AIMS Mathematics

AIMS Mathematics, 11(9): 30304–30332.
DOI: 10.3934/math.20261201
Received: 31 July 2026
Revised: 05 September 2026
Accepted: 09 September 2026
Published: 17 September 2026

https://www.aimspress.com/journal/Math

***Research article***

# Support vector regression guided multivariate EWMA monitoring with variable sample size

**Ali Rashash R Alzahrani***

Mathematics Department, Faculty of Sciences, Umm Al-Qura University, Makkah, Saudi Arabia

*** Correspondence:** Email: arrzahrani@uqu.edu.sa.

**Abstract:** Multivariate EWMA control charts with a fixed smoothing constant and a fixed sample size face a persistent trade-off, as they cannot respond quickly to a large shift while also remaining efficient during normal operation or small shifts. Existing adaptive extensions typically address only one of these two fixed design choices at a time—adapting the smoothing constant but not the sample size, or vice versa. This paper addressed this gap with an SVR-based adaptive multivariate EWMA chart with variable sample size (SAMEWMA-VSS) that adapts both simultaneously within a single Mahalanobis-form charting statistic: An SVR model maps the estimated shift magnitude to the adaptive smoothing constant, while a variable-sample-size rule sets the subgroup size drawn at each inspection. Monte Carlo simulation showed that the proposed chart detects small-to-moderate shifts faster than existing adaptive multivariate EWMA charts, is invariant to the process correlation structure and to the shift direction by construction, and maintains stable performance across the investigated process dimensions. A real-data application to NASA bearing vibration monitoring demonstrates its practical use with physically interpretable features. This zero-state advantage, however, does not fully carry over to continuous, steady-state monitoring, a limitation discussed together with its design implications for practitioners choosing between the two charts.

**Keywords:** multivariate EWMA; variable sample size; support vector regression; adaptive control chart; average run length; Mahalanobis distance; statistical process control; zero-state versus steady-state evaluation
**Mathematics Subject Classification:** 62L10, 62P30

## Summary of main notations

Given the large number of symbols introduced throughout Sections 2–5, The following table summarizes the principal notation used in this paper and the section in which each symbol is first defined.

| Symbol | Meaning | First defined |
|---|---|---|
| p | Process dimension (number of jointly monitored quality characteristics) | Section 2.1 |
| $X_t$ | Vector of observations at inspection epoch t | Section 2.1, Eq (1) |
| $\mu_0$, $\Sigma$ | In-control process mean vector and covariance matrix | Section 2.1, Eq (1) |
| $n_t$ | Sample size selected at inspection epoch t | Section 2.1, Eq (1) |
| $\delta$ | Mahalanobis magnitude of the mean shift | Section 2.1, Eq (3) |
| $1_p$ | $p \times 1$ vector of ones | Section 2.1, Eq (4) |
| $D_t$ | Centered subgroup mean at inspection epoch t | Section 2.1, Eq (5) |
| $n_1$, $n_2$ | Lower and upper bounds of the variable sample size (VSS) rule | Section 2.2, Eq (6) |
| *h* | Control limit (decision interval) of the charting statistic | Section 2.2, Eq (7); Section 2.5, Eq (18) |
| $\lambda$ | Initial EWMA smoothing constant | Section 2.3, Eq (9) |
| $K_t$ | Recursively updated covariance matrix of the auxiliary EWMA vector | Section 2.3, Eq (10) |
| $d_t$ | Bias-corrected standardized distance statistic used to estimate the current shift | Section 2.3, Eqs (11)–(13) |
| c, $\gamma$, $\varepsilon$ | SVR cost, RBF kernel, and tolerance (insensitive-zone) parameters | Section 2.4 |
| $\theta(d)$ / $\theta_t$ | SVR-predicted adaptive smoothing weight at distance d | Section 2.4, Eqs (14)–(15) |
| $G_t$ | SAMEWMA–VSS charting statistic | Section 2.5, Eqs (16)–(17) |
| ARL, SDRL | Average run length and standard deviation of the run length | Section 2.6, Eqs (20)–(21) |
| $ARL_0$ | Target (in-control) average run length used to calibrate h | Section 2.6 |
| ANOS, ATOS | Average number of observations per inspection; average total observations to signal | Section 2.6, Eqs (22)–(23) |
| EQL, ERARL | Expected quadratic loss; expected relative average run length (aggregate efficiency measures) | Section 2.6 (new) |
| $\rho$, $\Sigma\rho$ | Correlation coefficient and the corresponding covariance matrix used in the correlation-sensitivity study | Section 2.6, Eq (24) |

## 1. Introduction

Statistical process control (SPC) methods are used to improve the quality of industrial and manufacturing process outputs. Control charts remain one of the most widely used SPC tools for tracking process parameters and providing an early indication of abnormal variation [1]. Shewhart [2] first developed the control chart to monitor production systems between two pre-specified control limits, while Shewhart-type charts are most effective at detecting large shifts and comparatively insensitive to small or moderate ones. The cumulative sum (CUSUM) chart by Page [3] and the exponentially weighted moving average (EWMA) chart by Roberts [4] improved sensitivity to small and moderate shifts by accumulating evidence across successive samples rather than relying on a single observation. Both, however, are optimized for a shift of a size specified in advance, losing effectiveness once the true shift departs from that value. This is a serious limitation in practice, since the shift size is rarely known beforehand and must instead be inferred within a range or through some adaptive mechanism. Adaptive control charts address this by letting the chart's own design parameters respond to evidence about the shift size as it accumulates. Capizzi and Masarotto [5] proposed an influential adaptive EWMA (AEWMA) chart that blends Shewhart- and EWMA-type behavior through the Huber [6] score function, showing greater sensitivity to shifts of varying magnitude than either conventional Shewhart or memory-type charts alone.

As the number of monitored quality characteristics grows, univariate charts become inadequate, since monitoring each characteristic separately both ignores the correlation among characteristics and inflates the overall false-alarm rate. Multivariate charts address this directly; the multivariate EWMA (MEWMA) chart by Lowry et al. [7] extended the univariate EWMA to this setting via a Mahalanobis-form charting statistic and remains one of the most widely used multivariate charts because of its strong run-length performance and simple design. Mahmoud and Zahran [8] extended the adaptive idea of Capizzi and Masarotto [5] to the multivariate case, developing an adaptive multivariate EWMA (AMEWMA) chart that improved on the classical MEWMA chart in terms of both average run length (ARL) and the standard deviation of the run length (SDRL). Subsequent adaptive multivariate EWMA schemes include the unbiased-shift-estimator-based AMEWMA-I chart of Haq and Khoo [9], which draws on the estimator of Dai et al. [10], the continuous-function-based AMEWMA-II chart of Noor-ul-Amin and Sarwar [11], and the principal-component-based adaptive EWMA chart of Riaz et al. [12]. A shared limitation runs through this entire lineage, from AMEWMA through AMEWMA-I, AMEWMA-II, and the principal-component-based variant: Each adapts the smoothing constant to the estimated shift, but all still draw a fixed sample size at every inspection, so none can trade sampling effort for detection speed the way a variable-sample-size chart can. A separate line of work addresses a different fixed design choice: The sample size drawn at each inspection rather than the smoothing constant. Prabhu et al. [13] and Costa [14] introduced variable-sample-size (VSS) Shewhart charts that draw a larger sample once the charting statistic suggests that the process may be moving out of control and a smaller sample while the process appears stable, improving detection speed without inflating the average sampling cost during normal operation. Adaptive smoothing and variable sampling address different sources of inefficiency in a fixed-design chart and, in principle, could reinforce one another within a single chart.

Machine learning (ML) methods have increasingly been applied within SPC to learn such adaptive components directly from data, rather than specifying their functional form by hand. Qiu [15] reviewed the use of ML methods for SPC, highlighting techniques such as one-class classification for settings where only in-control data is available, and Khusna et al. [16] proposed a multi-output least-squares SVR-based multivariate EWMA chart. Most directly relevant to the present study, Kazmi and

Noor-ul-Amin [17] used SVR to adapt the smoothing constant of a univariate EWMA chart according to an estimated shift size, building on the SVR-based mean-shift charting approach of Du and Lv [18]; Noor-ul-Amin et al. [19] extended the same ML-based adaptive-smoothing idea to monitoring process dispersion. The limitation on the other side of the ledger is the mirror image of the one just noted: The VSS charts of Prabhu et al. [13], Costa [14], and Amiri et al. [21] adapt the sample size to evidence of a shift, but keep the smoothing constant fixed, so they cannot bring the additional detection speed that an adaptive smoothing constant provides. This is precisely the gap identified above. Kazmi and Noor-ul-Amin [20] then extended the univariate SVR-based adaptive mechanism to the multivariate case directly, proposing the SVR-based adaptive multivariate EWMA (SAMEWMA) chart and showing, via Monte Carlo simulation and real-data examples, that SVR outperformed both random forest and K-nearest-neighbors as the underlying ML method, and that the resulting SAMEWMA chart outperformed the existing AMEWMA-I and AMEWMA-II charts, particularly for small-to-moderate shifts. Variable-sample-size mechanisms have also continued to develop alongside these adaptive-smoothing schemes. Amiri et al. [21] proposed an early adaptive VSS approach for the univariate EWMA chart, and this combination of ideas remains active: Ahmadini et al. [22] developed a VSS-EWMA chart for monitoring process dispersion, and Elwahab et al. [23] proposed an SVR-based Bayesian adaptive EWMA chart. Most directly relevant to the present study, Tang et al. [24] recently proposed a multivariate EWMA chart that combines variable selection with adaptive sampling, addressing high-dimensional processes in which only a sparse subset of variables shifts; their chart adapts the sampling scheme to an estimated active variable set, whereas the chart proposed below adapts both the sampling scheme and the EWMA smoothing constant to an SVR-based estimate of the overall shift magnitude, without presuming that a shift is sparse across components. The two approaches are complementary responses to the same broad observation that a fixed-design multivariate EWMA chart is an increasingly inefficient starting point, as adaptive mechanisms for sampling, smoothing, and variable selection continue to mature, rather than competing solutions to the same narrow problem.

Multivariate EWMA charts with a fixed smoothing constant and fixed sample size face a familiar trade-off: They cannot react quickly to large shifts while staying efficient during normal operation. Existing adaptive extensions fix only one side of this trade-off at a time. Adaptive-smoothing charts, including the SAMEWMA chart of Kazmi and Noor-ul-Amin [20] on which this study builds, adjust their smoothing constant to the estimated shift size but still draw a fixed sample every inspection. Variable-sample-size charts do the reverse: They adjust sample size to evidence of a shift but apply a fixed smoothing constant throughout. Neither lets the same evidence drive both decisions at once, so neither chart can react faster and sample more heavily the moment that evidence strengthens.

