00	0.00		
0.9	−0.9	2	3	20	0.012	0.06	0.18	0.50	5.07	0.00	0.00		
0.9	−0.5	3	4	2	0.0033	0.31	0.64	0.10	14.80	0.00	0.00	0.46	0.82
0.9	−0.5	3	4	5	0.0046	0.29	0.65	0.20	10.00	0.00	0.00		
0.9	−0.5	3	4	10	0.007	0.23	0.54	0.30	7.56	0.00	0.00		
0.9	−0.5	3	4	20	0.0119	0.18	0.43	0.50	5.07	0.00	0.06		
0.9	0	2	3	2	0.003	0.12	0.40	0.10	14.80	0.00	0.00	0.14	0.46
0.9	0	2	3	5	0.0046	0.08	0.30	0.20	10.00	0.00	0.00		
0.9	0	2	3	10	0.0071	0.06	0.23	0.30	7.56	0.00	0.00		
0.9	0	2	3	20	0.0119	0.06	0.18	0.50	5.07	0.00	0.00		
0.9	0.5	2	3	2	0.003	0.10	0.35	0.20	10.00	0.00	0.00	0.14	0.46
0.9	0.5	2	3	5	0.0046	0.08	0.29	0.30	7.56	0.00	0.00		
0.9	0.5	2	3	10	0.0071	0.06	0.23	0.50	5.07	0.00	0.00		
0.9	0.5	2	3	20	0.012	0.06	0.18	0.75	3.54	0.01	0.10		
0.5	−0.9	2	3	2	0.0031	0.06	0.20	0.20	10.00	0.00	0.00	0.14	0.46
0.5	−0.9	2	3	5	0.0049	0.07	0.24	0.30	7.56	0.00	0.00		
0.5	−0.9	2	3	10	0.0073	0.06	0.21	0.50	5.07	0.00	0.00		
0.5	−0.9	2	3	20	0.0123	0.05	0.16	0.75	3.54	0.01	0.10		
0.5	0	1	2	2	0.003	0.02	0.11	0.20	10.00	0.00	0.00	0.02	0.14
0.5	0	1	2	5	0.0046	0.01	0.08	0.30	7.56	0.00	0.00		
0.5	0	1	2	10	0.0071	0.01	0.06	0.50	5.07	0.00	0.00		
0.5	0	1	2	20	0.0121	0.02	0.05	0.75	3.54	0.00	0.01		
0.5	−0.5	2	3	2	0.003	0.09	0.33	0.20	10.00	0.00	0.00	0.14	0.46
0.5	−0.5	2	3	5	0.0047	0.08	0.29	0.30	7.56	0.00	0.00		
0.5	−0.5	2	3	10	0.0072	0.06	0.22	0.50	5.07	0.00	0.00		
0.5	−0.5	2	3	20	0.012	0.05	0.16	0.75	3.54	0.01	0.10		

control ARLs than any of the CUSUMs or the Shewhart chart. The exception is that when $\mu = 5$, the Shewhart chart has a slightly lower ARL. Another general trend is that the performance of the autoregressive T^2 chart relative to the CUSUM improves as the mean shift magnitude increases. For example, consider the case where $\phi_1 = 0.9$ and $\theta_1 = 0$. For mean shifts of 2 and 3, the CUSUMs with K = 0.1 and 0.2 have the lowest ARL, whereas for larger mean shifts of 4 and 5, the autoregressive T^2 chart with $p = 2$ has substantially better ARL values than any of the CUSUMs. When $\phi_1 = 0.5$ and $\theta_1 = 0$ or -0.5 (i.e., when the autocorrelation is moderate) CUSUMs with appropriate K values perform substantially better for small to moderate mean shifts, and the autoregressive T^2 chart with $p = 2$ or 5 performs the best for large mean shifts.

One primary advantage of the autoregressive T^2 chart over the CUSUM chart is that, for a given ARMA model, the optimal choice of p does not seem to depend strongly on the magnitude of the mean

TABLE 2. Continued

		Mean Shift				T^2 P_1				CUSUM P_1		Shewhart P_1	
ϕ_1	θ_1	μ_3	μ_4	p	α	μ_3	μ_4	K	H	μ_3	μ_4	μ_3	μ_4
0.98	0	5	6	2	0.0037	0.96	1.00	0.30	7.56	0.00	0.03	0.97	1.00
0.98	0	5	6	5	0.0049	0.91	0.99	0.50	5.07	0.28	0.67		
0.98	0	5	6	10	0.0073	0.84	0.97	0.75	3.54	0.76	0.96		
0.98	0	5	6	20	0.0122	0.75	0.93	1.00	2.67	0.91	0.99		
0.9	−0.9	4	5	2	0.0034	0.42	0.70	0.10	14.80	0.00	0.00	0.82	0.97
0.9	−0.9	4	5	5	0.0048	0.56	0.86	0.20	10.00	0.00	0.00		
0.9	−0.9	4	5	10	0.0073	0.52	0.82	0.30	7.56	0.00	0.00		
0.9	−0.9	4	5	20	0.012	0.43	0.74	0.50	5.07	0.06	0.28		
0.9	−0.5	5	6	2	0.0033	0.89	0.98	0.10	14.80	0.00	0.00	0.97	1.00
0.9	−0.5	5	6	5	0.0046	0.91	0.99	0.20	10.00	0.00	0.00		
0.9	−0.5	5	6	10	0.007	0.84	0.97	0.30	7.56	0.00	0.03		
0.9	−0.5	5	6	20	0.0119	0.74	0.93	0.50	5.07	0.28	0.67		
0.9	0	4	5	2	0.003	0.77	0.96	0.10	14.80	0.00	0.00	0.82	0.97
0.9	0	4	5	5	0.0046	0.65	0.91	0.20	10.00	0.00	0.00		
0.9	0	4	5	10	0.0071	0.54	0.84	0.30	7.56	0.00	0.00		
0.9	0	4	5	20	0.0119	0.43	0.74	0.50	5.07	0.06	0.28		
0.9	0.5	4	5	2	0.003	0.70	0.93	0.20	10.00	0.00	0.00	0.82	0.97
0.9	0.5	4	5	5	0.0046	0.65	0.91	0.30	7.56	0.00	0.00		
0.9	0.5	4	5	10	0.0071	0.54	0.84	0.50	5.07	0.06	0.28		
0.9	0.5	4	5	20	0.012	0.43	0.74	0.75	3.54	0.38	0.76		
0.5	−0.9	4	5	2	0.0031	0.48	0.76	0.20	10.00	0.00	0.00	0.82	0.97
0.5	−0.9	4	5	5	0.0049	0.57	0.86	0.30	7.56	0.00	0.00		
0.5	−0.9	4	5	10	0.0073	0.51	0.82	0.50	5.07	0.06	0.28		
0.5	−0.9	4	5	20	0.0123	0.42	0.74	0.75	3.54	0.38	0.76		
0.5	0	3	4	2	0.003	0.40	0.76	0.20	10.00	0.00	0.00	0.46	0.82
0.5	0	3	4	5	0.0046	0.29	0.65	0.30	7.56	0.00	0.00		
0.5	0	3	4	10	0.0071	0.22	0.54	0.50	5.07	0.00	0.06		
0.5	0	3	4	20	0.0121	0.16	0.43	0.75	3.54	0.10	0.38		
0.5	−0.5	4	5	2	0.003	0.68	0.92	0.20	10.00	0.00	0.00	0.82	0.97
0.5	−0.5	4	5	5	0.0047	0.65	0.91	0.30	7.56	0.00	0.00		
0.5	−0.5	4	5	10	0.0072	0.54	0.84	0.50	5.07	0.06	0.28		
0.5	−0.5	4	5	20	0.012	0.42	0.74	0.75	3.54	0.38	0.76		

