



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/99097/>

Version: Accepted Version

Article:

Triantafyllopoulos, K. and Bersimis, S. (2016) Phase II control charts for autocorrelated processes. *Quality Technology and Quantitative Management*, 13 (1). pp. 88-108. ISSN: 1684-3703

<https://doi.org/10.1080/16843703.2016.1139844>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Phase II control charts for autocorrelated processes

K. Triantafyllopoulos and S. Bersimis

School of Mathematics and Statistics,

Hicks Building, University of Sheffield, S3 7RH, Sheffield, UK

email: kostas@sheffield.ac.uk, Tel: +44 114 2 3741

and Department of Statistics and Insurance Science,

University of Piraeus,

80, Karaoli and Dimitriou Street, 185 34 Piraeus, Greece,

email: sbersim@unipi.gr, Tel: +30 210 4142452

September 23, 2014

Abstract

A large amount of SPC procedures are based on the assumption that the process subject to monitoring consists of independent observations. Chemical processes as well as many non-industrial processes exhibit autocorrelation, for which the above mentioned control procedures are not suitable. This paper proposes a Phase II control procedure for autocorrelated and possibly locally stationary processes. A time-varying autoregressive (AR) model is proposed, which is capable of dealing with the autocorrelation as well as with local non-stationarities of the temporal process. Such non-stationarities are induced by the time-varying nature of the AR coefficients. The model is optimised during Phase I when it is assured that the process is in control and as a result the model describes accurately the process. The Phase II proposed control procedure is based on a comparison of the current time series model with an alternative model, measuring deviations from it. This comparison is carried out using Bayes factors, which help to establish the in-control or out-of-control state of the process in Phase II. Using the threshold rules of the Bayes factors we propose a binomial-type control procedure for the monitoring of the process. The methodology of this paper is illustrated using two data sets consisting of temperature

measurements at two different stages in the manufacturing of a plastic mould.

Some key words: autocorrelated processes, time series monitoring, binomial-type control procedures, statistical process control, time-varying autoregressive model, Phase II control charts.

1 Introduction

Statistical process control (SPC) was initially used to monitor a process, consisting of independent observations, collected usually over time. The first applications were motivated by manufacturing and industry more generally, where the objective was to use SPC in order to enable the final product to conform to predefined standards and to reduce process variability. The control chart, first suggested by Shewhart, became a standard tool in SPC and its various versions and extensions highlight perhaps its central role in SPC.

Nowadays, the wide availability of plethora of data sets for various processes, including non-industrial processes, has resulted in the requirement of constructing control charts for autocorrelated data (Woodall and Montgomery, 1999). On one hand processes sampled in high frequency induces autocorrelation and on the other hand some processes such as chemical processes and environmental processes inherently introduce autocorrelation when they are sampled. The effect of autocorrelation in standard control charts has been studied by various authors. Considering cumulative sum (CUSUM) charts, Johnson and Bagshaw (1974) and Bagshaw and Johnson (1975) conclude that even small levels of autocorrelation, generated by simple autoregressive and moving average time series models, may well fail to detect out of control signals, if a conventional CUSUM chart is used. As many authors have pointed out (VanBrackle and Reynolds, 1997; Schmid, 1997; Psarakis and Papaleonida, 2007), this has the implication that the control limits need to be tightened in order to take into account autocorrelation. Alwan and Roberts (1988) presented a control procedure for the residuals in order to treat the effect of correlation. This control procedure monitors the residuals of an adequate model that describes the data. Several other studies consider monitoring the residuals after fitting a time series model to processes that exhibit autocorrelation and may be stationary or non-stationary (Hawkins, 1991; Montgomery and Mastrangelo, 1991; Kim et

al., 2012).

The following areas indicate the wide interest in monitoring autocorrelated processes: continuous product manufacturing operations (Psarakis and Papaleonida, 2007; Box et al., 2009; Zhang and Pollard, 1994; and Box et al., 2009), finance (Frisen, 2007), public-health monitoring and outbreak detection (Woodall, 2006; Shmueli and Burkom, 2010; Schioler and Frishen, 2012), environmental monitoring (Lund and Seymour, 1999; Pan and Chen, 2008), network monitoring and intrusion detection (Ye and Chen, 2001), education and E-learning (Karvounidis et al., 2012).

Considering autocorrelated processes (either industrial or non-industrial) a key difficulty for devising control charts is the definition of the in-control state; Alwan and Roberts (1988) discuss this difficulty. Due to autocorrelation and time-series dynamics such processes may exhibit certain non-standard variation over time even in the in-control state. As a result the overall mean of the process is not a good indicator of the stability of the process. Many of the above mentioned processes exhibit local stationarity or stability, which in short implies that from one point of time to another the process may change just slightly, but overall the process changes quite significantly. This suggests that control should not be based on an overall mean, but instead it should focus on the time-to-time conformance to a given model, hence the in-control state should suggest agreement with such a given model, while out-of-control state should favour deviations or shifts from this. This naturally suggests fitting a *null* model, using Alwan and Roberts (1988) way of thinking, and forming an alternative or model of deviations from the null; hence defining in-control and out-of-control states according to whether support is cast on the null or the alternative model.

In this paper we set-up such an approach for Phase II control, suggesting as null model a time-varying autoregressive (AR) time series model. The process is assured to be in-control in Phase I during which the time series model is optimised, its parameters are estimated and calibrated, hence this model describes the process well. The autoregression is able to model many time series data and the time-varying AR coefficients are responsible for driving possible local non-stationarities. The model is then casted in convenient state-space form and Bayesian inference is employed in order to fit the model to the data. The goodness of fit is assessed by

residual analysis in Phase I. As mentioned above we consider as in-control state the process when the time series model fit is adequate, or when the residuals follow independently a standard normal distribution. An alternative model is set, which quantifies possible deviations in mean and variance of this distribution. Then the proposed control procedure is based on a comparison at each time of these two models. Such a model comparison is facilitated (a) by using Bayes factors to create an index of similarity or distance of the two models and (b) by developing a binomial-type control procedure, based on the empirical guidelines and interpretation rules of the Bayes factors that have been used in the long-run in various fields of statistics (Kass and Raftery, 1995; Robert, 2007).

The remainder of the paper is organised as follows. Section 2 describes the time series model and discusses Bayesian inference and forecasting. The next section describes Bayes factors and proposes the control strategy and control rules developed in the paper. In Section 4 two processes consisting temperature measurements in two stages in the production of a plastic mould are used to illustrate the methodology. Conclusions are provided in the following section, while the appendix details technical details of the computation of the average run length (ARL) of the proposed control procedure.

2 Time-varying autoregressive model

2.1 Model set-up

Suppose that observations y_1, y_2, \dots, y_n are collected over time and that $\{y_t\}$ forms a univariate time series, so that y_i and y_j are serially correlated, for $i \neq j$; in other words $\{y_t\}$ exhibits autocorrelation. The principle aim of this section is to propose a mechanism for describing the time series, which will later enable us to detect an out of control state and the general evaluation of the process $\{y_t\}$.