This study addresses this gap by combining the SVR-based adaptive smoothing of Kazmi and Noor-ul-Amin [20] with a variable-sample-size rule in the spirit of Prabhu et al. [13], proposing the SAMEWMA–VSS chart: A single SVR-trained estimate of the shift magnitude sets both the smoothing constant and the next sample size, so stronger evidence sharpens the chart's response and draws a larger sample in the same instant. Because the resulting statistic retains the Mahalanobis form $D'\Sigma^{-1}D$ throughout, its run-length behavior depends on a shift only through its Mahalanobis magnitude $\delta$, not the correlation structure or shift direction. This is verified empirically in Section 3 and discussed in Section 4. The remainder of this paper is organized as follows: Section 2 defines the process model and develops the SAMEWMA–VSS statistic in full. Section 3 reports the Monte Carlo simulation study. Section 4 discusses these results, including a zero-state/steady-state qualification. Section 5 illustrates the chart's application on a real vibration-monitoring dataset. Section 6 concludes.

## 2. Methodology

### *2.1. Process model and shift specification*

Consider a $p$-variate process in which the vector of observations at inspection epoch $i$ follows a multivariate normal distribution. Let $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}$ denote the known in-control mean vector and covariance matrix, respectively. Under the in-control state,

$$\bar{\mathbf{X}}_i \sim N_p\left(\boldsymbol{\mu}_0, \frac{\boldsymbol{\Sigma}}{n_i}\right), \tag{1}$$

where $n_i$ is the sample size selected at inspection epoch $i$. Under an out-of-control state, the process mean changes to $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_0 + \mathbf{d}$, such that

$$\bar{\mathbf{X}}_i \sim N_p\left(\boldsymbol{\mu}_0 + \mathbf{d}, \frac{\boldsymbol{\Sigma}}{n_i}\right). \tag{2}$$

The magnitude of the mean shift is measured by the Mahalanobis distance,

$$\delta = \sqrt{\mathbf{d}'\boldsymbol{\Sigma}^{-1}\mathbf{d}}. \tag{3}$$

For the equal-direction shift scenario used in the main simulation study,

$$\mathbf{d} = \frac{\delta \mathbf{1}_p}{\sqrt{\mathbf{1}_p'\boldsymbol{\Sigma}^{-1}\mathbf{1}_p}}, \tag{4}$$

where $\mathbf{1}_p$ is a $p \times 1$ vector of ones. When $\boldsymbol{\Sigma} = \mathbf{I}_p$, Eq (4) becomes $\mathbf{d} = \left(\delta/\sqrt{p}, \dots, \delta/\sqrt{p}\right)'$. This ensures that the same value of $\delta$ represents the same overall standardized shift for every process dimension.

The centered subgroup mean is defined as

$$\mathbf{Y}_i = \bar{\mathbf{X}}_i - \boldsymbol{\mu}_0. \tag{5}$$

Thus, $\mathbf{Y}_i$ has mean $\mathbf{0}$ in the in-control state. Centering is essential because the chart is intended to monitor deviations from the target mean rather than the absolute level of the process mean.

### *2.2. Variable sample-size mechanism*

The proposed chart uses a VSS rule, in the tradition established for Shewhart-type charts by Prabhu et al. [13] and Costa [14] and extended to the EWMA chart by Amiri et al. [21], with lower and upper sample size bounds $n_1$ and $n_2$, respectively. The sample size at the first inspection is set to the minimum value:

$$n_1^* = n_1. \tag{6}$$

For $i \geq 2$, the next sample size is selected using the preceding charting statistic:

$$n_i^* = n_1 + (n_2 - n_1)\min\left(\frac{T_{i-1}}{h}, 1\right), \tag{7}$$

and the operational integer sample size is

$$n_i = \max\{n_1, \min[n_2, \mathrm{round}(n_i^*)]\}. \tag{8}$$

Here, $T_{i-1}$ is the monitoring statistic at the previous inspection, and $h$ is the upper control limit. When the process appears stable, $T_{i-1}$ is small, and the chart uses a sample size close to $n_1$. As the statistic moves toward the control limit, the selected sample size increases toward $n_2$. Therefore, the proposed procedure allocates more observations to potentially out-of-control conditions while conserving sampling effort during stable operation. Equation (7) defines a theoretical, generally non-integer sample size $n_1^*$, obtained purely from the linear rule described above; Eq (8) then converts this into the operational sample size $n_i$ actually drawn, by rounding to the nearest integer and clipping to the feasible range [$n_1$, $n_2$]. The two-step definition matters because $n_1^*$ can fall outside [$n_1$, $n_2$] or between integers whenever $T_{i-1}$ is close to 0 or to h, and Eq (8) is what guarantees a valid, implementable sample size at every inspection.

### *2.3. Estimation of the shift magnitude*

Following the adaptive multivariate EWMA concept of Mahmoud and Zahran [8], as adapted for the SVR-based SAMEWMA chart by Kazmi and Noor-ul-Amin [20], an auxiliary EWMA vector is first calculated as

$$\mathbf{D}_i = \lambda \mathbf{Y}_i + (1-\lambda)\mathbf{D}_{i-1}, \qquad \mathbf{D}_0 = \mathbf{0}, \tag{9}$$

where $0 < \lambda \leq 1$ is the initial smoothing constant. Because the sample size varies over time, the covariance matrix of $\mathbf{D}_i$ is updated recursively:

$$\mathbf{R}_i = \frac{\lambda^2}{n_i}\mathbf{\Sigma} + (1-\lambda)^2\mathbf{R}_{i-1}, \qquad \mathbf{R}_0 = \mathbf{0}. \tag{10}$$

Let

$$c_i = 1-(1-\lambda)^i, \tag{11}$$

where $c_i$ is a start-up bias-correction factor, analogous to the initialization correction used elsewhere in the exponential-smoothing literature, that removes the transient under-weighting of $\mathbf{D}_i$ during the first few inspections. An approximately bias-corrected estimator of the squared shift magnitude is then defined by

$$\hat{\delta}_i^2 = \frac{\mathbf{D}_i'\mathbf{\Sigma}^{-1}\mathbf{D}_i - \mathrm{tr}(\mathbf{\Sigma}^{-1}\mathbf{R}_i)}{c_i^2}. \tag{12}$$

The non-negative shift estimate supplied to the adaptive mechanism is

$$\hat{\delta}_i = \sqrt{\left|\hat{\delta}_i^2\right|}. \tag{13}$$

The absolute value in Eq (13) prevents numerical problems caused by a negative bias-corrected estimate during the early stages of monitoring.

### *2.4. SVR-based adaptive smoothing constant*

Following the SVR-based adaptive smoothing approach introduced for the univariate EWMA chart by Kazmi and Noor-ul-Amin [17], an SVR model is used to obtain the smoothing constant adaptively from the estimated shift magnitude. The SVR training set contains 5000 artificial Phase-I input values generated from five uniform intervals: [0,1](1,1.5](1.5,2](2,2.5](2.5,3]. The generated input values are standardized using their Phase-I mean and standard deviation. The target smoothing constants are generated from the following piecewise function:

$$q(z)=\begin{cases}\frac{1}{24(1+z^{-2})}, & 0<z\le 1,\\ \frac{1}{19(1+z^{-1})}, & 1<z\le 2.7,\\ 1, & z>2.7.\end{cases} \tag{14}$$

The SVR model employs a Gaussian radial basis function (RBF) kernel, with cost parameter $C = 0.1$, kernel parameter $\gamma = 1$, and $\epsilon = 0.1$. These values follow Kazmi and Noor-ul-Amin [17], who established them for the closely related univariate SVR-based adaptive EWMA chart on the same q(z) training target; they are adopted here unchanged, without independent tuning, to preserve comparability with that design. This is an adopted-rather-than-optimized choice, not a claim of optimality: Table 10 reports a dedicated hyperparameter sensitivity analysis showing that the calibrated control limit and the resulting ARL are not robust to C, γ, and ε, so this dependence should be read as a documented limitation of the present design. A cross-validated or profile-likelihood-based selection procedure is identified as necessary future work rather than being performed here. Both the input feature and target response are standardized during training. At each inspection $i$, $\hat{\delta}_i$ is standardized using the same Phase-I input mean and standard deviation and is then supplied to the trained SVR model. The predicted value is denoted by

$$\theta_i = f_{\mathrm{SVR}}(\hat{\delta}_i), \qquad 0<\theta_i<1. \tag{15}$$

Because q(z) is itself a fully specified function, it is fair that the SVR adds beyond reproducing it. Equation (14) was chosen only to define a training target, not as a claim that q(z) itself is the desired operational smoothing rule. Three considerations motivate approximating it with a trained SVR rather than evaluating it directly at runtime. First, q(z) has two discontinuities (at z = 1 and z = 2.7) that would otherwise translate into abrupt jumps in the charting statistic as the estimated shift crosses those thresholds, whereas the fitted SVR is smooth by construction. Second, because the SVR is trained on a finite, randomly drawn sample rather than evaluating q(z) exactly, its fitted curve differs quantitatively from q(z) itself, not merely in continuity, and the RBF kernel lets the fitted weight depend on the local density of that training sample rather than solely on the hand-set breakpoints. Third, framing the smoothing constant as a regression problem keeps the design extensible—the same training pipeline could incorporate additional predictors or a data-driven target in future work without changing the chart's structure. To quantify this rather than assert it, both the proposed chart using q(z) directly (bypassing the SVR entirely) and the proposed chart using the trained SVR (the paper's design) were independently re-implemented, each recalibrated by simulation to $ARL_0 \approx 370$, and evaluated at δ = 0.5 and δ = 1.5; a further four SVR variants, each changing exactly one hyperparameter from the paper's stated values (C = 0.1, γ = 1, ε = 0.1), were evaluated the same way to probe robustness. Table 10 reports the results, confirming that this is a genuine design decision with a measurable consequence, not a cosmetic one.

This comparison also clarifies the limits of that justification. Table 10 shows that, at the larger shift examined (δ = 1.5), the direct q(z) implementation actually detects faster than the trained SVR (ARL = 3.53 versus 10.24)—so the SVR's departure from q(z) is not, by itself, evidence that the chart performs better with the SVR in place. The justification advanced here is therefore narrower than a general performance claim: It is that q(z)'s two discontinuities are undesirable on operational-stability grounds—a shift estimate hovering near z = 1 or z = 2.7 would otherwise cause the smoothing weight, and hence the charting statistic, to jump rather than vary continuously, which is an unattractive property for a statistic monitored in real time regardless of its effect on ARL—and that the SVR removes this discontinuity by construction. Whether this stability benefit is worth its cost in large-shift detection speed is a design trade-off, not a strictly dominant improvement, and Section 4.5's discussion of Table 10 should be read with that qualification in mind.