shift whereas the optimal choice of K for the CUSUM does. As evidence, consider the case when $\phi_1 = 0.9$ and $\theta_1 = -0.5$. The autoregressive T^2 chart with $p = 5$ provides the lowest ARL for mean shifts of magnitude 3, 4, 5, or 6. In contrast, the optimal choices of K for the CUSUM range from approximately 0.1 to something substantially larger when the mean shift magnitude varies from 3 to 6. Of the K values considered, K = 0.5 provided the lowest ARL for $\mu = 6$. The Shewhart chart, however, had a much lower ARL. Since a Shewhart chart is the limiting form of a CUSUM as K becomes large, it is clear that the optimal CUSUM has K much larger than 0.5. The optimal K for $\mu = 3$ performs poorly when $\mu = 6$, and vice versa. Similar observations apply to almost all of the cases considered. For $\phi_1 = 0.5$ and $\theta_1 = -0.9$, the optimal p value for the autoregressive T^2 chart was approximately 20 for all mean shifts ranging from $\mu = 2$ to 5 (for $\mu = 5$, $p = 10$ provided a slightly lower ARL, however). In contrast, the optimal K for the CUSUM varied from 0.3 for $\mu = 2$ to very large (i.e., the Shewhart limiting form) for $\mu = 5$.

Invariance, with respect to the mean shift magnitude, of the optimal value of p for the autoregressive

TABLE 3. Comparison of P_5 Values for Various Magnitude Mean Shifts and Various ARMA(1,1) Model Parameters. (All charts have an in-control ARL of 500.)

		Mean Shift				T^2 P_5				CUSUM P_5		Shewhart P_5	
ϕ_1	θ_1	μ_1	μ_2	p	α	μ_1	μ_2	K	H	μ_1	μ_2	μ_1	μ_2
0.98	0	3	4	2	0.0037	0.44	0.79	0.30	7.56	0.01	0.03	0.47	0.82
0.98	0	3	4	5	0.0049	0.39	0.73	0.50	5.07	0.05	0.17		
0.98	0	3	4	10	0.0073	0.31	0.63	0.75	3.54	0.17	0.48		
0.98	0	3	4	20	0.0122	0.24	0.51	1.00	2.67	0.30	0.67		
0.9	−0.9	2	3	2	0.0034	0.06	0.19	0.10	14.80	0.00	0.00	0.31	0.78
0.9	−0.9	2	3	5	0.0048	0.17	0.53	0.20	10.00	0.00	0.00		
0.9	−0.9	2	3	10	0.0073	0.31	0.80	0.30	7.56	0.00	0.01		
0.9	−0.9	2	3	20	0.012	0.32	0.82	0.50	5.07	0.01	0.05		
0.9	−0.5	3	4	2	0.0033	0.32	0.65	0.10	14.80	0.00	0.00	0.49	0.84
0.9	−0.5	3	4	5	0.0046	0.46	0.81	0.20	10.00	0.00	0.00		
0.9	−0.5	3	4	10	0.007	0.41	0.76	0.30	7.56	0.00	0.01		
0.9	−0.5	3	4	20	0.0119	0.32	0.65	0.50	5.07	0.03	0.11		
0.9	0	2	3	2	0.003	0.13	0.42	0.10	14.80	0.00	0.00	0.15	0.47
0.9	0	2	3	5	0.0046	0.13	0.39	0.20	10.00	0.00	0.00		
0.9	0	2	3	10	0.0071	0.11	0.32	0.30	7.56	0.00	0.02		
0.9	0	2	3	20	0.0119	0.09	0.25	0.50	5.07	0.03	0.10		
0.9	0.5	2	3	2	0.003	0.18	0.52	0.20	10.00	0.00	0.07	0.18	0.54
0.9	0.5	2	3	5	0.0046	0.22	0.61	0.30	7.56	0.04	0.28		
0.9	0.5	2	3	10	0.0071	0.19	0.55	0.50	5.07	0.17	0.60		
0.9	0.5	2	3	20	0.012	0.15	0.44	0.75	3.54	0.26	0.72		
0.5	−0.9	2	3	2	0.0031	0.09	0.27	0.20	10.00	0.00	0.01	0.27	0.71
0.5	−0.9	2	3	5	0.0049	0.16	0.48	0.30	7.56	0.01	0.07		
0.5	−0.9	2	3	10	0.0073	0.23	0.67	0.50	5.07	0.06	0.24		
0.5	−0.9	2	3	20	0.0123	0.22	0.67	0.75	3.54	0.11	0.38		
0.5	0	1	2	2	0.003	0.05	0.29	0.20	10.00	0.00	0.01	0.04	0.20
0.5	0	1	2	5	0.0046	0.04	0.27	0.30	7.56	0.00	0.09		
0.5	0	1	2	10	0.0071	0.04	0.22	0.50	5.07	0.03	0.26		
0.5	0	1	2	20	0.0121	0.04	0.17	0.75	3.54	0.05	0.34		
0.5	−0.5	2	3	2	0.003	0.15	0.44	0.20	10.00	0.00	0.02	0.17	0.51
0.5	−0.5	2	3	5	0.0047	0.16	0.51	0.30	7.56	0.02	0.12		
0.5	−0.5	2	3	10	0.0072	0.14	0.45	0.50	5.07	0.08	0.32		
0.5	−0.5	2	3	20	0.012	0.11	0.34	0.75	3.54	0.13	0.43		

T^2 chart has important practical significance, since it is generally difficult to identify one single mean shift magnitude of interest. It is desirable that a test be effective at detecting a variety of mean shift magnitudes. Guidelines for selecting the optimal value of p for the autoregressive T^2 chart are provided in a following section.

Another advantage of the autoregressive T^2 chart over the CUSUM chart is that the probability of detecting the mean shift immediately after it occurs is usually much larger for the T^2 chart. Tables 2 and 3 show the probability of detecting the mean shift on the first observation (denoted by P_1) and on or before the fifth observation (denoted by P_5) after the occurrence of the shift. For many of the cases, P_1 for the autoregressive T^2 chart is almost as large as that for the Shewhart chart. In this sense, the autoregressive T^2 chart possesses desirable properties of both the Shewhart chart (fast detection of large shifts) and the CUSUM (good protection against small shifts). For many of the models, a single T^2 chart with an