A general purpose approach to the control of autocorrelated processes is to consider a time series model that describes well the process and then to devise some control procedure, usually on the residuals of the fitted model. Following this approach we consider the autoregressive

model of order p with time-varying parameters, defined as

$$y_t = \phi_{1t}y_{t-1} + \phi_{2t}y_{t-2} + \cdots + \phi_{pt}y_{t-p} + \varepsilon_t, \quad t > p, \quad (1)$$

where the innovations sequence $\{\varepsilon_t\}$ is a white noise process, i.e. $\{\varepsilon_t\}$ is an i.i.d. series of random variables with zero mean and some variance σ^2 . In this paper we assume that ε_t follows a normal distribution, for all t , so that $\varepsilon_t \sim N(0, \sigma^2)$. The parameters of this model are the time-varying autoregressive (AR) coefficients $\phi_{1t}, \phi_{2t}, \dots, \phi_{pt}$ together with the variance of the innovations σ^2 . Each of ϕ_{it} is assumed to follow a random walk process, i.e.

$$\phi_{it} = \phi_{it-1} + \nu_{it}, \quad i = 1, 2, \dots, p, \quad (2)$$

where ν_{it} is a white noise process. The random walk of ϕ_{it} is motivated by the assertion that locally at time t , $\phi_{it} \approx \phi_{it-1}$, i.e. that locally the AR coefficients are assumed to slowly vary, if the variance of ν_{it} is small; in other words locally the AR coefficients are expected to be time-invariant, i.e. $E(\phi_{it}) = E(\phi_{it-1})$, where $E(\cdot)$ denotes expectation. The transition equation (2) allows the above local expectation constancy from $t-1$ to t , but with increased variance, i.e. $\text{Var}(\phi_{it}) \geq \text{Var}(\phi_{it-1})$.

The above model can be put in state-space form by defining

$$y_t = \mathbf{x}_t^\top \boldsymbol{\phi}_t + \varepsilon_t \quad \text{and} \quad \boldsymbol{\phi}_t = \boldsymbol{\phi}_{t-1} + \boldsymbol{\nu}_t, \quad (3)$$

where $\mathbf{x}_t = [y_{t-1}, y_{t-2}, \dots, y_{t-p}]^\top$ is a design vector, the state vector is composed of the AR coefficients $\boldsymbol{\phi}_t = [\phi_{1t}, \phi_{2t}, \dots, \phi_{pt}]^\top$ and the sequence of random vectors $\{\boldsymbol{\nu}_t\}$ is assumed to be a p -dimensional Gaussian white noise, with some covariance matrix \mathbf{V}_t . It is assumed that at time $t = p$ *a priori* $\boldsymbol{\phi}_p$ is independent of ε_t and $\boldsymbol{\nu}_t$, for any $t > p$ and $\boldsymbol{\phi}_p$ follows a Gaussian prior, while $1/\sigma^2$ (the precision of ε_t) follows a gamma prior $\phi_p \sim N_p(\mathbf{m}_p, \sigma^2 \mathbf{C}_p)$ and $1/\sigma^2 \sim G(n_p/2, n_p S_p/2)$, for some mean vector \mathbf{m}_p , some covariance matrix $\sigma^2 \mathbf{C}_p$ and positive parameters n_p, S_p .

It is worth pointing out that if we force $\mathbf{V}_t = \mathbf{0}$ (the $p \times p$ zero matrix), then the model is reduced to an autoregressive model with time-invariant AR coefficients, since $\phi_{it} = \phi_{i,t-1}$ with probability 1, for all t and $i = 1, 2, \dots, p$. The model, however, benefits by setting a non-zero \mathbf{V}_t so that $\boldsymbol{\phi}_t$ may be slowly varying allowing AR coefficients being locally appropriate and

changing reflecting on the dynamics of $\{y_t\}$. The specification of \mathbf{V}_t may be facilitated by the use of a discount factor δ ($0 < \delta \leq 1$), so that $\text{Var}(\boldsymbol{\phi}_t) = \delta^{-1}\text{Var}(\boldsymbol{\phi}_{t-1}) \geq \text{Var}(\boldsymbol{\phi}_{t-1})$. If δ is close to 0, \mathbf{V}_t is large and this results in erratic shocks in $\boldsymbol{\phi}_t$. Both of these two extremes are to be avoided; typical values of δ are in the range of 0.7 to 0.99 (West and Harrison, 1997, Chapter 6).

2.2 Bayesian inference

Based on the observed data y_1, y_2, \dots, y_n , with information at time t : $y_{1:t} = [y_1, y_2, \dots, y_t]^\top$, a Bayesian version of the Kalman filter may be applied. At time t , the posterior distribution of $\boldsymbol{\phi}_t$, given σ^2 and $y_{1:t}$ is $\boldsymbol{\phi}_t \mid \sigma^2, y_{1:t} \sim N_p(\mathbf{m}_t, \sigma^2 \mathbf{C}_t)$ where \mathbf{m}_t and \mathbf{C}_t are updated by the recurrent relationships

$$\begin{aligned} \mathbf{m}_t &= \mathbf{m}_{t-1} + \mathbf{k}_t e_t^* \quad \text{and} \quad \mathbf{C}_t = \delta^{-1} \mathbf{C}_{t-1} - \mathbf{k}_t \mathbf{k}_t^\top q_t, \\ e_t^* &= y_t - \mathbf{x}_t^\top \mathbf{m}_{t-1}, \quad q_t = \delta^{-1} \mathbf{x}_t^\top \mathbf{C}_{t-1} \mathbf{x}_t + 1, \quad \mathbf{k}_t = \frac{\mathbf{C}_{t-1} \mathbf{x}_t}{\mathbf{x}_t^\top \mathbf{C}_{t-1} \mathbf{x}_t + \delta}. \end{aligned} \tag{4}$$

\mathbf{k}_t is the Kalman gain and e_t^* is the one-step forecast error.