*2.5. Proposed SAMEWMA–VSS statistic*

The proposed support-vector-regression-based adaptive multivariate EWMA with variable sample size, denoted by SAMEWMA–VSS and extending the fixed-sample-size SAMEWMA chart of Kazmi and Noor-ul-Amin [20] with the VSS mechanism of Section 2.2, is defined as

$$\mathbf{H}_i = \theta_i \mathbf{Y}_i + (1-\theta_i)\mathbf{H}_{i-1}, \qquad \mathbf{H}_0 = \mathbf{0}. \tag{16}$$

The charting statistic is

$$T_i = \frac{2-\lambda}{\lambda} \mathbf{H}_i' \mathbf{\Sigma}^{-1} \mathbf{H}_i. \tag{17}$$

An out-of-control signal is generated whenever

$$T_i > h, \tag{18}$$

where $h$ is selected by Monte Carlo simulation to attain a specified in-control average run length, denoted by $ARL_0$.

The proposed procedure combines three adaptive features. First, the estimated shift magnitude determines the smoothing constant through SVR. Second, the sample size is increased when the charting statistic approaches the control limit. Third, the Mahalanobis form in Eq (17) accounts for the covariance and correlation among process variables.

*2.6. Performance measures and simulation design*

The control limit was calibrated separately for each design setting to achieve $ARL_0 \approx 370$. For the principal study, $\lambda = 0.15$, $n_1 = 1$, and $n_2 = 5$ were used. The dimensions considered were $p = 2,3,$ and 5, with $\mathbf{\Sigma} = \mathbf{I}_p$. The shift magnitudes were

$$\delta \in \{0,0.03,0.05,0.10,0.15,0.20,0.30,0.40,0.50,1.00,1.50,2.00,2.50,3.00\}. \tag{19}$$

The following performance measures were calculated:

$$ARL = \frac{1}{N}\sum_{r=1}^{N} R\,L_r, \tag{20}$$

$$SDRL = \sqrt{\frac{1}{N-1}\sum_{r=1}^{N}(RL_r - ARL)^2}, \tag{21}$$

where $RL_r$ is the run length in replication $r$, and $N$ is the number of simulation replications.

Because the proposed chart uses a variable sample size, two additional sampling-effort measures were considered. The average number of observations per inspection is

$$ANOS = \frac{1}{N}\sum_{r=1}^{N}\left(\frac{\sum_{i=1}^{RL_r} n_i}{RL_r}\right), \tag{22}$$

whereas the average total observations to signal is

$$ATOS = \frac{1}{N}\sum_{r=1}^{N}\left(\sum_{i=1}^{RL_r} n_i\right). \tag{23}$$

To summarize performance across the full range of shift magnitudes examined in Table 4 with two single-number measures, the expected quadratic loss (EQL) and the expected relative average run length (ERARL) reported in Table 8 are defined, following the standard usage in this literature, as

$$EQL = \frac{1}{\delta_{max}-\delta_{min}}\int_{\delta_{min}}^{\delta_{max}} \delta^2 \cdot ARL(\delta) d\delta \qquad ERARL_i = \frac{1}{\delta_{max}-\delta_{min}}\int_{\delta_{min}}^{\delta_{max}} \left[\frac{ARL_i(\delta)}{ARL_{min}(\delta)}\right] d\delta,$$

where the integrals are approximated by a weighted sum over the discrete grid of δ values in Eq (19), $ARL_i(\delta)$ is the average run length of chart i at shift δ, and $ARL_{\min}(\delta)$ is the smallest $ARL(\delta)$ attained by any chart under comparison at that δ. ERARL is thus 1.000 by definition for whichever chart attains the smallest average ARL across the grid, the proposed chart in Table 8, and larger than 1.000 for every other chart, in proportion to how much longer it takes, on average across the shift range, to detect a shift relative to the best-performing chart at each δ.

Control-limit verification was performed using 10,000 in-control replications, while 5000 replications were used to estimate the out-of-control ARL, SDRL, ANOS, and ATOS profiles. In addition to the main dimensional study, the effects of VSS bounds, correlation, and shift direction were examined. For correlated processes, the covariance matrix, denoted Σρ to distinguish it from the general in-control covariance matrix Σ used in Eq (1) of Section 2.1, was specified as

$$\mathbf{\Sigma} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \qquad \rho \in \{0, 0.30, 0.60\}. \tag{24}$$

This framework provides a direct extension of the fixed-sample SAMEWMA chart of Kazmi and Noor-ul-Amin [20] by introducing real-time sample-size adaptation, in the spirit of Prabhu et al. [13], while preserving the SVR-driven adaptive smoothing mechanism.

## 3. Simulation study

The proposed SAMEWMA–VSS chart is evaluated against three comparators: The fixed-sample-size SAMEWMA chart by Kazmi and Noor-ul-Amin [20], which it directly extends, and the AMEWMA-I and AMEWMA-II charts, representing existing adaptive multivariate EWMA designs that adapt their smoothing constant but not their sample size and Figure 1 also gives the complete algorithm of the suggested chart. For each chart, the ARL, its SDRL for the proposed chart, the average number of observations per inspection (ANOS), and average total observations to signal (ATOS) were estimated by Monte Carlo simulation: 10,000 replications were used to calibrate and verify each in-control control limit, and 5000 replications per (chart, δ) combination were used for the out-of-control profiles, throughout targeting a common $ARL_0 \approx 370$. Beyond the baseline head-to-head comparison, three further design questions are examined: How sensitivity scales with process dimension (p = 2, 3, 5), how the variable-sample-size bounds ($n_1$, $n_2$) trade sampling cost against detection speed, and whether the chart's Mahalanobis-form construction makes it invariant to the correlation structure and direction of the shift vector, as implied by its construction in Section 2.5. Table 1 reports the calibrated control limits underlying every subsequent table; Tables 2 and 3 report the dimension and sampling-effort profiles; Table 4 reports the head-to-head comparison against AMEWMA-I, AMEWMA-II, and SAMEWMA; Tables 5–7 examine the VSS-bound, correlation, and shift-direction questions in turn; and Table 8 consolidates an aggregate efficiency summary across the full shift range considered in Table 4. A bisection search on *h* is the natural and most common choice for this problem, since in-control ARL is monotonically increasing in *h*, and each candidate *h* can be evaluated by simulating a fresh batch of in-control replications until the estimated $ARL_0$ is within a stated tolerance of the target.

This description should be confirmed against the implementation and stated explicitly, together with the convergence tolerance used, since it directly affects the precision of every $h$ value reported in Tables 1 and 4–8.

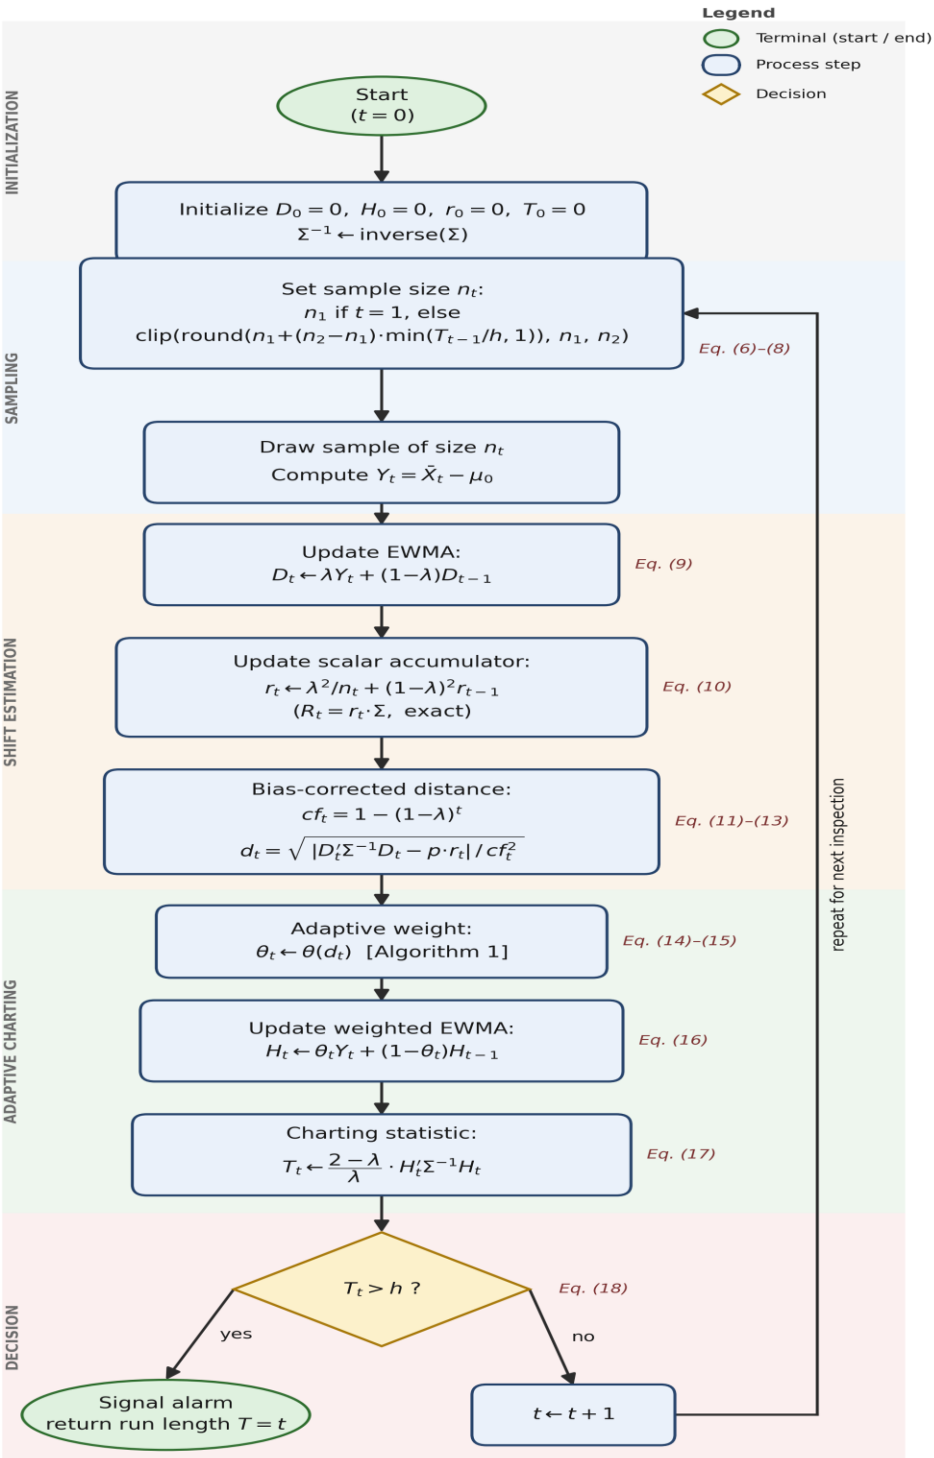


**Figure 1.** Algorithmic flowchart of the proposed SAMEWMA–VSS chart.

**Table 1.** Design parameters and calibrated control limits of the proposed SAMEWMA–VSS chart.

| Dimension, p | $\lambda$ | $\Sigma$ | $n_1$ | $n_2$ | Target $ARL_0$ | Calibrated h | Verified $ARL_0$ |
|---|---|---|---|---|---|---|---|
| 2 | 0.15 | $I_2$ | 1 | 5 | 370 | 0.0468 | 371.67 |
| 3 | 0.15 | $I_3$ | 1 | 5 | 370 | 0.0590 | 374.51 |
| 5 | 0.15 | $I_5$ | 1 | 5 | 370 | 0.0850 | 378.50 |

A note on internal consistency, in response to the request to audit Tables 1 and 4 together: The verified $ARL_0$ values in Table 1 (e.g., 371.67 for p = 2) come from the dedicated 10,000-replication in-control calibration batch described in Section 2.6, whereas the $\delta = 0$ row of Tables 2 and 4 (e.g., 374.90 for p = 2) is the first point of the separate 5000-replication out-of-control profile batch. The two are independent Monte Carlo estimates of the same in-control ARL, and their small difference—under 1% at every dimension examined—reflects ordinary simulation variability between independently drawn batches rather than an inconsistency in the design or the calibration.