TABLE 3. Continued

		Mean Shift				T^2 P_5				CUSUM P_5		Shewhart P_5	
ϕ_1	θ_1	μ_3	μ_4	p	α	μ_3	μ_4	K	H	μ_3	μ_4	μ_3	μ_4
0.98	0	5	6	2	0.0037	0.96	1.00	0.30	7.56	0.08	0.20	0.97	1.00
0.98	0	5	6	5	0.0049	0.94	0.99	0.50	5.07	0.45	0.77		
0.98	0	5	6	10	0.0073	0.89	0.98	0.75	3.54	0.81	0.97		
0.98	0	5	6	20	0.0122	0.80	0.95	1.00	2.67	0.92	0.99		
0.9	−0.9	4	5	2	0.0034	0.43	0.70	0.10	14.80	0.00	0.00	0.98	1.00
0.9	−0.9	4	5	5	0.0048	0.88	0.99	0.20	10.00	0.00	0.01		
0.9	−0.9	4	5	10	0.0073	0.99	1.00	0.30	7.56	0.02	0.06		
0.9	−0.9	4	5	20	0.012	0.99	1.00	0.50	5.07	0.17	0.48		
0.9	−0.5	5	6	2	0.0033	0.89	0.98	0.10	14.80	0.00	0.00	0.98	1.00
0.9	−0.5	5	6	5	0.0046	0.97	1.00	0.20	10.00	0.00	0.01		
0.9	−0.5	5	6	10	0.007	0.96	1.00	0.30	7.56	0.04	0.11		
0.9	−0.5	5	6	20	0.0119	0.90	0.99	0.50	5.07	0.35	0.71		
0.9	0	4	5	2	0.003	0.78	0.96	0.10	14.80	0.00	0.00	0.82	0.97
0.9	0	4	5	5	0.0046	0.74	0.94	0.20	10.00	0.01	0.04		
0.9	0	4	5	10	0.0071	0.65	0.90	0.30	7.56	0.07	0.21		
0.9	0	4	5	20	0.0119	0.53	0.82	0.50	5.07	0.30	0.61		
0.9	0.5	4	5	2	0.003	0.85	0.98	0.20	10.00	0.37	0.80	0.88	0.99
0.9	0.5	4	5	5	0.0046	0.92	0.99	0.30	7.56	0.73	0.96		
0.9	0.5	4	5	10	0.0071	0.90	0.99	0.50	5.07	0.92	1.00		
0.9	0.5	4	5	20	0.012	0.81	0.98	0.75	3.54	0.96	1.00		
0.5	−0.9	4	5	2	0.0031	0.57	0.83	0.20	10.00	0.07	0.27	0.97	1.00
0.5	−0.9	4	5	5	0.0049	0.83	0.98	0.30	7.56	0.28	0.62		
0.5	−0.9	4	5	10	0.0073	0.96	1.00	0.50	5.07	0.58	0.88		
0.5	−0.9	4	5	20	0.0123	0.96	1.00	0.75	3.54	0.77	0.97		
0.5	0	3	4	2	0.003	0.73	0.97	0.20	10.00	0.18	0.68	0.57	0.90
0.5	0	3	4	5	0.0046	0.74	0.97	0.30	7.56	0.50	0.91		
0.5	0	3	4	10	0.0071	0.65	0.95	0.50	5.07	0.76	0.98		
0.5	0	3	4	20	0.0121	0.53	0.89	0.75	3.54	0.83	0.99		
0.5	−0.5	4	5	2	0.003	0.79	0.97	0.20	10.00	0.13	0.43	0.86	0.98
0.5	−0.5	4	5	5	0.0047	0.85	0.98	0.30	7.56	0.40	0.77		
0.5	−0.5	4	5	10	0.0072	0.82	0.98	0.50	5.07	0.70	0.93		
0.5	−0.5	4	5	20	0.012	0.71	0.94	0.75	3.54	0.80	0.97		

TABLE 4. Mean Shift Magnitudes Used in the Simulations, Expressed in Units of σ_x. (The values for μ_1, μ_2, μ_3, and μ_4 are given in Tables 1–3 for each of the eight ARMA(1,1) models.)

ϕ_1	θ_1	σ_x	μ_1/σ_x	μ_2/σ_x	μ_3/σ_x	μ_4/σ_x
0.98	0	5.03	0.60	0.80	0.99	1.19
0.9	−0.9	4.25	0.47	0.71	0.94	1.18
0.9	−0.5	3.36	0.89	1.19	1.49	1.78
0.9	0	2.29	0.87	1.31	1.74	2.18
0.9	0.5	1.36	1.47	2.21	2.95	3.68
0.5	−0.9	1.90	1.05	1.58	2.10	2.63
0.5	0	1.15	0.87	1.73	2.60	3.46
0.5	−0.5	1.53	1.31	1.96	2.62	3.27

appropriate choice of p performed comparably to the optimal CUSUM for small mean shifts and the Shewhart chart for large mean shifts.

An Autoregressive Decomposition of the T^2 Statistic

Consider an nth order autoregressive model fit to $y_t = x_t - \mu$, even though y_t may not truly be an ARMA process. The AR(n) model form is

$$y_t = \beta_{1,n} y_{t-1} + \beta_{2,n} y_{t-2} + \ldots + \beta_{n,n} y_{t-n} + e_{t,n}, \quad (5)$$

where $\beta_{1,n}, \beta_{2,n}, \ldots, \beta_{n,n}$ are the AR(n) parameters which are optimal in the sense of minimizing the variance of $e_{t,n}$, the residual error for the AR(n) model. Let σ_n^2 denote the variance of $e_{t,n}$. In Appendix B we show that $\boldsymbol{\Sigma}$ can be factored as

$$\boldsymbol{\Sigma} = \mathbf{B}^{-1'} \mathbf{D} \mathbf{B}^{-1}, \quad (6)$$

where

$$\mathbf{B} = \begin{pmatrix} 1 & -\beta_{1,1} & -\beta_{2,2} & \cdots & -\beta_{p-1,p-1} \\ & 1 & -\beta_{1,2} & \cdots & -\beta_{p-2,p-1} \\ & & 1 & & \vdots \\ & 0 & & \ddots & -\beta_{1,p-1} \\ & & & & 1 \end{pmatrix}$$

is an upper triangular matrix containing the parameters for the set of AR(n) models with $n = 1, 2, \ldots, p-1$, and

$$\mathbf{D} = \begin{pmatrix} \sigma_0^2 & & & \\ & \sigma_1^2 & & \\ & & \ddots & \\ & & & \sigma_{p-1}^2 \end{pmatrix} \quad (7)$$

is a diagonal matrix containing the variances of the residual errors for the various order AR(n) models. Here, $e_{t,0}$ is defined as y_t, and σ_0^2 is defined as γ_0, the variance of x_t. The factorization in Equation (6) is a scaled version of the Cholesky factorization of $\boldsymbol{\Sigma}$.

Substituting Equation (6) into Equation (2), it follows that

$$\begin{aligned} T_t^2 &= [\mathbf{X}_t - \boldsymbol{\mu}_0]' \left[\mathbf{B}^{-1'} \mathbf{D} \mathbf{B}^{-1}\right]^{-1} [\mathbf{X}_t - \boldsymbol{\mu}_0] \\ &= [\mathbf{X}_t - \boldsymbol{\mu}_0]' \left[\mathbf{B} \mathbf{D}^{-1} \mathbf{B}'\right] [\mathbf{X}_t - \boldsymbol{\mu}_0] \\ &= \sum_{n=0}^{p-1} \frac{e_{t-p+1+n,n}^2}{\sigma_n^2}. \end{aligned} \quad (8)$$

From Equation (8), the autoregressive T^2 control chart has a close relationship to residual-based control charts. The value T_t^2 is the sum of the squares of the residual errors for various order AR models, scaled by the residual variances (the residual for the n^{th} order model is delayed by $p - 1 - n$ observations). This includes the residual for the AR(0) model, which is just the original process y_t. Note that for this interpretation to be valid it is not necessary that y_t be an ARMA process. Similar T^2 decompositions for true multivariate processes are well known (see e.g., Hawkins (1993) and Mason, Tracy, and Young (1997)).

When the process is in-control, each of the p terms in the summation in Equation (8) follows a chi-square distribution with one degree-of-freedom. This follows from the fact that each $e_{i,n}$, being a linear combination of Gaussian random variables, is itself Gaussian. The square of a zero-mean Gaussian random variable, divided by its variance, is a chi-square random variable with one degree-of-freedom. Furthermore, we show in Appendix B that the covariance matrix of the $e_{i,n}$ terms in Equation (8) is the diagonal matrix in Equation (7). Hence, they are uncorrelated. Since they are Gaussian, they are also independent, and so are their squared values. Thus, all p terms in Equation (8) are chi-square random variables with one degree-of-freedom and are independent of each other. Thus, T_t^2 is chi-square distributed with p degrees-of-freedom.