With the gamma prior of $1/\sigma^2$, the posterior distribution of $1/\sigma^2$ is the gamma distribution $1/\sigma^2 \mid y_{1:t} \sim G(n_t/2, n_t S_t/2)$, with parameters $n_t/2$ and $n_t S_t/2$, where where $n_t = n_{t-1} + 1$ and $n_t S_t = n_{t-1} S_{t-1} + e_t^{*2}/q_t$. It follows that the posterior distribution of σ^2 is an inverse gamma with parameters $n_t/2$ and $n_t S_t/2$ and so $\sigma^2 \mid y_{1:t}$ concentrates about its mode $\hat{\sigma}_t^2 = n_t S_t (n_t + 2)^{-1}$ asymptotically degenerating. The mode of $\sigma^2 \mid y_{1:t}$ is a better estimator than the mean, because the posterior distribution of $\sigma^2 \mid y_{1:t}$ is not symmetric and this choice is in line with standard practices in Bayesian analysis (Robert, 2007). Conditional on $\sigma^2 = \hat{\sigma}_t^2$ the distribution of the residuals is $e_{t+1} \mid \sigma^2 = \hat{\sigma}_t^2, y_t \sim N(0, \hat{\sigma}_t^2 q_{t+1})$. Model assessment is usually carried out by the standardised residuals, defined at each time t as

$$e_t = \frac{e_t^*}{\sqrt{q_t}}, \tag{5}$$

which distribution is $e_t \mid y_{1:t-1} \sim N(0, 1)$ and hence when the model fit is good one would expect 95% of the residuals to fall within the ± 1.96 quantiles of the standard normal distribution $N(0, 1)$. This can be exploited in order to construct chi-square tests, for more details of which the reader is referred to West and Harrison (1997) and to references therein. In

addition to the above the residuals are serially independent (Durbin and Koopman, 2012), i.e. e_i is independent of e_j , for any $i \neq j$.

3 Control procedure

Consider the situation that the process of interest $\{y_t\}$ is subject to monitoring. SPC is usually concerned with the detection of shifts from the mean or the variance of an observed process. However, since this paper deals with monitoring of locally stationary processes, it is expected that even in the in-control state of the process, some temporal changes in the mean and / or in the variance will incur (e.g. departing from stationarity), but not considered as significant to cause an out of control signal. Hence, comparing and contrasting the process with a constant overall mean and variance is not appropriate here; instead we want to take into account for small shifts, not from an overall constant mean or variance, but from a local mean, e.g. comparing the observed value y_t with a projection in that time, such as the one-step ahead forecast mean at t . This naturally leads us to define control procedures that measure deviations from what would be the distribution of the standardised residuals if the model fit were perfect.

In principle, if the model fit is good, then according to the discussion above, $e_t \mid y_{1:t-1} \sim N(0, 1)$. Thus deviations from the $N(0, 1)$ distribution of the residuals will signify cause of alarm. Thus, it is natural to compare the density $e_t \mid y_{1:t-1} \sim N(0, 1)$ versus alternatives, measuring deviations from this. This observation is leading to the consideration of the null model $\mathcal{M}_0 : e_t \mid y_{1:t-1} \sim N(0, 1)$ and the alternative model $\mathcal{M}_A : e_t \mid y_{1:t-1} \sim N(\mu, \kappa^2)$, where μ measures deviations from the zero mean and κ measures deviations from the standard deviation of $N(0, 1)$. Such a comparison can commence by considering Bayes factors (Robert, 2007). It is of some interest to explore what is the effect to the original process implied by out of control signals on the residuals. Suppose that there is some out of control signal at time t , that is e_t follows a $N(\mu, \kappa^2)$, for some $\mu \neq 0$ and / or $\kappa \neq 1$. Then

$$e_t^{(\text{new})} = \frac{y_t - \hat{f}_t - \mu}{\kappa \sqrt{q_t}} \sim N(0, 1),$$

from which it follows

$$y_t \mid y_{1:t-1} \sim N(f_t + \mu, \kappa^2 q_t),$$

where $f_t = \mathbf{x}^\top \mathbf{m}_{t-1}$ is the one-step ahead forecast mean of y_t . As we see this updated forecast distribution (which is now accurate as $e_t^{(\text{new})}$ follows a $N(0, 1)$ distribution, indicates that y_t has shifted by a mean of μ and a variance of κ^2 times the original variance q_t . In order to explore graphically this even further, we simulate 100 observations from a time-varying AR(1) model (in-control process) and produce the in-control residuals (see solid lines in Figure 1 in both panels of the plot). Subsequently, for the out of control signals first we simulate 100 values of 0 and 1 from a Bernoulli distribution with probability of success $p = 0.1$ and then we add to the existing residuals shifts of size randomly generated by a $N(1, 2)$ distribution at the randomly generated points of “1” (that is we inflate the residuals by 10% shifts each having size simulated from a normal distribution). The randomly generated at time-points 54, 59, 72, 75, 77, 84, 88 (7 points while 3 more were ignored because the added shift outliers was negligible). From the implied values of the residuals e'_t (inflated by the shifts) we determine the implied values of the out of control process $y_t = f_t + \sqrt{q_t}e'_t$; the values of y'_t are indicated by the dashed line in the top panel of Figure 1 and the solid points depict the added points (5 being outliers, while two of them have lower magnitude). We see from Figure 1 that the added shifts (possible outliers) in the residuals depict shifts or outliers in the original process of similar magnitude as that in the residuals.

The Bayes factor provides a way to formally compare two competing models. The application of Bayes factors is similar to testing a full model vs reduced model (as in likelihood ratio tests or deviance analysis in generalised linear models) in classical statistics. However, Bayes factors exhibit several key advantages over the classical tests, e.g. one model does not have to be nested within the other and sequential application – suitable for time series and Phase II analysis – are easily accommodated (West, 1986). In this paper, given a set of data $y_{1:t-1}$, we compare the current model \mathcal{M}_0 (our time series model specified in the previous section) and \mathcal{M}_A , an alternative model specification, which measures deviations from \mathcal{M}_0 , as




Figure 1: Plot of simulated in-control observations together with 5 out of control points (top panel) and respective residuals (bottom panel); the solid lines indicate the simulated data (top panel) and the respective in-control residuals (bottom panel) and the dashed line show the out of control data (top panel) and out of control residuals. The solid points indicate the added out of control points and the horizontal lines in the bottom panel show the ± 1.96 confidence intervals of the $N(0, 1)$ distribution at 5% significance level.

B_t	Evidence against H_0
1 to 3.2	Not worth more than a bare mention
3.2 to 10	Substantial
10 to 100	Strong
> 100	Decisive

Table 1: Jeffreys rules for the interpretation of the Bayes factor.

already mentioned above. From the definition of Bayes factor

$$B_t = \frac{p(e_t | y_{1:t-1}, \mathcal{M}_A)}{p(e_t | y_{1:t-1}, \mathcal{M}_0)}, \quad (6)$$

it is apparent that

- if $B_t < 1$ we have evidence that \mathcal{M}_0 fits to data better than \mathcal{M}_A , while
- if $B_t > 1$ we have evidence that \mathcal{M}_A fits to data better than \mathcal{M}_0 .

Obviously, if $B_t = 1$, the two models are inseparable, in the sense they both produce the same error distribution.

In the context of Bayesian thinking B_t has been used as a tool for decision making. Note that classical hypothesis testing gives one hypothesis (or model) preferred status (the “null hypothesis” H_0), and only considers evidence against it while B_t offers an excellent tool for making a decision about which is the underlying model that produced the data (not simply against a hypothesis). In general, according to the literature (see e.g. Kass and Raftery, 1995) a value of $B_t < 1$ means that \mathcal{M}_0 is more strongly supported by the data under consideration than \mathcal{M}_A . Harold Jeffreys (Jeffreys, 1935,1961) gave a scale for interpretation, summarised in Table 1, of the strength of evidence B_t gives in favour or against \mathcal{M}_0 . Jeffreys’ scale has been used successfully for decades in decision making (Kass and Raftery, 1995; Robert, 2007).