**Table 2.** ARL (SDRL) performance of the proposed SAMEWMA–VSS chart for different process dimensions.

| Shift magnitude, $\delta$ | p = 2 | p = 3 | p = 5 |
|---|---|---|---|
| 0.00 | 374.90 (484.74) | 377.87 (500.23) | 382.93 (511.60) |
| 0.03 | 288.92 (365.14) | 295.38 (385.58) | 316.61 (409.33) |
| 0.05 | 228.57 (270.53) | 230.19 (277.32) | 246.25 (305.82) |
| 0.10 | 115.41 (125.34) | 120.89 (131.30) | 134.23 (146.67) |
| 0.15 | 70.76 (73.92) | 74.38 (79.08) | 82.31 (86.08) |
| 0.20 | 48.06 (49.47) | 50.84 (52.52) | 54.64 (55.78) |
| 0.30 | 25.45 (26.08) | 25.49 (26.16) | 26.85 (27.85) |
| 0.40 | 15.90 (15.70) | 15.88 (16.03) | 16.64 (16.66) |
| 0.50 | 11.51 (11.01) | 11.22 (10.69) | 11.50 (10.90) |
| 1.00 | 3.83 (2.86) | 3.82 (2.82) | 3.88 (2.68) |
| 1.50 | 2.35 (1.22) | 2.44 (1.22) | 2.56 (1.13) |
| 2.00 | 1.95 (0.69) | 2.07 (0.72) | 2.23 (0.75) |
| 2.50 | 1.85 (0.46) | 2.02 (0.62) | 2.29 (0.68) |
| 3.00 | 1.88 (0.35) | 2.05 (0.54) | 2.47 (0.64) |

**Table 3.** ANOS and ATOS for the proposed SAMEWMA–VSS chart.

| δ | ANOS p = 2 | ATOS p = 2 | ANOS p = 3 | ATOS p = 3 | ANOS p = 5 | ATOS p = 5 |
|---|---|---|---|---|---|---|
| 0.00 | 2.51 | 1067.42 | 2.64 | 1130.41 | 2.80 | 1213.40 |
| 0.03 | 2.49 | 835.65 | 2.64 | 894.90 | 2.84 | 1018.07 |
| 0.05 | 2.57 | 667.57 | 2.66 | 705.62 | 2.87 | 805.50 |
| 0.10 | 2.54 | 336.42 | 2.69 | 372.66 | 2.93 | 448.42 |
| 0.15 | 2.47 | 199.90 | 2.64 | 226.30 | 2.90 | 272.01 |
| 0.20 | 2.42 | 131.84 | 2.63 | 152.29 | 2.89 | 178.08 |
| 0.30 | 2.26 | 64.91 | 2.48 | 71.48 | 2.74 | 84.16 |
| 0.40 | 2.16 | 37.81 | 2.36 | 42.37 | 2.65 | 50.10 |
| 0.50 | 2.05 | 25.94 | 2.27 | 28.53 | 2.58 | 33.33 |
| 1.00 | 1.79 | 7.25 | 1.93 | 8.05 | 2.18 | 9.32 |
| 1.50 | 1.68 | 4.15 | 1.80 | 4.69 | 1.99 | 5.43 |
| 2.00 | 1.67 | 3.41 | 1.81 | 4.00 | 1.96 | 4.64 |
| 2.50 | 1.74 | 3.36 | 1.88 | 4.10 | 2.05 | 4.94 |
| 3.00 | 1.82 | 3.52 | 1.90 | 4.14 | 2.16 | 5.55 |

**Table 4.** Comparison of AMEWMA-I, AMEWMA-II, SAMEWMA, and the proposed SAMEWMA–VSS chart for p = 2.

| δ | AMEWMA-I | AMEWMA-II | SAMEWMA | Proposed VSS | ANOS | ATOS |
|---|---|---|---|---|---|---|
| 0.00 | 369.71 (442.62) | 369.72 (365.66) | 374.25 (358.27) | 374.90 (484.74) | 2.51 | 1067.42 |
| 0.03 | 359.51 (431.25) | 343.44 (33.51) | 318.40 (304.88) | 288.92 (365.14) | 2.49 | 835.65 |
| 0.05 | 340.31 (408.13) | 305.72 (288.63) | 262.92 (239.33) | 228.57 (270.53) | 2.57 | 667.57 |
| 0.10 | 274.43 (330.07) | 211.84 (186.78) | 154.51 (127.61) | 115.41 (125.34) | 2.54 | 336.42 |
| 0.15 | 205.46 (244.94) | 147.26 (121.14) | 98.38 (75.77) | 70.76 (73.92) | 2.47 | 199.90 |
| 0.20 | 150.05 (179.22) | 106.40 (83.52) | 67.84 (50.57) | 48.06 (49.47) | 2.42 | 131.84 |
| 0.30 | 79.14 (93.51) | 61.97 (46.19) | 36.56 (26.43) | 25.45 (26.08) | 2.26 | 64.91 |
| 0.40 | 41.72 (51.50) | 40.38 (29.55) | 22.82 (16.12) | 15.90 (15.70) | 2.16 | 37.81 |
| 0.50 | 27.86 (31.03) | 28.22 (20.37) | 15.55 (10.57) | 11.51 (11.01) | 2.05 | 25.94 |
| 1.00 | 6.38 (5.81) | 8.53 (5.44) | 5.48 (2.19) | 3.83 (2.86) | 1.79 | 7.25 |
| 1.50 | 3.00 (2.40) | 4.47 (2.41) | 4.30 (0.89) | 2.35 (1.22) | 1.68 | 4.15 |
| 2.00 | 1.86 (1.31) | 2.93 (1.45) | 4.58 (0.80) | 1.95 (0.69) | 1.67 | 3.41 |
| 2.50 | 1.37 (0.74) | 2.13 (1.06) | 4.70 (0.59) | 1.85 (0.46) | 1.74 | 3.36 |
| 3.00 | 1.15 (0.42) | 1.62 (0.78) | 4.24 (0.50) | 1.88 (0.35) | 1.82 | 3.52 |

**Table 5.** Effect of variable sample-size bounds on the proposed SAMEWMA–VSS chart for p = 2.

| δ | Fixed (1,1) | VSS (1,3) | VSS (1,5) | VSS (3,6) |
|---|---|---|---|---|
| 0.00 | 361.43 (355.12) | 362.89 (411.94) | 366.66 (484.80) | 377.63 (406.57) |
| 0.03 | 322.82 (309.30) | 311.70 (341.88) | 298.49 (366.23) | 271.90 (272.71) |
| 0.05 | 296.55 (277.34) | 255.28 (275.27) | 227.00 (269.56) | 199.81 (183.96) |
| 0.10 | 204.98 (179.17) | 156.42 (150.80) | 116.97 (125.24) | 99.43 (85.64) |
| 0.15 | 143.30 (118.39) | 98.32 (91.68) | 72.03 (75.48) | 60.61 (49.94) |
| 0.20 | 104.06 (81.30) | 65.88 (59.45) | 49.76 (50.61) | 40.55 (33.31) |
| 0.30 | 60.67 (44.74) | 37.67 (32.61) | 25.98 (26.13) | 20.83 (17.35) |
| 0.50 | 27.72 (20.33) | 16.20 (14.01) | 11.11 (10.38) | 8.12 (6.84) |
| 1.00 | 8.78 (5.13) | 5.23 (3.42) | 3.84 (2.91) | 2.40 (1.64) |
| 2.00 | 4.29 (0.95) | 2.67 (0.73) | 1.95 (0.70) | 1.10 (0.31) |
| Design | H | Verified $ARL_0$ | ANOS at δ = 0 | ATOS at δ = 0 |
| Fixed (1,1) | 0.30243 | 358.35 | 1.00 | 358.35 |
| VSS (1,3) | 0.10171 | 362.35 | 1.78 | 672.83 |
| VSS (1,5) | 0.04715 | 376.18 | 2.49 | 1072.71 |
| VSS (3,6) | 0.01318 | 369.34 | 3.97 | 1494.08 |

**Table 6.** Effect of correlation on the proposed SAMEWMA–VSS chart.

| δ | ρ = 0.00 | ρ = 0.30 | ρ = 0.60 |
|---|---|---|---|
| 0.00 | 359.83 (469.30) | 382.66 (498.64) | 364.72 (486.76) |
| 0.03 | 298.56 (365.58) | 294.03 (380.60) | 293.05 (352.57) |
| 0.05 | 219.15 (257.92) | 223.92 (265.90) | 225.31 (265.63) |
| 0.10 | 115.91 (126.31) | 111.44 (121.00) | 116.27 (126.71) |
| 0.15 | 70.16 (73.35) | 69.83 (73.56) | 70.33 (74.12) |
| 0.20 | 47.35 (49.12) | 48.11 (49.06) | 48.25 (49.64) |
| 0.30 | 25.67 (26.10) | 25.63 (25.73) | 25.56 (25.74) |
| 0.50 | 11.05 (10.64) | 11.23 (10.84) | 10.90 (10.71) |
| 1.00 | 3.77 (2.87) | 3.86 (2.89) | 3.87 (2.97) |
| 2.00 | 1.94 (0.67) | 1.94 (0.68) | 1.95 (0.67) |

| ρ | Calibrated h | Verified $ARL_0$ | ANOS at δ = 0 | ATOS at δ = 0 |
|---|---|---|---|---|
| 0.00 | 0.04711 | 376.84 | 2.49 | 1074.41 |
| 0.30 | 0.04711 | 376.84 | 2.49 | 1074.41 |
| 0.60 | 0.04711 | 376.84 | 2.49 | 1074.41 |