The relationship to residual-based control charts becomes even more apparent when x_t is an AR(1) process $x_t - \mu = \phi_1(x_{t-1} - \mu) + a_t$. Since the optimal AR($n$) model for $n = 1, 2, \ldots, p-1$ is just the true AR(1) model, $\mathbf{B}$ and $\mathbf{D}$ in Equation (6) have the structure

$$\mathbf{B} = \begin{pmatrix} 1 & -\phi_1 & 0 & \cdots & 0 \\ & 1 & -\phi_1 & \cdots & \vdots \\ & & 1 & \vdots & 0 \\ & 0 & & \ddots & -\phi_1 \\ & & & & 1 \end{pmatrix}$$

and

$$\mathbf{D} = \begin{pmatrix} \sigma_x^2 & & & \\ & \sigma_a^2 & & \\ & & \ddots & \\ & & & \sigma_a^2 \end{pmatrix} \quad (9)$$

where $\sigma_x^2 = \gamma_0$ is the variance of x_t. Consequently, the decomposition in Equation (8) becomes

$$T_t^2 = \frac{(x_{t-p+1} - \mu_0)^2}{\sigma_x^2} + \sum_{i=0}^{p-2} \frac{e_{t-i}^2}{\sigma_a^2},$$

TABLE 5. Relationship Between the Optimal p (Denoted p^*) and the AR Model Order (Denoted n^*) Above Which the Parameter Magnitudes Drop Below 0.1. (The values for p^* are approximate, from Table 1.)

ϕ_1	0.98	0.90	0.90	0.90	0.90	0.50	0.50	0.50
θ_1	0.00	−0.90	−0.50	0.00	0.50	−0.90	0.00	−0.50
p^*	2	20	5	2	2−5	20	2	5
$n^* + 1$	2	18	5	2	4	18	2	5
$\beta_{1,19}$	0.98	1.80	1.40	0.90	0.40	1.40	0.50	1.00
$\beta_{2,19}$	0.00	−1.61	−0.70	0.00	0.20	−1.25	0.00	−0.50
$\beta_{3,19}$	0.00	1.44	0.35	0.00	0.10	1.12	0.00	0.25
$\beta_{4,19}$	0.00	−1.29	−0.17	0.00	0.05	−1.00	0.00	−0.12
$\beta_{5,19}$	0.00	1.15	0.09	0.00	0.03	0.89	0.00	0.06
$\beta_{6,19}$	0.00	−1.02	−0.04	0.00	0.01	−0.79	0.00	−0.03
$\beta_{7,19}$	0.00	0.90	0.02	0.00	0.01	0.70	0.00	0.02
$\beta_{8,19}$	0.00	−0.80	−0.01	0.00	0.00	−0.62	0.00	−0.01
$\beta_{9,19}$	0.00	0.70	0.01	0.00	0.00	0.55	0.00	0.00
$\beta_{10,19}$	0.00	−0.61	0.00	0.00	0.00	−0.48	0.00	0.00
$\beta_{11,19}$	0.00	0.53	0.00	0.00	0.00	0.42	0.00	0.00
$\beta_{12,19}$	0.00	−0.46	0.00	0.00	0.00	−0.36	0.00	0.00
$\beta_{13,19}$	0.00	0.39	0.00	0.00	0.00	0.30	0.00	0.00
$\beta_{14,19}$	0.00	−0.32	0.00	0.00	0.00	−0.25	0.00	0.00
$\beta_{15,19}$	0.00	0.26	0.00	0.00	0.00	0.20	0.00	0.00
$\beta_{16,19}$	0.00	−0.20	0.00	0.00	0.00	−0.16	0.00	0.00
$\beta_{17,19}$	0.00	0.14	0.00	0.00	0.00	0.11	0.00	0.00
$\beta_{18,19}$	0.00	−0.08	0.00	0.00	0.00	−0.07	0.00	0.00
$\beta_{19,19}$	0.00	0.03	0.00	0.00	0.00	0.03	0.00	0.00

where e_j is the residual error for the true model at observation number j and is given by $e_j = (x_j - \mu_0) - \phi_1(x_{j-1} - \mu_0)$. Consequently, when the process is AR(1), T_t^2 is a moving average of the squares of the residuals added to the square of $(x_{t-p+1} - \mu_0)/\sigma_x$.

Autoregressive T^2 Chart Design

To implement the autoregressive T^2 control chart, the user must specify p and α. We first discuss guidelines for selecting p. As discussed previously, the simulation results indicate that the value of p that provides the lowest out-of-control ARL does not depend strongly on the size of the mean shift. In other words, p can be selected based only on the autocovariance structure of the process (equivalently, on the ARMA parameters if x_t can be represented as an ARMA process).

Table 5 shows what were somewhat subjectively determined to be the "optimal" p values (from the four values of p (2, 5, 10, and 20) that were considered) for the eight different ARMA(1,1) models used in the simulation results of Tables 1 through 3. The optimal values for p are denoted by p^*. For some models, a single value of p was optimal for all mean shifts considered. For others, different values of p provided a lower out-of-control ARL, depending on the size of the mean shift. It was generally true, however, that for a given model a single value of p was close to optimal for all mean shifts considered. For example, for $\phi_1 = 0.5$ and $\theta_1 = -0.5$, the optimal values of p for mean shifts of size 2, 3, 4, and 5 were 20, 10, 5, and 5, respectively. For the smaller mean shifts of size 2 and 3, however, the ARL for $p = 5$ was only slightly higher than for $p = 10$ and 20. In this case, $p^* = 5$ was designated the "optimal" value.

For a given model, there appears to be a simple relationship between p^* and the model characteristics. Table 5 also shows the parameters, defined in Equation (5), of a high order AR(n) model (i.e., $n = 19$) fit to the autocovariance function of each ARMA model. More specifically, the AR(19) model parameters are the solution to the Yule-Walker equations of order 19 (Box et al., 1994). The parameters can be easily determined by finding $\mathbf{B}$ in the Cholesky factorization in Equation (6) with $p = 20$. The AR(19) parameters are then contained in the last column of $\mathbf{B}^{-1}$. If the model is invertible, the parameters of a high order AR model will decay to zero as the lag increases (Box et al., 1994). Let n^* be the lag after which the magnitude of the parameters drops below some small value, say 0.1. In other words, n^* is defined as the smallest integer such that $|\beta_{j,\infty}| < 0.1$ for all $j > n^*$. Thus, n^* is essentially the AR model order that is needed to "capture" the dynamics of the process. From the discussion on the autoregressive T^2 decomposition of the previous section, if p is set as $n^* + 1$, then $\boldsymbol{\Sigma}$ will also capture the dynamics of the process. Setting $p^* = n^* + 1$ appears to be an effective rule-of-thumb for selecting the optimal value of p. This follows from Table 5, which shows that for all eight models considered, p^* is close to $n^* + 1$. As discussed above, n^* can be found by forming $\boldsymbol{\Sigma}$ with a relatively large p (e.g., $p = 20$ or 30), taking its Cholesky factorization, and inspecting the last column of $\mathbf{B}^{-1}$. Note that if the process truly is AR(n), then $p^* = n + 1$. For example, $p^* = 2$ for each of the three AR(1) models.