Returning to the context of SPC we suggest that if \mathcal{M}_0 describes the process in the in-control state and \mathcal{M}_A describes the process in the out-of-control state, we can then setup a control procedure based on the classification of B_t . In fact, using this classification (Table 2) we may arrange consecutive values of B_t into three categories. The suggested allocation

Category	B_t value	Interpretation
0	< 3.2	In-control
1	3.2 to 10	Substantial evidence that the process is out-of-control
2	> 10	Strong evidence that the process is out-of-control

Table 2: The proposed control procedure.

is identical to Jeffreys' rules, with only difference being that Jeffreys' last two categories are merged. This merge is considered since Jeffreys' last category relates to values of B_t that are too extreme causing complexity for the user, without offering significant benefits in monitoring. The proposed classification is only one of the possible arrangement. In order to define more complex procedures more than three categories can be considered, however, the additional complexity may not be appealing to the practitioner. The proposed categories and related control procedure results in a simple and effective monitoring approach, as evidenced in the illustration in the next section.

The probability of assigning a value of B_t in one of these categories is unknown since the distribution of B_t is unknown. Nevertheless, still we may calculate accurately the corresponding probabilities by simulation. Then, after calculating the probabilities of assigning a value of B_t to a category we may define appropriate control rules. In this paper, we propose the following composite rule: The process is out-of-control if

- a value of B_t assigned to category 2 appears in the sequence, or
- two values of B_t assigned to category 1 appear in the sequence, in a window of length $k = 4$.

The length k of the window is a parameter of the control procedure and may be changed accordingly depending on desired sensitivity set-up by the modeller. A smaller value of k gives small sensitivity, while larger values of k increase the sensitivity of the procedure. The ARL of this procedure may be studied using the Markov chain approach, as described The method is described in the appendix.

In the remaining of the section we proceed to an extensive exploration of the performance of the new control procedure based on the rules given in Table 1. As it is already mentioned

above, these rules were first proposed by Jeffreys (1935) in his pioneering paper on Bayes factors and later reported in his book (Jeffreys, 1961). They have been used extensively and documented widely within the practice of the Bayesian paradigm; for an excellent exposition see Kass and Raftery (1995) and Robert (2007, Section 5.2).

Since the rules given in Table 1 have been tested in practice in many fields, we engage them for the first time in SPC after appropriately modified them (see Table 2). Thus, in the sequel we provide numerical results related to the performance of the control procedure by assuming the control limits as fixed in the way Jeffreys proposed. The exploration of the rules performance is done using the ARL (average run length).

In Table 3 we give the ARL values for the case that an event that affects the time series is the cause of an out-of-control state. It is evident that under the assumption that the process is in-control the corresponding ARL will tend to infinity (equation (6) takes the value 1, with probability 1, if $\mathcal{M}_A = \mathcal{M}_0$, which is the case of an in-control procedure). This is similar to testing a null hypothesis H_0 against an alternative that is exactly equal to the null $H_A \equiv H_0$; in such a case the ARL is equal to infinity as you would require infinite observations in order to reject H_0 (since $H_A \equiv H_0$). In other words the value of infinity in the ARL depicts that no matter how many observations we have, there will be no chance to reject the null hypothesis (model \mathcal{M}_0). This is a very significant characteristic of the control procedure since if we compare the model with itself the probability of rejecting the model should be zero. In the context of Bayesian hypothesis testing, a discussion on the uncertainty related to the parameters of the alternative hypothesis can be found in Aitkin (1997) and Aitkin et al. (2005). As we define shifts-departures from the null model (i.e. we specify the alternative distribution shifting from $N(0, 1)$), we observe that the process is getting out of control (because the true distribution is not the null any more) and the ARL decreasing from infinity to lower values. As we can see the proposed procedure can identify very fast (ARL=2.02) a shift in the original process that causes deviations to the mean of the residuals equal to two sigma. In case that a shift in the original process appears and causes deviations to the mean of the residuals equal to three sigma then the proposed procedure will identify almost immediately (ARL=1.28). The results of Table 3 suggest that the rules given in Table

State	Distribution	mean	variance	ARL
in-control	Normal	0.00	1.0	∞
out-of-control	Normal	0.25	1.0	1349.74
	Normal	0.50	1.0	978.92
	Normal	1.00	1.0	9.23
	Normal	2.00	1.0	2.02
	Normal	3.00	1.0	1.28

Table 3: ARL when the residuals are affected by mean shifts.

State	Distribution	mean	variance	ARL
in-control	Normal	0.00	1.0	∞
out-of-control	Normal	0.00	2.0	33.40
	Normal	0.00	3.0	9.02
	Normal	0.00	4.0	5.58
	Normal	0.00	9.0	2.69

Table 4: ARL when the residuals are affected by variance shifts.

2 lead to an efficient monitoring procedure.

Table 4 contains the ARL values in the case of an event that affects the time series is causing an out-of-control state characterised by aberrant fluctuations of the residuals. As we can see the proposed procedure can identify fast (ARL=33.40) a shift in the original process that doubles the corresponding variance of the residuals while this expected time is event smaller in the case that the shift in the original process triples the corresponding variance of the residuals.

In Table 5 ARL values are provided in the case of an event that affects the time series is causing an out-of-control state characterised by aberrant fluctuations of the residuals as well as to a shift in the mean of the model deviations. The ARL values are quite small justifying once more the overall good performance of the proposed method.

Finally, in Table 6 ARL values are provided in the case that an event that affects the time series is causing an out-of-control state characterised by a change in the distribution of the

State	Distribution	mean	variance	ARL
in-control	Normal	0.00	1.0	∞
out-of-control	Normal	1.00	2.0	5.87
	Normal	2.00	2.0	2.25
	Normal	1.00	4.0	3.94
	Normal	2.00	4.0	2.33

Table 5: ARL when the residuals are affected by both mean and variance shifts

State	Distribution	mean	variance	ARL
in-control	Normal	0.00	1.0	∞
out-of-control	t_{15}	0.00	1.0	957.92
	t_{10}	0.00	1.0	276.34
	t_5	0.00	1.0	52.86
	t_3	0.00	1.0	21.12

Table 6: ARL when the residual distribution changes.

residuals. As an alternative model we consider the standard Student t -distribution t_ν with various values of degrees of freedom ν . As the degrees of freedom decrease, the underlying process of the residuals shifts away from the model of the normal distribution. As we can see in the case the degrees of freedom are equal to 3 the ARL is equal to 21.12. The method appears to be sensitive to the disturbance distribution; in particular the values of the ARL increase as the degrees of freedom increase (for 15 degrees of freedom the ARL is equal to 957.92). However, we observe that the t distribution considered as an alternative model has mode 0 and scale equal to 1; so in other words the alternative model considers deviations of zero mean and of variance $\nu/(\nu-2)$, which for large degrees of freedom are close to 1 (same as in the null model). Thus, the main deviation from the null model comes in the tails between the t and the normal distribution. When we consider a t alternative model with mean and scale shifts, i.e. a $t_\nu(1, 2)$, then the ARL is much lower even for $\nu = 15$ (results not shown).