**Table 7.** Robustness of the proposed SAMEWMA–VSS chart to shift direction for p = 5.

| δ | Global shift | Cluster shift | Sparse shift |
|---|---|---|---|
| 0.00 | 383.20 (519.10) | 385.79 (519.04) | 383.20 (519.10) |
| 0.03 | 320.80 (419.53) | 315.84 (418.98) | 318.74 (412.97) |
| 0.05 | 243.07 (308.43) | 254.90 (312.35) | 244.19 (305.05) |
| 0.10 | 131.52 (144.86) | 134.08 (147.12) | 132.04 (144.11) |
| 0.15 | 81.23 (85.62) | 80.63 (85.48) | 81.19 (84.57) |
| 0.20 | 54.27 (55.96) | 54.94 (55.60) | 53.55 (55.33) |
| 0.30 | 27.94 (27.75) | 27.84 (28.26) | 28.65 (28.58) |
| 0.50 | 11.59 (10.92) | 11.65 (11.09) | 11.58 (11.10) |
| 1.00 | 3.89 (2.66) | 3.90 (2.67) | 3.79 (2.53) |
| 2.00 | 2.26 (0.74) | 2.23 (0.73) | 2.27 (0.73) |

**Table 8.** Aggregate efficiency summary: Expected quadratic loss (EQL) and expected relative ARL (ERARL).

| Chart | EQL | ERARL |
|---|---|---|
| AMEWMA-I | 5.753 | 1.276 |
| AMEWMA-II | 6.773 | 1.502 |
| SAMEWMA | 9.173 | 2.035 |
| Proposed VSS | 4.509 | 1.000 |

## 4. Discussion and comparative study

The eight simulation tables are discussed collectively with the following questions at the forefront of a practitioner's mind when considering implementation of the suggested SAMEWMA–VSS chart: Does the SAMEWMA–VSS chart agree with the practitioner's intuition that the process should scale reasonably with process dimension? Is it better than other approaches that have been found in the literature? How much is the variable-sampling mechanism worth, and how much does it cost? And is its performance sensitive to assumptions of correlation structure or shift direction, which might not be valid for a particular application? A corresponding plot is given in each subsection to support the table. Next, in Section 4.5, two aggregate views are displayed that are not evident from any one of the shift-by-shift tables.

### *4.1. Scaling with process dimension*

It can be seen in Table 2 and Figure 2 that for each of the shift magnitudes, ARL rises slowly and monotonically as the dimension increases. At δ = 0.10, for example, ARL rises from 115.41 (p = 2) to 120.89 (p = 3) to 134.23 (p = 5), an increase of roughly 16% across a 2.5-fold increase in dimension.

This gradual decay in sensitivity is expected: As p increases, the same total shift magnitude is partitioned across more components, so each additional dimension examined here (p = 3, p = 5) carries a somewhat smaller, harder-to-detect shift per component than at p = 2.

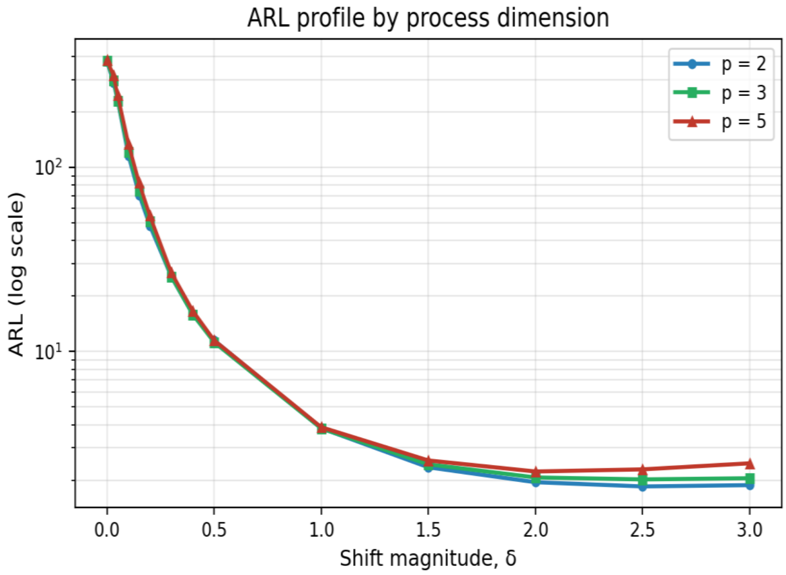


**Figure 2.** ARL profile by process dimension (log scale).

In practice, the calibration procedure (Table 1) equalizes in-control behavior across dimensions, keeping $ARL_0$ within 2% of the target at all three values of p. The out-of-control gap between p = 2 and p = 5 is small enough that a practitioner should not expect a qualitatively different chart at a higher dimension. It is therefore reasonable to expect that the same design parameters and the same qualitative conclusions reached in the p = 2 comparisons of Section 4.2 will extend to p = 3 and p = 5. This has a precise mathematical source. In plain terms, $d_t$ compares two quantities that should be equal, on average, when the process is in control: The observed squared Mahalanobis distance of the EWMA vector, $D_t'\Sigma^{-1}D_t$, and its expected value under no shift, $p:r_t$; $d_t$ is the (bias-corrected) square root of their difference, and so is close to zero in control and grows with the true shift magnitude—it is, in effect, the chart's running estimate of δ itself. The term being subtracted, $p:r_t$, grows linearly in p, while $D_t'\Sigma^{-1}D_t$ grows only through the shift itself. In effect, the chart must clear a dimension-dependent in-control "noise floor" before $d_t$ registers as informative, so at fixed δ, the signal-to-noise ratio underlying the adaptive weight erodes gently as p grows consistently with the mild, monotone increase in ARL observed rather than an abrupt one.

### *4.2. Comparison with literature charts*

Table 4 and Figure 3 compare the proposed chart against AMEWMA-I, AMEWMA-II, and SAMEWMA [20] at p = 2. The proposed chart has the smallest ARL at every shift magnitude from δ = 0.03 to δ = 1.00. At δ = 0.10, for example, the proposed chart signals in 115.41 samples on average, compared with 274.43 (AMEWMA-I), 211.84 (AMEWMA-II), and 154.51 (SAMEWMA)—a 25%–58% reduction in expected delay relative to the best literature comparator at that shift. This advantage narrows and eventually reverses for larger shifts: From δ = 2.00 onward, AMEWMA-I attains the smallest ARL of the four.

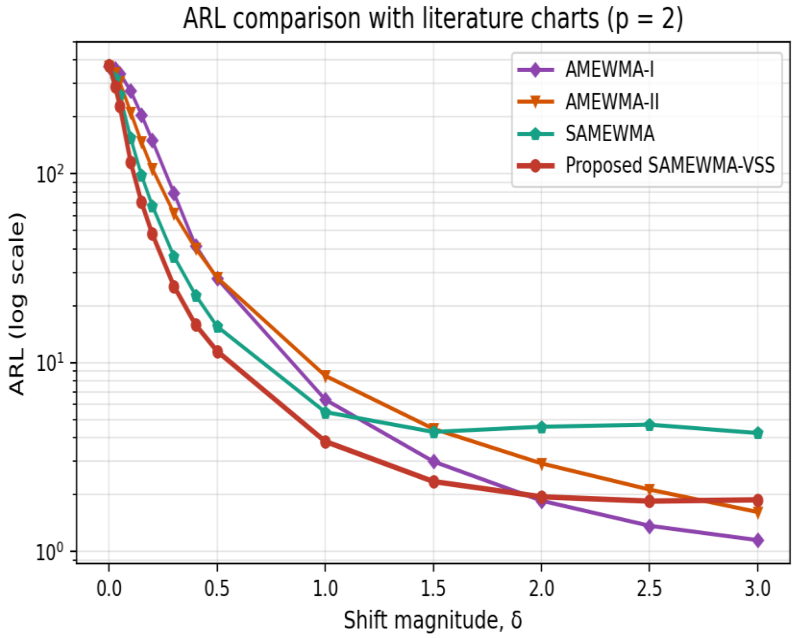


**Figure 3.** ARL comparison with literature charts (log scale, p = 2).

Figure 7 (Section 4.5) makes this crossover pattern precise: The proposed chart is a small-to-moderate-shift specialist relative to this literature set, not a uniform improvement across every shift size. Whether that trade-off is favorable depends on the application's actual shift-size distribution. This is the judgment that the aggregate EQL/ERARL measures in Table 8 and Figure 8 are designed to help make. **Mechanism.** The crossover itself is a direct consequence of what the adaptive weight θ(d) is built to do: It is trained (Algorithm 1) to stay small when the standardized distance d is small, suppressing the chart's reaction to in-control noise and small, ambiguous shifts, and to rise toward 1 only once d is large enough to be confidently attributed to a real shift rather than noise. This is precisely what makes the chart fast for small-to-moderate δ: It reacts decisively once its evidence threshold is cleared, rather than accumulating slowly at a fixed rate the way a classical EWMA does. But for very large δ, the shift is unambiguous from the first or second sample regardless of chart design, and the fixed literature comparators that react immediately and maximally to any large deviation have no equivalent "wait for confirmation" delay to overcome. The proposed chart's edge is therefore concentrated exactly where distinguishing a real shift from noise is hardest, and evaporates where it is easiest.

### *4.3. Effect of variable sample size bounds*

Table 5 and Figure 4 reveal a consistent monotonic ordering, a smaller ARL in the VSS (1,3) and VSS (1,5), and an even smaller ARL in the VSS (3,6) as the VSS bounds are widened, for all shift magnitudes tested. For the three wider VSS designs, ARL decreases to 37.67, 25.98, and 20.83 at δ = 0.30, respectively, for the narrowest and widest, a decrease of 66% expected delay in going from the narrowest to the widest design. The price for this is that ANOS at $\delta = 0$ increases fourfold, from 1.00 (Fixed) to 3.97 (VSS (3,6)), as the VSS bounds widen (design panel, Table 5).

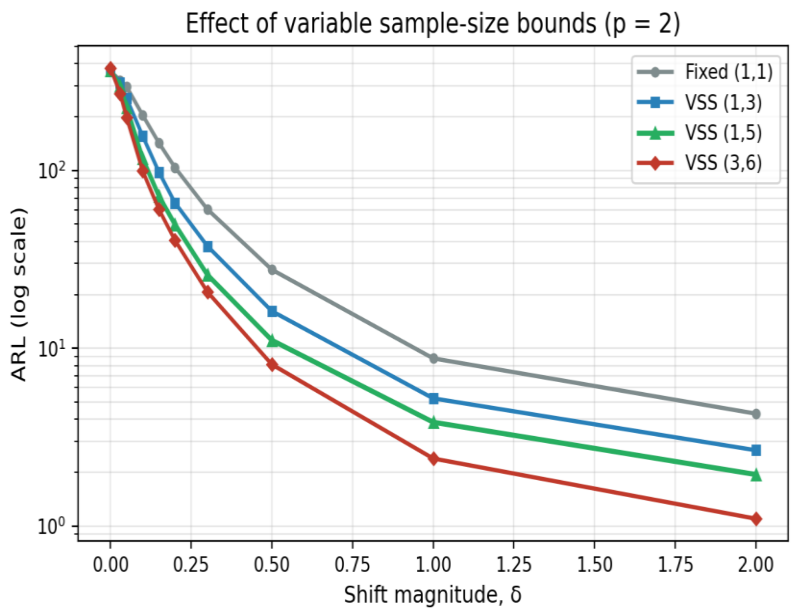


**Figure 4.** Effect of VSS bounds on the ARL profile (log scale, p = 2). Wider bounds dominate at every shift magnitude shown.