In the preceding method for finding n^* we assume that the true process autocovariance function (or, equivalently, the true ARMA model describing the process) is accurately known, and then attempt to approximate the process characteristics using a sufficiently high-order AR model. In situations where the autocovariance or ARMA model is estimated from limited data and cannot be assumed to be accurate, one could still apply this method using the estimated ARMA model. In this situation, however, the following alternative method for finding n^* may be preferable. The alternative method is to fit various order AR models to the data, until a further increase in the model order (beyond n^*) is no longer statistically significant; p^* would again be set as $n^* + 1$.

After selecting p, one potential means of selecting α is to fix the false alarm probability. If the threshold for T_t^2 is set as the $1 - \alpha$ percentile of the chi-square distribution with p degrees-of-freedom, then clearly the false alarm probability will be α for any given isolated time. There is no simple relationship, however, between α and the in-control ARL. The sequence T_t^2 can have high autocorrelation, in particular if p is large, and the in-control ARL may be substantially larger than $1/\alpha$. This is evident from Table 1, where the in-control ARL (denoted by ARL_0) was 500 for all cases.

Empirically, we have observed that for a given ARMA model (i.e., a given autocovariance structure) and p, the relationship between ARL_0 and α is very close to log-linear. In other words,

$$\log(\mathrm{ARL}_0) \cong c_0 - c_1 \log(\alpha), \tag{10}$$

where c_0 and c_1 are constants that depend on p and the ARMA model parameters. Tables 6 and 7 give

TABLE 6. Values for c_0 for Various p and ARMA(1,1) Parameters (To Be Used in Equation (10))

		(ϕ, θ)							
		(0.98,0)	(0.9,−0.9)	(0.9,−0.5)	(0.9,0)	(0.9,0.5)	(0.5,−0.9)	(0.5,−0.5)	(0.5,0)
	2	1.069	0.864	0.833	0.709	0.706	0.624	0.572	0.521
	3	1.202	1.009	0.977	0.900	0.900	0.835	0.818	0.792
	4	1.316	1.133	1.093	1.061	1.052	0.996	1.029	0.998
	5	1.410	1.250	1.219	1.217	1.201	1.167	1.225	1.203
	6	1.505	1.385	1.349	1.357	1.323	1.309	1.371	1.343
	7	1.618	1.509	1.482	1.486	1.462	1.460	1.490	1.464
p	8	1.695	1.631	1.600	1.606	1.554	1.577	1.597	1.570
	9	1.774	1.738	1.726	1.723	1.659	1.702	1.695	1.673
	10	1.854	1.834	1.830	1.832	1.746	1.807	1.798	1.765
	12	1.983	1.975	1.968	1.978	1.901	1.962	1.928	1.900
	14	2.097	2.093	2.077	2.090	2.028	2.086	2.043	2.019
	16	2.198	2.212	2.170	2.191	2.140	2.189	2.149	2.136
	18	2.296	2.320	2.266	2.299	2.273	2.293	2.242	2.239
	20	2.381	2.407	2.341	2.377	2.364	2.378	2.325	2.323

TABLE 7. Values for c_1 for Various p and ARMA(1,1) Parameters (To Be Used in Equation (10))

		(ϕ, θ)							
		(0.98,0)	(0.9,−0.9)	(0.9,−0.5)	(0.9,0)	(0.9,0.5)	(0.5,−0.9)	(0.5,−0.5)	(0.5,0)
p	2	0.918	0.941	0.940	0.951	0.950	0.971	0.972	0.975
	3	0.913	0.938	0.938	0.944	0.944	0.964	0.955	0.957
	4	0.909	0.935	0.937	0.938	0.937	0.957	0.942	0.945
	5	0.905	0.931	0.934	0.930	0.931	0.948	0.929	0.933
	6	0.901	0.925	0.927	0.922	0.926	0.939	0.920	0.924
	7	0.897	0.916	0.917	0.912	0.920	0.926	0.913	0.917
	8	0.893	0.907	0.905	0.903	0.915	0.914	0.906	0.910
	9	0.889	0.897	0.892	0.893	0.909	0.903	0.899	0.906
	10	0.887	0.888	0.884	0.886	0.904	0.895	0.895	0.901
	12	0.883	0.879	0.881	0.879	0.896	0.887	0.890	0.895
	14	0.879	0.872	0.878	0.874	0.888	0.882	0.887	0.890
	16	0.876	0.867	0.877	0.870	0.881	0.878	0.884	0.887
	18	0.873	0.863	0.875	0.868	0.875	0.874	0.883	0.883
	20	0.870	0.860	0.874	0.865	0.869	0.871	0.881	0.880

the constants c_0 and c_1, respectively, for each of the ARMA models and for various values of p ranging from 2 to 20. The constants were identified empirically by running Monte Carlo simulations for the eight ARMA(1,1) models considered in Table 1, with ARL_0 ranging from 200 to 1000. For larger values of p, c_0 and c_1 are relatively insensitive to the ARMA parameters and depend predominantly on p. For example, for $p = 10$ and 20, the approximations

$$p = 10: \quad \log(\text{ARL}_0) \cong 1.813 - 0.892 \log(\alpha)$$
$$p = 20: \quad \log(\text{ARL}_0) \cong 2.364 - 0.871 \log(\alpha)$$

were always within 2.9% and 2.3%, respectively, of the true ARL_0. The suggested method for selecting the threshold for T_t^2 is to specify a desired ARL_0, determine the approximate c_0 and c_1 from Tables 6 and 7, determine α from Equation (10), and then set the threshold as the $1-\alpha$ percentile of the chi-square distribution with p degrees-of-freedom. We point out that the procedures for selecting the optimal p and the thresholds are based on a number of empirical approximations and limited simulation for a specific set of ARMA parameters. As such, they can only be expected to provide approximate in-control ARLs.

Robustness to Model Uncertainty and Unknown $\boldsymbol{\Sigma}$

Up to this point, only the case where $\boldsymbol{\Sigma}$ is assumed known has been considered. Essentially, this means that enough data are available so that an estimate, denoted by $\widehat{\boldsymbol{\Sigma}}$, is sufficiently close to $\boldsymbol{\Sigma}$. It is often the case that the available data is insufficient to warrant this assumption. In this event, T_t^2 of Equation (2) with $\boldsymbol{\Sigma}$ replaced by $\widehat{\boldsymbol{\Sigma}}$ does not follow a chi-square distribution. The primary problem with this is that the actual in-control ARL may then be substantially shorter than what is desired, resulting in too many false alarms. This is no less a problem in residual-based control charts, and is analogous to having insufficient data to accurately estimate the ARMA model parameters, the model order, and σ_a^2. If the estimated ARMA model is inaccurate, then the residuals will no longer be uncorrelated, their variance will differ from $\widehat{\sigma}_a^2$, and the incontrol ARL of (for example) a CUSUM on the residuals will not be as desired. Adams and Tseng (1998), Lu and Reynolds (1999), and Apley and Shi (1999) have demonstrated the effects of model estimation errors on the ARLs for various residual-based charts. Adams and Tseng (1998) provide guidelines for determining the sample size (from which the model parameters are estimated) such that the effects are negligible.

Although the adverse effects of model estimation errors on the ARLs of residual-based control charts are known, no approaches for modifying the control limits to account for the model uncertainty have been proposed. A natural extension of the autoregressive T^2 control chart, when $\boldsymbol{\Sigma}$ must be estimated from limited data, provides one means of accommodating model uncertainty.