From the above results it is evident that the proposed method performs excellent in a wide range of situations. In order to further evaluate the performance of the proposed chart we

State	Distribution	mean	variance	ARL_{b_t}	ARL_{e_t}
in-control	Normal	0.0	1.0	∞	370.4
out-of-control (mean shifts)	Normal	0.50	1.0	983.00	155.22
	Normal	0.6208	1.0	113.33	113.37
	Normal	0.75	1.0	31.96	81.35
	Normal	1.00	1.0	9.37	43.82
	Normal	2.00	1.0	2.14	6.30
	Normal	3.00	1.0	1.29	2.00
out-of-control (variance shifts)	Normal	0.00	2.0	33.10	29.49
	Normal	0.00	3.0	9.08	12.00
	Normal	0.00	4.0	5.66	7.49
	Normal	0.00	9.0	2.72	3.15
out-of-control (both mean and variance shifts)	Normal	1.00	2.0	5.95	12.37
	Normal	2.00	2.0	2.31	4.17
	Normal	1.00	4.0	3.98	5.52
	Normal	2.00	4.0	2.37	3.17

Table 7: Comparison of the proposed control chart with Alwan and Roberts (1988) control chart. Shown is the in-control ARL and out-of-control ARL for mean and variance shifts.

conduct a comparative study with a classical control scheme for autocorrelated data, proposed by Alwan and Roberts (1988). Table 7 shows in-control and out-control ARL for the proposed chart (ARL_{b_t}) and that of Alwan and Roberts (1988) (ARL_{e_t}), for mean shifts, variance shifts and shifts in both the mean and the variance. The ARLs are computed using Monte Carlo simulations of size 100,000. Since the proposed chart is based on the relative comparison of an in-control state and an out of control state (residual distributions), it has the implication of an infinite ARL when both states are in control (i.e. the alternative hypothesis is the same as the in control process). This makes the chart not directly comparable with other procedures which exhibit finite in-control ARL, such as that of Alwan and Roberts (1988). Nevertheless, Table 7 illustrates that the proposed control chart has very good performance

for a wide range of out of control shifts (alternative hypotheses); indeed for that range of out of control scenaria, it outperforms the classical chart mentioned above. In particular we note that for mean shifts exceeding 0.6208σ has smaller in-control ARL compared to that of Alwan and Roberts (1988). For the proposed control chart, the table highlights a very steep drop of the ARL in relation to increases in the mean shifts; the ARL of Alwan and Roberts (1988) does not compare well with this. For example, for a mean shift of 1σ , the ARL of the proposed control chart is 9.37, while the ARL of the classical chart is 43.82. A similar strong performance is observed when variance shifts and shifts in both the mean and variance are considered. The proposed control chart offers the advantage to monitor deterioration in both the mean and the variance, with notable simplicity within a single chart.

4 Production data set

In this section we consider two data sets consisting of temperature values (in $^{\circ}\text{C}$) of 2 components in the production process of a plastic mould. The data we consider here are two of the variables of the production data described in some detail in Pan and Jerrett (2004).

4.1 Production data: variable 1

Figure 2 shows 276 temperature values of the first component mentioned above; we observe that the data exposes a clear time-series autocorrelation (this is obvious by the local variation of the data and can be explored formally by looking at the sampled autocorrelation function (ACF)). Indeed this figure and the ACF indicates that the data is a non-stationary time series, with level (or time-windowed mean) growing for about the first 120 time-points, then decaying until about $t = 150$, from which point they are stabilised to approximately a constant level. For this data set we consider Phase I up to time $t = 220$ (in the sense that the time series model is estimated) and Phase II (time $t = 221 - 276$); the right vertical line in Figure 2 indicates the end of Phase I. For Phase I analysis we have excluded the first 20 observations needed for the model to train to the data set; this is indicated in Figure 2 by the left vertical

Production data: variable 1




Figure 2: Plot of the 1st variable of the production data. The solid points in Phase II indicate out of control signals.

line. Hence there are 200 initially fitted observations in Phase I and 56 observations in Phase II.

We start by discussing the model fit in Phase I. Model (1) is a good candidate model because it describes stationary time series (as it is autoregressive model), but incorporating local departures from stationarity, which are described by the random walk evolution of the AR coefficients; in other words this can be seen as a locally-stationary time series model, which matches to the data, as evidenced from Figure 2. Moreover, this claim is backed empirically from the goodness of fit results, discussed next.

For the implementation of the model in Phase I the normal and inverse gamma priors for β_8 and σ^2 are adopted and the initial values $\mathbf{m}_8 = [1, 1, \dots, 1]^\top$, $\mathbf{C}_8 = 1000\mathbf{I}$ (\mathbf{I} being the 8×8 identity matrix) are picked for the prior of ϕ_8 and $n_8 = 1/100$ and $S_8 = 1$ are picked for the precision $1/\sigma^2$. The above prior setting suggests a weakly informative or vague Bayesian prior specification, depicted by \mathbf{C}_8 and n_8 ; in one hand, the large prior covariance matrix $\mathbf{C}_8 = 1000\mathbf{I}$ and the low degrees of freedom $n_8 = 1/100$ in another hand, are designed to input vague or weak information into the posterior learning. We note that with $n_8 \approx 0$ the posterior estimate of σ^2 is close to its maximum likelihood estimate. For detailed discussion on vague prior specification the reader is referred to Robert (2007).

It remains to select the values of the discount factor δ (responsible for the dynamic evolution of the time-varying AR coefficients) and the order of the AR model p . We used the following criteria to choose a model:

- the mean of squared standardised residuals $\text{MSR} = 200^{-1} \sum_{t=21}^{200} e_t^2$ should be close to 1, where e_t are the standardised residuals defined in (5);
- the residuals in Phase I should be independent, or at least uncorrelated. From the theory of state space models (Durbin and Koopman, 2013) it follows that the residuals must be independent; hence the empirical results should back the theory, if the model fit is good.