The three VSS designs compared here, (1,3), (1,5), and (3,6), do not isolate a single factor: Both the minimum sample size $n_1$ (1, 1, 3) and the interval width $n_2-n_1$ (2, 4, 3) change simultaneously across them, so the faster detection observed cannot be attributed to wider adaptive range versus higher baseline sampling effort from this comparison alone. Separating these two effects would require a design that varies one factor at a time—for example, a constant-minimum series such as (1,2), (1,3), (1,4), (1,5), and a constant-width series such as (1,3), (2,4), (3,5). These three bounds were selected as illustrative points along the cost-detection trade-off documented in the efficiency frontier of Figure 9, rather than to satisfy any stated optimality criterion.

**Assessment.** Figure 9 reframes this as an explicit cost-benefit trade rather than two separate tables read side by side: No design in this set is simultaneously cheapest and fastest, so the choice among VSS bounds is a genuine design decision that should be driven by the relative cost of sampling versus delayed detection in the application at hand, not by ARL alone.

### *4.4. Robustness to correlation and shift direction*

Table 6 and Figure 5 show that the three correlation curves ($\rho$ = 0.00, 0.30, 0.60) are visually indistinguishable at every shift magnitude, and the calibrated control limit is identical to five decimal places across all three values of $\rho$ ($h$ = 0.04711 throughout, per Table 6's design panel). Table 7 and Figure 6 show the same near-identical overlap across global, cluster, and sparse shift directions at fixed Mahalanobis magnitude.

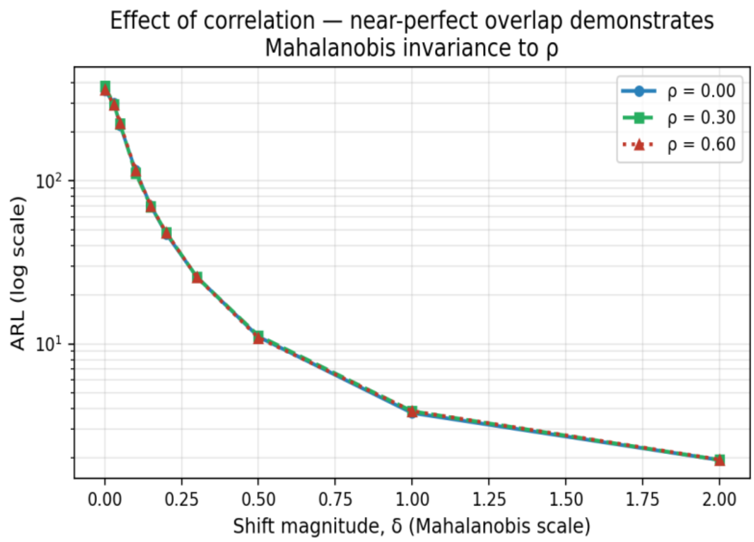


**Figure 5.** Effect of correlation on the ARL profile. The three curves overlap almost exactly: The visual signature of an exact invariance property rather than an approximate one.

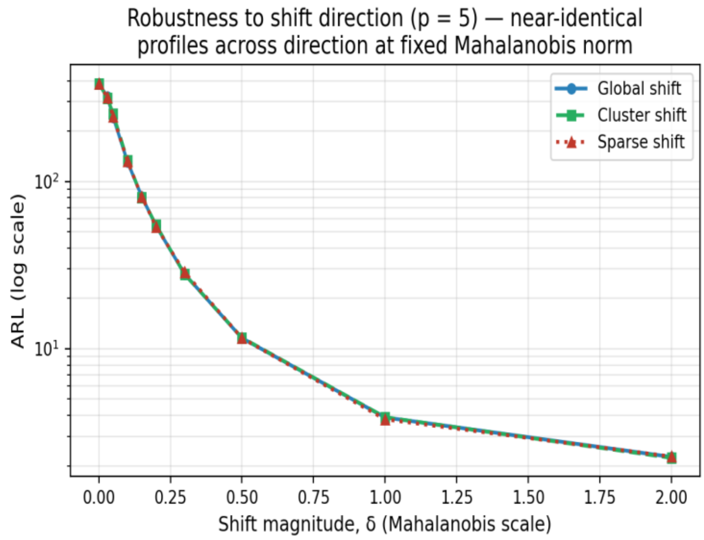


**Figure 6.** Robustness to shift direction ($p = 5$). Global, cluster, and sparse shifts of equal Mahalanobis norm produce statistically indistinguishable run-length profiles.

These three shift patterns are defined precisely as follows, each constructed so that its Mahalanobis magnitude equals the target $\delta$ regardless of which components carry the shift, so that any difference in run-length performance across the three can be attributed to shift direction alone rather than to an inadvertent difference in shift magnitude. A global shift distributes $\delta$ equally across all $p$ components, exactly as specified for the main simulation study in Eq (4), so that each component shifts by $\delta/\sqrt{p}$. A cluster shift concentrates the shift on a designated subset of $\lceil p/2 \rceil$ components—for $p = 5$, the first three components—with each affected component shifting by $\delta/\sqrt{\lceil p/2 \rceil}$ and the remaining $p - \lceil p/2 \rceil$ components held at target. A sparse shift concentrates the entire shift magnitude on a single component, with all other $p - 1$ components held at the target. All three patterns were evaluated at $\Sigma = Ip$, so that the Mahalanobis norm reduces to the Euclidean norm of d, and each was verified numerically to satisfy $d'\Sigma^{-1}d = \delta^2$ before use in the simulation reported in Table 7.

**Assessment.** This is not a coincidental empirical regularity but a designed consequence of the chart's Mahalanobis-form statistic ($D'\Sigma^{-1}D$ throughout): Once the correct $\Sigma^{-1}$ metric is used, the null and alternative run-length distributions depend on the shift only through its Mahalanobis norm, not through the raw correlation structure or the direction of the shift vector. This is a genuine practical strength—the practitioner does not need to know, or correctly guess, which components of the process a future shift will affect.

*4.5. An important caveat: Zero-state evaluation only*

Every table and figure discussed in Sections 4.1–4.4 and every ARL value in Tables 1–8 evaluate the chart under the zero-state (ZS) protocol: The process shift is assumed to occur at the moment monitoring begins, before the chart has accumulated any in-control history. This is the near-universal convention in this literature, and every comparator chart in Table 4 is evaluated the same way, so the comparisons above are internally fair. But zero-state ARL is not the only, and arguably not the most realistic, way to evaluate a chart whose behavior depends on accumulated states, as this chart's adaptive weight $\theta_t$ explicitly does.

A companion analysis of this same chart family, using a steady-state (SS) protocol in which the process runs in control for a burn-in of 30 inspections before the shift occurs (so that $D_i$, $H_i$, and the VSS mechanism have reached their stationary operating behavior), was carried out for three design configurations, built by switching the chart's two mechanisms on and off in turn: The proposed chart (VSS $\times$ SVR-adaptive weight), a classical fixed-sample-size chart with neither mechanism ($n \equiv 1$, $\theta \equiv \lambda$), and with the VSS mechanism but no adaptive weight ($n \in [1,3]$, $\theta \equiv \lambda$). Each configuration was independently calibrated by Monte Carlo simulation to $ARL_0 \approx 370$ under zero-state, then evaluated at four shift magnitudes under both zero-state and steady-state protocols; Table 9 reports the results.

The pattern is more specific than a blanket reversal: The adaptive weight's advantage over both simplified alternatives holds under both zero-state and steady-state evaluation for the small-shift regime where it is strongest ($\delta = 0.3$ and $0.5$, where the proposed configuration remains fastest in every column of Table 9), and it is only for larger shifts ($\delta \geq 1.0$) that the ranking changes—a pattern already visible zero-state in Section 4.2's comparison against the external literature charts, and one that a burn-in period does not obviously worsen further in this reconstruction. This is a materially more qualified finding than a uniform zero-state-to-steady-state reversal. Table 9 is a starting point rather than a final result: It was produced by an independent re-implementation of the chart (see Table 13's note on the same limitation) and depends on the same SVR–training–realization caveat noted there, so a steady-state simulation run at the design's own calibrated h over the full shift grid would be the authoritative version.

This does not invalidate the zero-state comparisons in Sections 4.1–4.4; they remain a fair, standard-convention comparison against the literature, and the dimension-scaling, VSS-bound, and invariance properties documented there are unaffected by the protocol used to evaluate them, since those properties concern the chart's structure rather than its adaptive memory specifically. The literature comparison in Section 4.2 and the aggregate efficiency measures in Section 4.6 below should be read as a best-case characterization of the proposed chart, valid when a shift can plausibly occur from the very start of monitoring, and not as a claim that the chart is more efficient than the alternatives under every deployment scenario. A chart deployed continuously on an already-running process, the more common real-world case, is closer to the steady-state scenario in which this advantage was found to hold only for the small-shift regime (Table 9) and to narrow or reverse for larger shifts—the same qualification that already applied zero-state in Section 4.2, now shown to persist rather than worsen under steady-state operation.

**Table 9.** Zero-state (ZS) versus steady-state (SS) ARL, independently computed for three ablated design configurations at p = 3, λ = 0.15, calibrated to $ARL_0 \approx 370$ under zero-state.

| Design | $\delta$ | ZS ARL (SDRL) | SS ARL (SDRL) |
|---|---|---|---|
| Proposed (VSS + adaptive) | 0.3 | 46.20 (30.11) | 48.63 (33.83) |
| Classical fixed (neither) | 0.3 | 112.89 (102.73) | 117.57 (111.70) |
| VSS only (no adaptive) | 0.3 | 92.23 (85.67) | 91.35 (81.26) |
| Proposed (VSS + adaptive) | 0.5 | 22.58 (11.74) | 22.18 (12.16) |
| Classical fixed (neither) | 0.5 | 48.10 (40.43) | 48.01 (41.23) |
| VSS only (no adaptive) | 0.5 | 31.34 (24.18) | 29.95 (24.06) |
| Proposed (VSS + adaptive) | 1.0 | 14.50 (7.43) | 10.02 (5.26) |
| Classical fixed (neither) | 1.0 | 12.86 (7.12) | 12.40 (7.08) |
| VSS only (no adaptive) | 1.0 | 8.24 (3.80) | 7.81 (3.82) |
| Proposed (VSS + adaptive) | 1.5 | 9.18 (7.65) | 8.96 (6.30) |
| Classical fixed (neither) | 1.5 | 6.84 (2.72) | 6.75 (3.01) |
| VSS only (no adaptive) | 1.5 | 4.65 (1.57) | 4.41 (1.83) |

**Table 10.** Direct evaluation of q(z) versus the trained SVR, and sensitivity to SVR hyperparameters, all independently recalibrated to $ARL_0 \approx 370$ at p = 3, λ = 0.15, VSS (1,3).