Suppose N observations $\{x_1, x_2, \ldots, x_N\}$ from Phase I operation (Sullivan and Woodall (1996)) are available, from which $\boldsymbol{\Sigma}$ is to be estimated. Consider, as an estimate of $\boldsymbol{\Sigma}$,

$$\widehat{\boldsymbol{\Sigma}} = \frac{1}{N-p+1} \sum_{t=p}^{N} [\mathbf{X}_t - \widehat{\boldsymbol{\mu}}_0] [\mathbf{X}_t - \widehat{\boldsymbol{\mu}}_0]',$$

where $\widehat{\boldsymbol{\mu}}_0 = [\widehat{\mu}_0 \ \widehat{\mu}_0 \ldots \widehat{\mu}_0]'$, with $\widehat{\mu}_0$ the usual sample average of $\{x_1, x_2, \ldots, x_N\}$. The suitability of the above estimate of $\boldsymbol{\Sigma}$ requires that the process is in-control when the Phase I data are collected. When this is not the case, some form of rational subgrouping is recommended when estimating $\boldsymbol{\Sigma}$. Sullivan and Woodall (1996), to which the reader is referred for details, provides an excellent discussion of strategies for more robustly estimating $\boldsymbol{\Sigma}$ during Phase I operation. In related work, Boyles (2000) discusses estimating ARMA model parameters when assignable causes are present during Phase I.

Assuming that our estimate of $\boldsymbol{\Sigma}$ is used and is suitable, the T^2 statistic would be

$$T_t^2 = [\mathbf{X}_t - \widehat{\boldsymbol{\mu}}_0]' \, \widehat{\boldsymbol{\Sigma}}^{-1} [\mathbf{X}_t - \widehat{\boldsymbol{\mu}}_0]. \tag{11}$$

The distribution of this statistic is better approximated as an F-distribution (scaled by an appropriate constant), which takes into account the uncertainty in $\widehat{\boldsymbol{\Sigma}}$, than as a chi-square distribution. The F-distribution is still an approximation, however, since $\{\mathbf{X}_p, \mathbf{X}_{p+1}, \ldots, \mathbf{X}_N\}$ are not independent of each other, and $\widehat{\boldsymbol{\Sigma}}$ will not exactly follow a Wishart distribution. Approximating by an F-distribution, the threshold for T_t^2 would be

$$\frac{(N-p+1)p}{N-2p+2} F(p, N-2p+2, 1-\alpha), \tag{12}$$

where $F(j,k,\nu)$ denotes the ν percentile of the F-distribution with j numerator degrees-of-freedom and k denominator degrees-of-freedom. As N gets larger, the expression in Equation (12) approaches the $1-\alpha$ percentile of the chi-square distribution with p degrees-of-freedom. This is reasonable, since as N gets large, $\widehat{\boldsymbol{\Sigma}}$ approaches $\boldsymbol{\Sigma}$. As N becomes smaller, the value in Equation (12) increases monotonically, and the threshold for T_t^2 gets larger. This serves to prevent an in-control ARL that is substantially shorter than desired.

There exists an analogous decomposition of Equation (11) for the case that $\boldsymbol{\Sigma}$ is estimated, similar to the decomposition in Equation (8) that was derived for $\boldsymbol{\Sigma}$ known. The form of the decomposition is identical to Equation (8), except that the $e_{j,n}$ terms are the residual errors for AR(n) models whose coefficients (the $\beta_{j,n}$'s from Equation (5)) are *estimated* from the Phase I data using linear least squares. Likewise, each σ_n^2 term is the sample variance of the AR(n) model residuals over the Phase I data from which the parameters are estimated. The proof is a straightforward extension of that in Appendix B, where the correlation inner-product associated with minimum mean square error predictors is replaced by the sample correlation inner-product associated with linear least squares regression. Mason et al. (1997) and Mason and Young (1999) present similar decompositions for multivariate processes.

Table 8 shows Monte Carlo simulation results that demonstrate how control limits based on Equation (12) help in preventing an undesirably short in-control ARL. The actual in-control ARL for various residual-based CUSUM and autoregressive T^2 charts are shown when the necessary quantities (the ARMA model for the CUSUM chart and $\boldsymbol{\Sigma}$ for the T^2 chart) are estimated from limited data with $N = 50$ and $N = 100$. For each case, 20,000 Monte Carlo trials were run. For each trial, N observations were generated from the true ARMA model (with parameters given in Table 8), the model was estimated from these data, and the test was implemented on

TABLE 8. True ARL_0 for Residual-Based CUSUM and T^2 Charts When the Model is Estimated from $N = 50$ and $N = 100$ Observations. (The desired ARL_0 in all cases is 500.)

true parameters			CUSUM ARL_0		T^2 ARL_0		
ϕ_1	θ_1	N	$\text{K} = 0.2\widehat{\sigma}_a$	$\text{K} = 0.5\widehat{\sigma}_a$	$p = 2$	$p = 5$	$p = 10$
0.9	0	50	340	265	499	585	701
0.9	0	100	400	338	467	472	454
0.5	−0.5	50	372	275	707	704	918
0.5	−0.5	100	402	352	578	595	555

additional data generated from the true model until the test signaled a false alarm. For the T^2 charts, only $\boldsymbol{\Sigma}$ was estimated. For the residual-based CUSUMs, the model order, the parameters, and σ_a^2 were estimated. For simplicity, μ_0 was assumed to be zero without loss of generality. To estimate the model parameters, the Matlab "ARMAX" routine, which is based on nonlinear least squares, was used. Various ARMA(n, m) models were fit to the data, and Akaike's Final Prediction Error criterion (Ljung (1987)) was used to estimate model order. The set of models that were fitted was AR(1), AR(2), AR(3), AR(4), ARMA(1,1), ARMA(2,1), ARMA(3,2), and ARMA(4,3). In all cases, the tests were designed with a desired ARL_0 of 500 in mind. The CUSUMs were designed as if the estimated model equaled the true model. For reference values of $\text{K} = 0.2\widehat{\sigma}_a$ and $0.5\widehat{\sigma}_a$, the decision thresholds were $\text{H} = 10.0\widehat{\sigma}_a$ and $5.07\widehat{\sigma}_a$, respectively. For the T^2 charts with $p = 2$, 5, and 10, the values of α in Table 1 were used. The thresholds for T_t^2 were then set according to Equation (12).

For $N = 50$, the CUSUM ARL_0's were substantially lower than 500, in particular when $\text{K} = 0.5$. Even for $N = 100$, the CUSUM ARL_0's were much lower than 500. Although the T^2 ARL_0's were at times substantially larger than 500, they were never substantially lower. Consequently, when there is model uncertainty due to limited data, the autoregressive T^2 chart with threshold given by Equation (12) can be considered more robust with respect to excessive false alarms. Use of Equation (12) should be viewed only as an approximate means of widening the control limits to account for model uncertainty, since the assumed F-distribution of (11) is only an approximation.

Illustrative Example

Consider the chemical process described in Montgomery and Mastrangelo (1991), which is represented as an AR(1) process with $\phi_1 = 0.847$. For simplicity, assume that $\sigma_a^2 = 1.0$, and that the data is centered so that the in-control mean is zero. Suppose one wishes to apply an autoregressive T^2 chart with $\text{ARL}_0 = 500$. We assume the model is accurate enough that the χ^2 control limit can be used. Two design parameters, p and α, must be selected.