Out of several possible combinations of the models (for discount factor $\delta = 0.5, 0.56, 0.6, 0.7, 0.8, 0.9, 1$ and AR order $p = 1, 2, 3, 4, 5, 6, 7, 8, 9, 10$) the model with $\delta = 0.56$ and $p = 8$

was the only one to fulfil both of these two requirements (for this model $\text{MSR} = 1.018$ and all lags of the ACF of the residuals fall inside the $\pm 0.141 = 2/\sqrt{200}$ confidence intervals, hence suggesting the residuals are uncorrelated. As p increases the log-likelihood increases, indicating that a high order AR model is suitable. This could give some support to considering moving average time series models, but their estimation is not as parsimonious as the AR, because of unobserved components or innovations, for a discussion of which the reader is referred to Triantafyllopoulos and Nason (2007). We also observe that the log-likelihood increases for low values of δ . The worst model is obtained for $\delta = 1$, which is not surprising as this is just a time-invariant AR model (there is not change in the AR coefficients) and this model should be appropriate only to stationary time series data. In contrast with the suggestions of a discount factor close to 1 (West and Harrison, 1997), our results suggest the use of a low discount factor, which is capable of capturing the dynamics of the AR coefficients and addressing the issue of non-stationarity of the data, apparent in Figure 2. However, if δ is too low, then this introduces significant variability to the posterior estimates of the AR coefficients ϕ_t , since \mathbf{V}_t has very large values. Here we have settled for a low enough discount factor $\delta = 0.56$ as mentioned above.

Figure 3 plots the (standardised) residuals against the ± 1.96 quantiles of the normal distribution. We observe that there are exactly 5% (10 out of 200) residuals outlying the ± 1.96 quantiles of the $N(0,1)$ confirming the goodness of fit. The left panel of Figure 4 shows the normal probability plot of the residuals in Phase I; we note the 10 (5%) outliers which are depicted also in Figure 3. It is noted that if the first 40 observations are used as a training data set, then the process is much more in control in Phase I (see the right panel of Figure 4). It seems that the low negative residuals in Phase I occurred in the beginning of the process and could be attributed to model training. Finally, we perform a chi-square test based on the test statistic $Q = \sum_{t=21}^{220} e_t^2$. Under the null hypothesis of $e_t \sim N(0,1)$ independently, $Q \sim \chi_{199}^2$ and with observed $Q = 203.65$ the p -value is $P(Q > 203.65) = 0.396$, hence there





Figure 3: Standardised residuals against ± 1.96 quantiles of the $N(0, 1)$ distribution.

QQ-plot of residuals (time 21-220)



QQ-plot of residuals (time 41-220)




Figure 4: QQ normal plot of the standardised residuals in Phase I.

is no evidence to suggest that e_t does not follow a $N(0, 1)$ in Phase I. All the above establish the time series model in Phase I.

Proceeding now to Phase II, with the specification of the Bayes factor as detailed in the section above and defined in equation (6), is

$$B_t = \frac{p(e_t | y_{1:t-1}, \mathcal{M}_A)}{p(e_t | y_{1:t-1}, \mathcal{M}_0)} = \kappa^{-1} \exp[2^{-1}e_t^2 - 2^{-1}\kappa^{-2}(e_t - \mu)^2].$$

Noting that the alternative model has a $N(\mu, \kappa^2)$, measuring deviations from $N(0, 1)$ of the residuals, we need to specify μ and κ . Obviously, if μ, κ are too close to 0,1, respectively, the two models will be harder to separate; on the other hand if μ, κ are large (in modulus), then the control procedure may fail to detect small shifts in mean or variance. In this paper for illustration purposes we consider two alternative scenarios (a) $N(1.5, 1)$ (increase in mean by 1.5 unit, 150%) and (b) $N(0, 1.3)$ increase in variance by 30%. Other alternatives were considered (not shown here), but gave no out of control signal or the same out of control signals as the above. We note that in Phase II we do not know which alternative models may be true, hence to declare the process to be in-control a range of alternative models needs to be considered and no signal to be issued. However, a case that appears many times in practice is that the practitioner knows a set of possible out-of-control states. Thus, the practitioner may test for the deviation distributions corresponding to these specific out-of-control states. In other words, if an alternative model gives a Bayes factor in category 0 this gives only support in favour of the in-control hypothesis for the given alternative model.

Table 8 shows the control classification of Table 2 for Phase II of the production data using the Bayes factor with the alternative model $e_t \sim N(1.5, 1)$, i.e. $\mu = 1.5$ and $\kappa = 1$. Out of 56 time points in Phase II there are 3 out of control signals (indicated in Table 8 by either 1 or 2). Highlighted in Table 8 are the first two signals that each concerns a different sub-rule, for $t = 245$ (sub-rule 1) and $t = 254, 270$ (sub-rule 2). In a comparison with an EWMA control chart in the residuals in Phase II, the proposed methodology seems to be even more sensitive than the EWMA; the latter gave an out of control signal on the 257th observation. The three solid points in Figure 2 indicate the out of control signals in Phase II at times 245, 254 and 270.

The Bayes factor using the alternative model $e_t \sim N(0, 1.3)$, i.e. $\mu = 0$ and $\kappa = 1.3$ did

Time	221	222	223	224	225	226	227	228	229	230	231	232	233	234
Category	0	0	1	0	0	0	0	0	0	0	0	0	0	0
Time	235	236	237	238	239	240	241	242	243	244	245	246	247	248
Category	0	0	0	0	0	0	0	1	0	0	1	0	0	0
Time	249	250	251	252	253	254	255	256	257	258	259	260	261	262
Category	0	0	0	0	0	2	0	0	0	0	0	0	0	0
Time	263	264	265	266	267	268	269	270	271	272	273	274	275	276
Category	0	0	0	0	0	0	0	2	0	0	0	0	0	0

Table 8: Control in Phase II (time $t = 221 - 276$) for the production data. Out of control signals are those with associated counts 1 or 2.

not indicate any out of control signals in Phase II. Considering monitoring both the mean and the variance, again several alternative model needs to be considered. Here we note that all of these models and in particular the model $e_t \sim N(1.5, 1.3)$ gave identical results as in the mean-only case, i.e. $\mu = 1.5$ and $\kappa = 1$. Consulting the above we conclude that any possible deviations from the in-control state of the process are due to mean shifts. The proposed control procedure can be used to identify source of deterioration in an autocorrelated process. The proposed control procedure identified that the process is out-of-control. This fact was also identified by the EWMA control chart as well as by the control charts proposed by Pan and Jarrett (2004) and Triantafyllopoulos (2006). It is evident that the proposed procedure identifies the corresponding out-of-control cases providing an autocorrelation adjusted control procedure with excellent performance.

4.2 Production data: variable 2

In this section we consider 200 temperature values of the 2nd component of the production data. In order to investigate whether the process is in control or not, we follow a similar analysis as in variable 1 above. Figure 5 shows the data; we use the first 150 observations in Phase I and the last 50 in Phase II. We are going to fit the time-varying autoregressive

Production data: variable 2




Figure 5: Plot of the 2nd variable of the production data. The solid points in Phase II indicate out of control signals.

(TVAR) model of Section 2.2 in Phase I. We observe that the data in Phase I appears to be much more stable than the equivalent data in phase I for the 1st component of the production data, considered in Section 4.1 above.