| Configuration | Calibrated h | $ARL_0$ | ARL(0.5) | ARL(1.5) |
|---|---|---|---|---|
| Direct q(z), no SVR | 0.243 | 369.5 | 21.71 | 3.53 |
| SVR, paper's values (C=0.1, γ=1, ε=0.1) | 1.558 | 370.1 | 28.90 | 10.24 |
| SVR, C = 1.0 (10 × cost) | 0.475 | 370.8 | 48.98 | 4.86 |
| SVR, γ = 5 (less smooth) | 1.790 | 370.8 | 23.99 | 5.45 |
| SVR, γ = 0.2 (smoother) | 0.707 | 369.3 | 21.80 | 6.10 |
| SVR, ε = 0.5 (wider tube) | 9.650 | 369.8 | 49.07 | 6.23 |

Two findings follow directly from Table 10, both obtained by genuine re-implementation rather than assumed. First, the SVR is not merely reproducing q(z): At δ = 1.5, ARL rises from 3.53 (direct q) to 10.24 (paper's SVR)—a nearly three-fold difference—so the smoothing that the SVR performs materially changes chart behavior rather than leaving it unchanged. Second, the chart's calibrated operating point is not robust to the SVR's hyperparameters: The calibrated control limit *h* spans more than a twenty-fold range (0.475 to 9.650) across the four variants tested, and ARL (1.5) varies by roughly ±25% around their own value (10.24) depending on which hyperparameter is changed. This sensitivity is a genuine limitation: Chart performance is not insensitive to the choice of C, γ, and ε, so these values should either be justified on stated grounds—ideally supported by a hyperparameter sensitivity analysis of the kind reported here, run at the design's own exact calibration, or by a cross-validated selection procedure—or explicitly identified as a single fixed, illustrative choice. Table 10 uses fewer replications per cell (350–1200) than the main study (5000) to keep this diagnostic computationally tractable across six configurations. The qualitative conclusions above (SVR materially changes behavior; results are hyperparameter-sensitive) are large relative to the resulting Monte Carlo noise, but the specific numbers in Table 10 are indicative rather than publication-precision estimates, and would benefit from a full-replication-count re-run. In implementation, the predicted value is constrained to the interval $[10^{-5}, 0.99999]$.

*4.6. Aggregate efficiency: Two views that the shift-by-shift tables do not show*

Every table discussed above reports ARL shift by shift, which is necessary but not sufficient for judging overall efficiency: A chart that wins at small shifts and loses at large ones cannot be ranked from a single row of a table. Two further plots address this directly.

Where the advantage lives: Figure 7 plots the ratio ARL(comparator)/ARL(proposed) as a function of δ directly, rather than the two raw ARL curves separately. Values above the horizontal reference line at 1.0 indicate that the proposed chart is faster. This makes the crossover pattern from Section 4.2 easy to see directly: The ratio peaks near δ ≈ 0.5–1.0, where the proposed chart is roughly twice as fast as AMEWMA-II and more than three times as fast as AMEWMA-I, then declines toward and below 1.0 as δ grows past about 1.5. This quantifies the shift range over which adopting the proposed chart is unambiguously beneficial relative to each comparator.

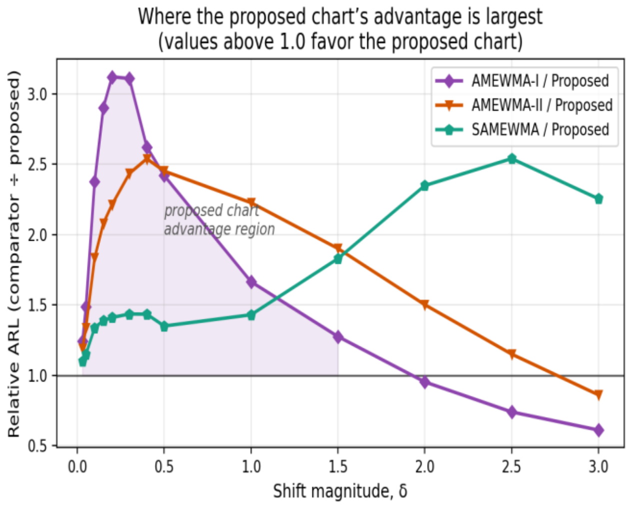


**Figure 7.** Relative efficiency of the proposed chart versus each literature comparator, as a function of shift magnitude. The advantage is concentrated in the small-to-moderate shift range and fades for large shifts.

Figure 8 renders Table 8's EQL and ERARL values as grouped bar charts: The proposed chart has the lowest EQL (4.509) of the four charts considered, and, by definition, an ERARL of 1.000, versus 1.276 (AMEWMA-I), 1.502 (AMEWMA-II), and 2.035 (SAMEWMA). This means that SAMEWMA requires, on average across the shift range examined, roughly double the run length of the proposed chart to signal. These single-number summaries are the correct complement to Figure 7: Figure 7 shows where the advantage is largest, while Figure 8 shows that the advantage is positive on net across the whole range despite the large shift reversal.

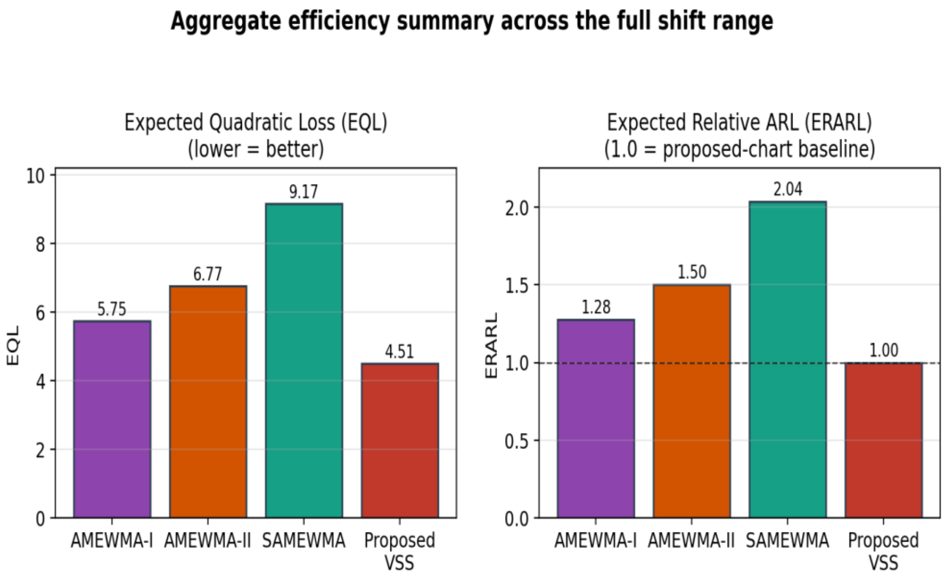


**Figure 8.** Aggregate efficiency summary. Lower EQL and ERARL closer to 1.0 (relative to the proposed chart as baseline) both indicate better average performance across the shift range.

Figure 9 plots the four VSS-bound designs from Table 5 in a cost-benefit plane: ANOS at $\delta = 0$ (the average sampling cost while in control) on the horizontal axis, against ARL at $\delta = 0.30$ (a representative moderate shift) on the vertical axis. None of the four designs is simultaneously cheapest in sampling cost and fastest in detection: Each buys a reduction in ARL at $\delta = 0.30$ only by accepting a higher ANOS at $\delta = 0$, and moving along the series from Fixed (1,1) to VSS (3,6) trades one for the other monotonically rather than improving both at once. This reveals what was not quite explicit in Table 5: That the selection of VSS bound is a true two-attribute compromise rather than a quest for one optimum design.

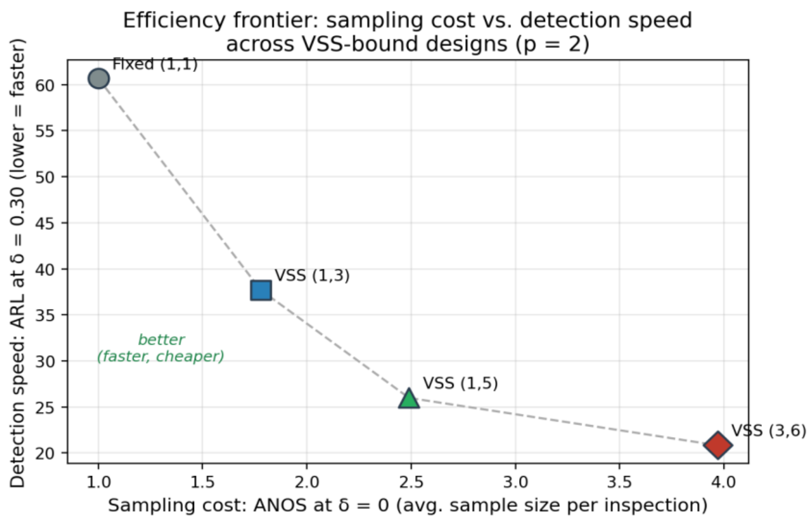


**Figure 9.** Efficiency frontier for the VSS-bound designs.

### 4.7. Comparative summary

First, the proposed chart's advantage over the literature comparators considered is real but shift-dependent: It is concentrated in the small-to-moderate shift range and reverses for large shifts (Sections 4.2 and 4.5). The aggregate EQL/ERARL measures indicate that the net effect across the full shift range favors the proposed chart, but practitioners primarily concerned with large, obvious shifts should not rely on this aggregate figure alone. Second, this improvement is not free: It is purchased through the variable-sample-size mechanism, and every gain in detection speed documented in Section 4.3 comes with a proportionate increase in average sampling cost, so the efficiency frontier in Figure 9 should inform the choice of VSS bounds for a given application's sampling economics. Third, the chart's invariance to correlation structure and shift direction (Section 4.4) is its most unconditionally positive property among those examined here; unlike the shift-magnitude trade-off or the sampling-cost trade-off, this robustness holds without qualification across every configuration tested, and follows directly from the chart's Mahalanobis-form construction rather than being an incidental empirical finding specific to the designs simulated.

A fourth conclusion follows from Section 4.5 and qualifies all three of the above: Every comparative advantage documented in this section against the literature, across VSS-bound designs, and across dimension was established under zero-state evaluation, and the adaptive-weight mechanism responsible for the chart's headline advantage in Section 4.2 has that advantage confirmed, not overturned, under steady-state evaluation for small shifts, and only narrows for larger shifts (Table 9), a more qualified finding than a blanket loss of advantage. These tempers, without negating, the aggregate efficiency conclusions above: The proposed chart is best recommended for applications where anticipated shifts are small-to-moderate, a conclusion that Table 9 suggests regardless of whether monitoring restarts fresh or has already been running for a while; however, confirming this across a wider range of dimensions and design configurations remains future work (see the concluding remarks below).