Applying Equation (4) with $\theta_1 = 0$ gives

$$\gamma_k = \frac{\sigma_a^2}{1 - \phi_1^2}\phi_1^k.$$

For arbitrary p, the covariance matrix would therefore be of the form

$$\boldsymbol{\Sigma} = \frac{\sigma_a^2}{1 - \phi_1^2}\begin{pmatrix} 1 & \phi_1 & \phi_1^2 & \cdots & \phi_1^{p-1} \\ \phi_1 & 1 & \phi_1 & \cdots & \phi_1^{p-2} \\ \phi_1^2 & \phi_1 & 1 & \cdots & \phi_1^{p-3} \\ \vdots & \vdots & & \ddots & \vdots \\ \phi_1^{p-1} & \phi_1^{p-2} & \phi_1^{p-3} & \cdots & 1 \end{pmatrix}$$

with $\phi_1 = 0.847$. Taking the Cholesky factor of $\boldsymbol{\Sigma}$ or using Equation (9) directly gives

$$\mathbf{B} = \begin{pmatrix} 1 & -0.847 & 0 & \cdots & 0 \\ & 1 & -0.847 & & \vdots \\ & & 1 & \ddots & 0 \\ & 0 & & \ddots & -0.847 \\ & & & & 1 \end{pmatrix}.$$

By inspection of the last column of $\mathbf{B}$, it is clear that $\beta_{j,\infty} = 0$ for $j > n^* = 1$. Using the guidelines presented earlier, $p = n^* + 1 = 2$ is recommended. The covariance matrix for $p = 2$ reduces to

$$\boldsymbol{\Sigma} = \frac{\sigma_a^2}{1 - \phi_1^2}\begin{pmatrix} 1 & \phi_1 \\ \phi_1 & 1 \end{pmatrix} = \begin{pmatrix} 3.54 & 3.00 \\ 3.00 & 3.54 \end{pmatrix}.$$

From Tables 6 and 7, $c_0 = 0.709$ and $c_1 = 0.951$ for a similar AR(1) model with $\phi_1 = 0.9$. Using these in Equation (10) gives $\alpha = 0.0031$ for an approximate ARL_0 of 500. The test threshold is $\chi^2(1 - \alpha, p) = 11.55$.

An autoregressive T^2 chart with $p = 2$ and $\alpha = 0.0031$ was applied to the simulated AR(1) process,

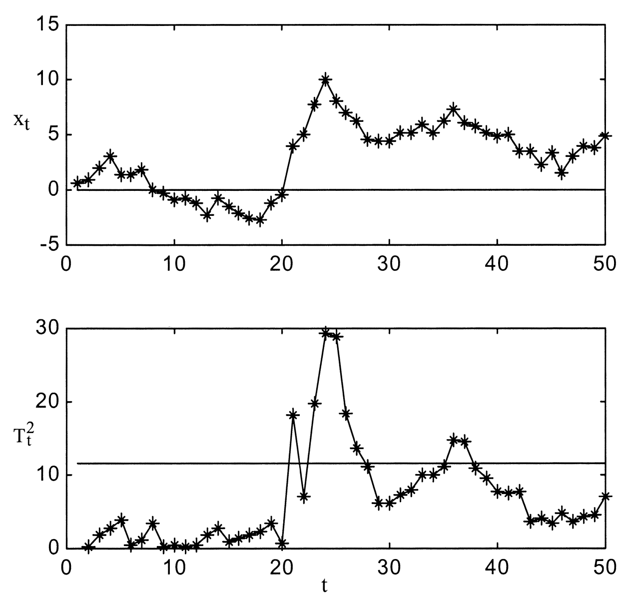

FIGURE 1. Autoregressive T^2 Chart Example.

where a mean shift of magnitude 4.5 was introduced at observation number 21. Since the process standard deviation is $\sigma_x = \gamma_0^{1/2} = 1.88$, the mean shift is approximately 2.4 standard deviation units. The simulated process and the T^2 statistics are shown in Figure 1. In this case, the T^2 chart signaled on the first observation following the mean shift.

Conclusions

A method for monitoring autocorrelated processes, termed the autoregressive T^2 chart, has been presented in this paper. The terminology results from the fact that the T^2 statistic can be decomposed into the sum of the squares of the residual errors for various order autoregressive models fit to the process. Hence, there is a close relationship between residual-based control charts and the autoregressive T^2 chart.

The performance of the autoregressive T^2 chart, in terms of the ARL, has been compared to residual-based CUSUM and Shewhart charts for a number of ARMA(1,1) processes. For certain ranges of the ARMA parameters, the autoregressive T^2 chart performs substantially better than either the CUSUM or Shewhart charts. In general, for moderate to large mean shifts, the autoregressive T^2 chart is superior for most of the processes considered. For small mean shifts the CUSUM chart is often superior, depending on the process parameters. Even in situations when the CUSUM has a lower out-of-control ARL, the autoregressive T^2 chart usually has a much higher probability of detecting the mean shift within the first few observations following the shift, and is similar to a residual-based Shewhart chart in this respect. One primary advantage of the autoregressive T^2 chart over a residual-based CUSUM is that a single T^2 chart design is often nearly optimal for a wide range of mean shift sizes. In contrast, the optimal CUSUM design almost always depends on the mean shift size of interest. A CUSUM chart that is optimal for small mean shifts may perform poorly for large mean shifts, and vice-versa. Guidelines for designing the autoregressive T^2 chart have been provided.

An additional advantage is that the autoregressive T^2 chart provides some robustness with respect to an excessive number of false alarms when there is large uncertainty in the process model. Specifically, when there are limited data available for estimating an ARMA process model, a residual-based CUSUM may have substantially lower in-control ARL than what is intended. The autoregressive T^2 chart possesses a natural mechanism for taking into account model uncertainty due to limited data. In the presence of model uncertainty the actual in-control ARL for the autoregressive T^2 chart is never substantially lower than the intended ARL in the examples considered. In some cases it is, however, substantially higher. There may exist better approximations to the distribution of T_t^2 in Equation (11) that will allow more accurate specification of the in-control ARL when data are limited. Although SPC for autocorrelated processes and residual-based control charts have been widely studied, how to modify the control chart to account for model uncertainty has not been studied and deserves more attention.

In this paper we have investigated the autoregressive T^2 chart performance only for the case of process mean shifts. Typically, it is desirable to detect variance changes as well. It is likely that the autoregressive T^2 chart would also be effective at detecting variance increases. This speculation follows from the decomposition of T_t^2 as the sum of the *squares* of the autoregressive residual errors.

Appendix A

This appendix provides an algorithm for calculating $\boldsymbol{\Sigma}$, when y_t follows the ARMA(n,m) model of Equation (3). Let g_j, $j = 0, 1, 2, \ldots$, denote the Green's function for the ARMA model. Green's function is essentially the output of the ARMA model when the input sequence a_t is a single pulse of unit magnitude at time $t = 0$ (see, for example, Pandit and Wu (1990)). It can be calculated recursively via

$$g_j = \begin{cases} \sum_{i=1}^{n} \phi_i g_{j-i} - \theta_j & 1 \le j \le m \\ \sum_{i=1}^{n} \phi_i g_{j-i} & m < j \end{cases},$$

with initial conditions $g_0 = 1$ and $g_j = 0$ for $j < 0$. It can be shown (Pandit and Wu, (1990)) that the autocovariance function of x_t is given by

$$\gamma_k = \sigma_a^2 \sum_{j=0}^{\infty} g_j g_{j+k}.$$

Since g_j decays exponentially for stable ARMA models, the infinite summation can be truncated.