Following a similar data experimentation as in Section 4.1, the fitted model in Phase I is a TVAR model (1) of order $p = 2$, with $\delta = 0.99$. This experimentation involves fitting several TVAR models in Phase I using the two criteria set out in the previous section (small MSR and independence of the residuals) for a range of values of the AR order p and the discount factor δ . We note that in contrast with the TVAR model of Section 4.1 (on the first component of the production data) the selected discount factor $\delta = 0.99$ is considerably high and close to one. This high discount factor allows for relatively small fluctuations of the time-varying AR coefficients reflecting on the relatively stable process in Phase I and being in line with the suggestions of West and Harrison (1997). As in the previous section, we use the first 20 observations as a training data set; Figure 5 depicts this as well as Phase II by the two vertical lines. Figure 6 shows the standardised residuals in Phase I together with the usual ± 1.96 confidence intervals; out of 150 observations in Phase I there are 8 outliers corresponding to a 5.3% of outliers, which is close to the ideal 5% of a good model fit. Figure 7 shows the normal probability plot of the standardised residuals; the left panel considers the residuals when the first 20 observations are used as a training data set and the right panel when the first 50 observations are used as a training data set. We note that when only 20 observations are used as a training data set there are some outliers, which is also depicted in the residual plot of Figure 6; the situation is considerably improved when the first 50 observations are used as a training data set (right panel of Figure 7).

Moving on to Phase II analysis, several alternative models $\mathcal{M}_A: N(\mu, \kappa^2)$ are considered in order to quantify the strength of evidence the Bayes factor B_t gives for the out of control behaviour of the process. Out of several alternative models the model $N(2, 1)$ gives 3 out of control signals (sub-rule 2) at time points $t = 160, 170$ and 185 , as it is illustrated in Table 9





Figure 6: Standardised residuals against ± 1.96 quantiles of the $N(0, 1)$ distribution.

QQ-plot of residuals (time 21-150)



QQ-plot of residuals (time 51-150)




Figure 7: QQ normal plot of the standardised residuals in Phase I.

Time	151	152	153	154	155	156	157	158	159	160	161	162	163	164
Category	0	0	0	0	0	0	0	0	0	2	0	0	0	0
Time	165	166	167	168	169	170	171	172	173	174	175	176	177	178
Category	0	0	0	0	0	0	2	0	0	0	0	0	0	0
Time	180	181	182	183	184	185	186	187	188	189	190	191	192	193
Category	0	0	0	0	0	2	0	0	0	0	0	0	0	0
Time	194	195	196	197	198	199	200	–	–	–	–	–	–	–
Category	0	0	0	1	0	0	0	–	–	–	–	–	–	–

Table 9: Control in Phase II (time $t = 151 - 200$) for the 2nd variable of the production data. Out of control signals are those with associated counts 1 or 2.

and in Figure 5 (solid points in Phase II). Other alternative models with different shifts from $N(0, 1)$ gave a subset of these 3 out of control signals or no out of control signal. Therefore, based on this analysis, we conclude that for this data set there are three out of control signals in Phase II driven by a shift of 200% in the mean of the standardised residuals, while the variance of the process seems to be stable. This is reinforced when considering both mean and variance shifts and finding out that the possible out of control signals are the same as those considering mean shifts only.

5 Concluding comments

This paper develops a methodology for monitoring autocorrelated processes. A time-varying autoregressive (AR) model is fitted to the data and the standardised residuals are obtained. The model is capable of addressing local non-stationarities via the dynamic evolution of the AR coefficients and so it can describe stable and non-stable processes. Several criteria of goodness of fit based on the residuals are discussed. After the model has been set up in Phase I, a control procedure based on the theory of Bayes factors is proposed. The main idea is to compare the performance of the time series model against some alternative, which quantifies departures from it. As a result a control procedure based on binomial-type statistics and taking into account the empirical rules and guidelines for Bayes factors is proposed. The

methodology is put into practice using two data sets consisting of temperature measurements in two stages of manufacturing of a plastic mould. It is believed that the developed methodology will be particularly appealing to non-industrial process monitoring, where non-stable or non-stationary processes are typical (finance, environmentrics and so forth; for references see the discussion in the introduction). A comparative study using simulation illustrates that the the proposed monitoring procedure outperforms a classical control chart, due to Alwan and Roberts (1988). Future research should consider a more extensive comparison study as well as multivariate extensions of the proposed monitoring scheme and it is our intentions to explore these lines of research in the near future.

Acknowledgements

This work is supported by the General Secretariat for Research and Technology (GSRT, Ministry of Education, Greece) research funding action “ARISTEIA II”. We are grateful to two anonymous referees who provided helpful comments on an earlier version of the paper.

Appendix

Calculation of the ARL

Several authors have exploited Markov chain techniques to study the run length distribution of control charts, see e.g. Bersimis et al. (2007).

Let $\{X_t, t = 0, 1, 2, \dots\}$ be a Markov chain defined over a finite state space $\Omega = \{a_1, a_2, \dots, a_s\}$ so that a_s is an absorbing state and the event $T \leq n$ is equivalent to the event $X_n = a_s$, while $X_{n-1} \neq a_s$. We denote by $\boldsymbol{\pi}_0^\top = [P(X_0 = a_1), P(X_0 = a_2), \dots, P(X_0 = a_s)]$ the (row) vector of initial probabilities of the Markov chain and by $\mathbf{\Lambda} = [P(X_t = a_j \mid X_{t-1} = a_i)]_{s \times s}$ its transition probability matrix. It is then clear that

$$P(T > n) = 1 - P(X_n = a_s) = 1 - \boldsymbol{\pi}_0^\top \mathbf{\Lambda}^n \mathbf{1}_s,$$

where $\mathbf{1}_s^\top = [0, 0, \dots, 0, 1]_{1 \times s}$ is the unit vector of \mathbb{R}^s , while the probability mass function of

the waiting time random variable T can be easily derived, on observing that

$$P(T = n) = P(T > n - 1) - P(T > n) = \boldsymbol{\pi}_0^\top \boldsymbol{\Lambda}^{n-1} (\boldsymbol{\Lambda} - \mathbf{I}) \mathbf{1}_s. \quad (\text{A-1})$$

The last formula offers an easy to use formula for the evaluation of the run length distribution of the proposed control chart. The methodology described briefly here is used here in order to calculate the ARL distribution. The first step is to acquire the in control distribution ($F_{in}(\cdot)$) of B_t . This can be done by simulation.

Having calculated the corresponding in-control distribution $F_{in}(\cdot)$ of B_t we may calculate the probabilities $p_0 = P(B_t < 3.2)$, $p_2 = P(B_t > 10)$ and $p_1 = 1 - p_0 - p_2$. Therefore, provided that the process is in-control, the charting procedure can be modelled by a sequence of multi-state trials with 3 possible outcomes in each trial and respective occurrence probabilities for each outcome p_0, p_1, p_2 .