Appendix B

This appendix provides a simple proof of Equation (7). Define $\mathbf{Y}_t = [y_{t-p+1}\ y_{t-p+2}\ \ldots\ y_t]'$ and $\mathbf{E}_t = [e_{t-p+1,0}\ e_{t-p+2,1}\ \ldots\ e_{t,p-1}]'$ with $e_{i,j}$ as in Equation

(5). By definition of $\mathbf{B}$, $\mathbf{E}_t$, and $\mathbf{Y}_t$, it follows that $\mathbf{E}_t = \mathbf{B}'\mathbf{Y}_t$, or

$$\mathbf{Y}_t = \mathbf{B}^{-1'}\mathbf{E}_t. \qquad \text{(B1)}$$

Let $\mathbf{E}_{t,j}$ and $\mathbf{Y}_{t,j}$ represent the j^{th} elements of $\mathbf{E}_t$ and $\mathbf{Y}_t$, respectively. Then, $\mathbf{E}_{t,j}$ is the residual error in the minimum mean square error prediction of $\mathbf{Y}_{t,j}$, as a linear function of $\mathbf{Y}_{t,1}, \mathbf{Y}_{t,2}, \ldots, \mathbf{Y}_{t,j-1}$. Thus, $\mathbf{E}_{t,j}$ is orthogonal to (i.e., uncorrelated with) $\mathbf{Y}_{t,1}, \mathbf{Y}_{t,2}, \ldots, \mathbf{Y}_{t,j-1}$. Obviously, $\mathbf{E}_{t,j}$ must also be orthogonal to $\mathbf{E}_{t,1}, \mathbf{E}_{t,2}, \ldots, \mathbf{E}_{t,j-1}$, since $\mathbf{E}_{t,i}$ $(i = 1, 2, \ldots, p)$ is a linear combination of $\mathbf{Y}_{t,1}, \mathbf{Y}_{t,2}, \ldots, \mathbf{Y}_{t,i}$. Thus, the covariance matrix of $\mathbf{E}_t$ is the diagonal matrix $\mathbf{D}$ in Equation (7).

From Equation (B1), the covariance matrix of $\mathbf{X}_t$ is

$$\begin{aligned} \boldsymbol{\Sigma} &= E\left[(\mathbf{X}_t - \boldsymbol{\mu}_0)(\mathbf{X}_t - \boldsymbol{\mu}_0)'\right] \\ &= E\left[\mathbf{Y}_t\mathbf{Y}_t'\right] \\ &= \mathbf{B}^{-1'} E\left[\mathbf{E}_t\mathbf{E}_t'\right]\mathbf{B}^{-1} \\ &= \mathbf{B}^{-1'}\mathbf{D}\mathbf{B}^{-1}. \end{aligned}$$

Acknowledgments

D. W. Apley was partially supported in this research by the State of Texas Advanced Technology Program under grant 000512-0289-1999. F. Tsung was supported by RGC Competitive Earmarked Research Grant HKUST6073/00E.

References

Adams, B. M. and Tseng, I. T. (1998). "Robustness of Forecast-Based Monitoring Schemes". *Journal of Quality Technology* 30, pp. 328–339.

Alwan, L. C. and Roberts, H. V. (1988). "Time-Series Modeling for Statistical Process Control". *Journal of Business and Economic Statistics* 6, pp. 87–95.

Alwan, A. J. and Alwan, L. C. (1994). "Monitoring Autocorrelated Processes Using Multivariate Quality Control Charts". *Proceedings of the Decision Sciences Institute Annual Meeting* 3, pp. 2106–2108.

Apley, D. W. and Shi, J. (1999). "The GLRT for Statistical Process Control of Autocorrelated Processes". *IIE Transactions* 31, pp. 1123–1134.

Boyles, R. A. (2000). "Phase I Analysis for Autocorrelated Processes". *Journal of Quality Technology* 32, pp. 395–409.

Box, G. E. P.; Jenkins, G. M.; and Reinsel, G. C. (1994). *Time Series Analysis: Forecasting and Control,* 3rd ed. Prentice Hall, Englewood Cliffs, N.J.

Berthouex, P. M.; Hunter, W. G.; and Pallesen, L. (1978). "Monitoring Sewage Treatment Plants: Some Quality Control Aspects". *Journal of Quality Technology* 10, pp. 139–149.

Chow, M. C.; Wu, S. M.; and Ermer, D. S. (1979). "A Time Series Control Chart for a Nuclear Reactor". *Proceedings of the 1979 Annual Reliability and Maintainability Symposium,* pp. 1–7.

English, J. R.; Krishnamurthi, M.; and Sastri, T. (1991). "Quality Monitoring of Continuous Flow Processes". *Computers and Industrial Engineering* 20, pp. 251–260.

Hawkins, D. M. (1993). "Regression Adjustment for Variables in Multivariate Quality Control". *Journal of Quality Technology* 25, pp. 170–182.

Johnson, R. A. and Bagshaw, M. (1974). "The Effect of Serial Correlation on the Performance of CUSUM tests". *Technometrics* 16, pp. 103–112.

Krieger, C. A.; Champ, C. W.; and Alwan, L. C. (1992). "Monitoring an Autocorrelated Process". Presented at the Pittsburgh Conference on Modeling and Simulation, Pittsburgh, PA.

Lin, W. S. and Adams, B. M. (1996). "Combined Control Charts for Forecast-Based Monitoring Schemes". *Journal of Quality Technology* 28, pp. 289–301.

Ljung, L. (1987). *System Identification: Theory for the User*. Prentice-Hall, Englewood Cliffs, NJ.

Lu, C. W. and Reynolds, M. R., Jr. (1999). "EWMA Control Charts for Monitoring the Mean of Autocorrelated Processes". *Journal of Quality Technology* 31, pp. 166–188.

Lu, C. W. and Reynolds, M. R., Jr. (2001). "CUSUM Charts for Monitoring an Autocorrelated Process". *Journal of Quality Technology* 33, pp. 316-334.

Mason, R. L.; Tracy, N. D.; and Young, J. C. (1997). "A Practical Approach for Interpreting Multivariate T^2 Control Chart Signals". *Journal of Quality Technology* 29, pp. 396–406.

Mason, R. L. and Young, J. C. (1999). "Improving the Sensitivity of the T^2 Statistic in Multivariate Process Control". *Journal of Quality Technology* 31, pp. 155–165.

Montgomery, D. C. (2001). *Introduction to Statistical Quality Control,* 4th ed. John Wiley and Sons, New York, NY.

Montgomery, D. C. and Mastrangelo, C. M. (1991). "Some Statistical Process Control Methods for Autocorrelated Data". *Journal of Quality Technology* 23, pp. 179–193.

Pandit, S. M. and Wu, S. M. (1990). *Time Series and System Analysis with Applications*. Krieger Publishing Company, Malabar, FL.

Runger G. C.; Willemain, T. R.; and Prabhu, S. (1995). "Average Run Lengths for Cusum Control Charts Applied to Residuals". *Communications in Statistics: Theory and Methods* 24, pp. 273–282.

Sullivan, J. H. and Woodall, W. H. (1996). "A Comparison of Multivariate Control Charts for Individual Observations". *Journal of Quality Technology* 28, pp. 398–408.

Superville, C. R. and Adams, B. M. (1994). "An Evaluation of Forecast-Based Quality Control Schemes". *Communications in Statistics: Simulation and Computation* 23, pp. 645–661.

Vander Wiel, S. A. (1996). "Monitoring Processes That Wander Using Integrated Moving Average Models". *Technometrics* 38, pp. 139–151.

Vasilopoulos, A. V. and Stamboulis, A. P. (1978). "Modification of Control Chart Limits in the Presence of Data Correlation". *Journal of Quality Technology* 10, pp. 20–30.

Wardell, D. G.; Moskowitz, H.; and Plante, R. D. (1994). "Run- Length Distributions of Special-Cause Control Charts for Correlated Processes". *Technometrics* 36, pp. 3–17.

Key Words: *Autocorrelation, Hotelling's T^2 Statistic, Statistical Process Control, Time Series Analysis*.