According to the proposed rule, the events that give an out-of-control signal are: one value of B_t assigned to category 2 or two values of B_t assigned in category 1 in a window of length 4. These events may be described via the multi-state model given above by the family of patterns $\epsilon_a = \{2, 11, 101, 1001\}$. It is now easy to introduce a Markov chain $\{X_t, t = 0, 1, 2, \dots\}$ on the state space $\Omega = \{0, 1, 2, \epsilon_a\}$, where $\epsilon_a = \{2, 11, 101, 1001\}$ is the absorbing state of the chain, and 0, 1, 2 indicate that the last sample used for monitoring the process, produced a point in zones 0, 1, 2 respectively.

If we denote by $\boldsymbol{\pi}_0^\top = [P(X_0 = 0), P(X_0 = 1), P(X_0 = 10), P(X_0 = 100), P(X_0 = \epsilon_a)]$ the (row) vector of initial probabilities of the chain and by

$$\boldsymbol{\Lambda} = \left[\begin{array}{c|ccccc} & 0 & 1 & 10 & 100 & \epsilon_a \\ \hline 0 & p_0 & p_1 & 0 & 0 & p_2 \\ 1 & 0 & 0 & p_0 & 0 & p_1 + p_2 \\ 10 & 0 & 0 & 0 & p_0 & p_1 + p_2 \\ 100 & p_0 & 0 & 0 & 0 & p_1 + p_2 \\ \epsilon_a & 0 & 0 & 0 & 0 & 1 \end{array} \right]$$

its transition probability matrix, we may readily evaluate the probability distribution of the waiting time T by using (A-1).

References

- [1] Aitkin, M. (1997). The calibration of P-values, posterior Bayes factors and the AIC from the posterior distribution of the likelihood (with discussion). *Statistics and Computing*, 7, 253–272.
- [2] Aitkin, M., Boys, R. J. and Chadwick, T. (2005). Bayesian point null hypothesis testing via the posterior likelihood ratio. *Statistics and Computing*, 15, 217–230.
- [3] Alwan, L. C. and Roberts, H. V. (1988). Time-series modeling for statistical process control. *Journal of Business & Economic Statistics*, 6, 87–95.
- [4] Bagshaw, M. and Johnson, R. A. (1975). The effect of serial correlation on the performance of CUSUM tests II. *Technometrics*, 17, 73–80.
- [5] Box, G. E. P., Luceño, A. and Paniagua-Quinones, M. C. (2009). *Statistical Control by Monitoring and Adjustment*, 2nd edition. Wiley, New York.
- [6] Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*, 2nd edition. Oxford University Press, Oxford.
- [7] Bersimis, S., Psarakis, S. and Panaretos, J. (2007). Multivariate statistical process control charts: an overview. *Quality and Reliability Engineering International*, 23, 517–543.
- [8] Frisen, M. (2007) *Financial Surveillance*. Wiley, New York.
- [9] Hawkins, D. M. (1991). Multivariate quality control based on regression-adjusted variables. *Technometrics*, 33, 61–75.
- [10] Jeffreys, H. (1935). Some tests of significance, treated by the theory of probability. *Proceedings of the Cambridge Philosophy Society*, 31, 203–222.
- [11] Jeffreys, H. (1961). *Theory of Probability*, 3rd edition. Oxford University Press, Oxford.
- [12] Johnson, R. A. and Bagshaw, M. (1974). The effect of serial correlation on the performance of CUSUM tests. *Technometrics*, 16 (1), 103–112.

- [13] Karvounidis, T., Chimos, K., Bersimis, S. and Douligeris, C. (2012). An integrated self-evaluated framework for embedding web 2.0 technologies in the educational process, Proceedings of the IEEE Global Engineering Education Conference.
- [14] Kim, S. B., Jitpitaklert, W., Park, S.-K. and Hwang, S.-J. (2012). Data mining model-based control charts for multivariate and autocorrelated processes. *Expert Systems with Applications*, 39, 2073–2081.
- [15] Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795.
- [16] Lund, R. and Seymour, L. (1999). Assessing temperature anomalies for a geographical region: a control chart approach. *Environmetrics*, 8, 469–483.
- [17] Montgomery, D. C. and Mastrangelo, C. M. (1991). Some statistical process control methods for autocorrelated data. *Journal of Quality Technology*, 23, 179–204.
- [18] Pan, J.-N. and Chen, S.-T. (2008). Monitoring long-memory air quality data using ARFIMA model. *Environmetrics*, 19, 209–219.
- [19] Pan, X. and Jarrett, J. (2004). Applying state space to SPC: monitoring multivariate time series. *Journal of Applied Statistics*, 31, 397–418.
- [20] Psarakis, S. and Papaleonida, G. E. A. (2007). SPC procedures for monitoring autocorrelated processes. *Quality Technology and Quantitative Management*, 4, 501–540.
- [21] Robert, C. P. (2007). *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*, 2nd edition. Springer, New York.
- [22] Schioler, L. and Frisen, M. (2012). Multivariate outbreak detection. *Journal of Applied Statistics*, 39, 223–242.
- [23] Schmid, W. (1997). On EWMA charts for time series. In *Frontiers in Statistical Quality Control*, 5, 114–137.
- [24] Shmueli, G. and Burkom, H. (2010). Statistical challenges facing early outbreak detection in biosurveillance. *Technometrics*, 52, 39–51.

- [25] Triantafyllopoulos, K. (2006). Multivariate control charts based on Bayesian state space models. *Quality and Reliability Engineering International*, 22, 693–707.
- [26] Triantafyllopoulos, K. and Nason, G. P. (2007). A Bayesian analysis of moving average processes with time-varying parameters. *Computational Statistics and Data Analysis*, 52, 1025–1046.
- [27] VanBrackle, L. N. and Reynolds, M. R. (1997). EWMA and CUSUM control charts in the presence of autocorrelation. *Communications in Statistics: Simulation and Computation*, 26, 979–1008.
- [28] West, M. (1986). Bayesian model monitoring. *Journal of the Royal Statistical Society Series B*, 48, 70–78.
- [29] West, M. and Harrison, P. J. (1997). *Bayesian Forecasting and Dynamic Models*, 2nd edition. Springer, New York.
- [30] Woodall, W. H. (2006). The use of control charts in health care and public health surveillance (with discussion). *Journal of Quality Technology*, 38, 89–104.
- [31] Woodall W. H. and Montgomery D. C. (1999). Research issues and ideas in statistical process control. *Journal of Quality Technology*, 31, 376–386.
- [32] Ye, N. and Chen, Q. (2001). An anomaly detection technique based on a chi-squared statistic for detecting intrusions into information systems. *Quality and Reliability Engineering International*, 17, 105–112.
- [33] Zhang, N. F. and Pollard, J. P. (1994). Analysis of autocorrelations in dynamic processes. *Technometrics*, 36, 354–368.