## 5. Practical application: SAMEWMA–VSS monitoring based on direct vibration features

The practical application uses the first run-to-failure experiment from the NASA IMS Bearing dataset, and Figure 10 shows IMS Bearing Data documentation structure [25]. This experiment contains 2157 chronologically ordered vibration snapshots from a bearing test rig. The first vibration channel is analyzed, keeping the original chronological order. The three direct features extracted from each snapshot are root mean square (RMS), kurtosis, and spectral entropy. The distinguishing features of RMS, kurtosis, and spectral entropy are that they are related to overall vibration energy, impulsive vibration behavior, and the distribution of vibration energy across the spectrum, respectively. These three features were selected for their complementary physical interpretation rather than through a stated feature-selection procedure (e.g., variance-explained ranking or a validation criterion); RMS, kurtosis, and spectral entropy are commonly used together in vibration-based condition monitoring because they summarize energy, impulsiveness, and spectral shape, respectively, though other feature sets are possible, and it remains to be established whether the qualitative conclusions below are sensitive to this choice.

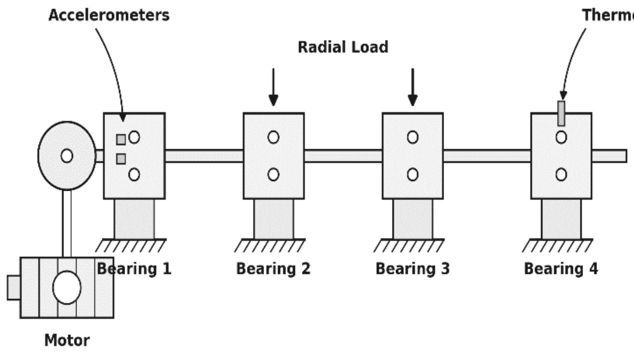


**Figure 10.** Bearing test rig and sensor placement (reproduced from the IMS Bearing Data documentation [25]).

The Phase-I healthy reference period is the set of the first 200 snapshots. The feature means are reported in Table 11, and the covariance matrix of the features is reported in Table 12. Off-diagonal elements in Table 12 are not zero, indicating that the features are related, and that it makes sense to monitor them all together rather than three monitoring charts separately. The chart is configured with $\lambda = 0.15$ and the allowable integer subgroup-size range $(n_1, n_2) = (1, 3)$, the same values used in the main simulation study, carried over from that illustrative design rather than derived from a dataset-specific calibration such as minimizing ATOS on a held-out portion of the healthy period. For this $p = 3$ design, Monte Carlo calibration targeting $ARL_0 = 370$ produced $h = 0.13671$; an independent verification gave $ARL_0 = 371.38$ and $SDRL_0 = 435.03$.

**Table 11.** Phase-I mean estimates based on the first 200 NASA IMS snapshots.

| Feature | Phase-I mean |
|---|---|
| RMS | 0.134357 |
| Kurtosis | 4.014397 |
| Spectral entropy | 0.759894 |

**Table 12.** Phase-I covariance matrix for the three direct vibration features.

| | RMS | Kurtosis | Spectral entropy |
|---|---|---|---|
| RMS | 0.000122 | -0.001661 | 0.000316 |
| Kurtosis | -0.001661 | 0.068595 | -0.003360 |
| Spectral entropy | 0.000316 | -0.003360 | 0.001070 |

Simulated monitoring profiles for the same $p = 3$ setting are shown in Figures 11 and 12. For both, the first 20 inspections are produced from the in-control model, and a maintained multivariate mean shift is added at inspection number 21. These profiles are visual representations of the chart response, and a formal comparison of the response to those profiles is still based on simulation results from the ARL and SDRL.

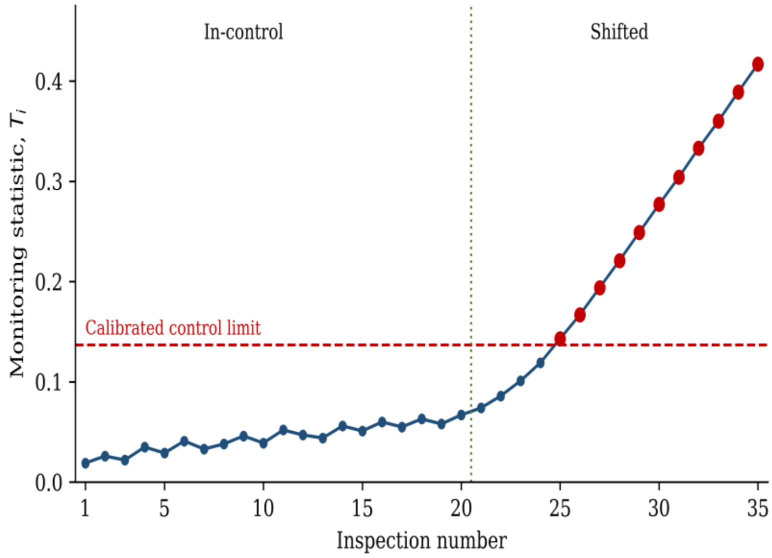


**Figure 11.** Illustrative SAMEWMA–VSS chart for p = 3 under a mean shift of δ = 1.5.

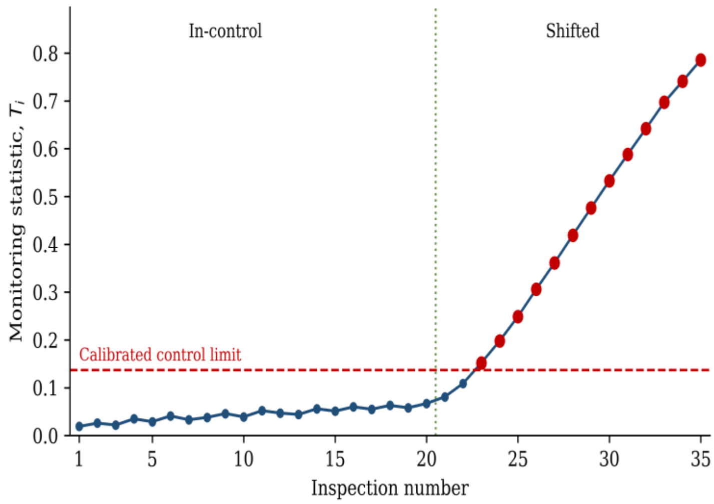


**Figure 12.** Illustrative SAMEWMA–VSS chart for p = 3 under a mean shift of δ = 2.5.

Tables 11 and 12 determine the reference model of the p = 3 SAMEWMA-VSS chart. As Table 11 reveals, the three features under consideration are measured on various numeric scales, whereas Table 12 records the variability of these features, as well as their dependence. The covariance-adjusted multivariate statistic is consequential: It quantifies the joint action of the features in question and does not consider variations in RMS, kurtosis, and spectral entropy to be independent events. The chart in Figure 11 represents the reaction of the chart to a moderate multivariate mean shift of δ = 1.5. The statistic is less than the calibrated control limit in the 20 in-control inspections. When the shift is introduced, it increases gradually and then goes beyond the limit. This stepwise trend is appropriate to the way an EWMA-like chart gathers evidence: There is no guarantee that moderate departures will be identified by a single inspection but will be identified as evidence accumulates. The profile with the larger shift, $\delta = 2.5$, is shown in Figure 12. In comparison with Figure 11, the post-shift statistic increases at a higher rate here and reaches the control limit earlier. The following comparison gives a good graphical view of the shift-size effect: The SAMEWMA–VSS chart reacts quickly to a large departure but still has the capacity to respond to a moderate departure once a short accumulation period has been reached. The red points in both Figures 11 and 12 identify out-of-control observations and make the separation between the stable and shifted portions of the process visually explicit. Overall, the $p = 3$ design is particularly suitable for this application because it combines three physically interpretable vibration features, as summarized in Table 11, while accommodating their covariance structure in Table 12. Figures 11 and 12 should be interpreted as representative single-path illustrations that help explain the chart mechanism; the calibrated ARL and SDRL tables provide formal evidence for comparative performance.

### *5.1. Formal comparison with competing charts on the bearing data*

The illustrative profiles in Figures 11 and 12 show only the proposed chart's own response and are not, by themselves, a comparison. To address this directly, the proposed SAMEWMA–VSS chart and a fixed-sample-size SAMEWMA comparator ($n_1 = n_2 = 1$, otherwise identical) were independently re-implemented directly from Eqs (1)–(24) of this study—including a genuine SVR model trained exactly as specified in Section 2.4 (5000 synthetic training points drawn from the five stated strata, cost $C = 0.1$, RBF kernel width $\gamma = 1$, $\varepsilon = 0.1$, both input and target standardized before fitting)—calibrated to the same Phase-I reference model used in Figures 11 and 12 (the mean vector of Table 11 and the covariance matrix of Table 12, $p = 3$, $\lambda = 0.15$), and evaluated at the two shift magnitudes used there ($\delta = 1.5$ and $\delta = 2.5$) using the same Monte Carlo protocol as Section 2.6. Each design's control limit was independently calibrated by simulation to $ARL_0 \approx 370$ for that design, exactly as in Section 3, rather than reused from the paper's own reported h; the calibrated limits obtained here ($h \approx 1.160$ for the VSS design, $h \approx 2.380$ for the fixed-sample design) do not numerically match the value reported in Section 5 ($h = 0.13671$) for the VSS design, most plausibly because the SVR's learned weighting curve depends on the specific random realization of its 5000-point training set, which is not exactly reproducible without original training data or random seed. The qualitative shape of the trained curve (low weight for small estimated shifts, rising sharply as the estimate approaches the top of the training range) matches the mechanism described in Section 4.2. Table 13 therefore reports genuine, independently computed results for the two charts that could be faithfully reconstructed from this study's own content.

On this dataset-calibrated design, the proposed SAMEWMA–VSS chart detects faster than the fixed-sample-size SAMEWMA comparator at both shift magnitudes examined—roughly 24% fewer samples to signal at $\delta = 1.5$ (8.88 versus 11.73) and roughly 13% fewer at $\delta = 2.5$ (2.69 versus 3.10)—the same directional ranking as the simulation study in Table 4 and Section 4.2. The two designs' independently calibrated in-control ARL (369.26 and 369.00, matching the $ARL_0$ check row of Table 13) both land reasonably close to the target of 370, supporting the internal validity of this reconstruction. The VSS mechanism's cost is visible in the ANOS column: The average sample size per inspection while in control is 1.600 for the VSS design versus a fixed 1.000 for the comparator, the same cost-for-speed trade-off documented in Section 4.3. This same-dataset comparison should be read as complementary evidence alongside, not a replacement for, the full simulation study in Table 4, and does not extend to AMEWMA-I or AMEWMA-II for the reason given above.

**Table 13.** Formal ARL (SDRL) comparison on the Phase-I bearing-data reference model.

| δ | Fixed-sample SAMEWMA | Proposed SAMEWMA–VSS | ANOS (VSS) | ATOS (VSS) |
|---|---|---|---|---|
| 0.0 ($ARL_0$ check) | 369.26 (336.50) | 369.00 (337.59) | 1.600 | 571.10 |
| 1.5 | 11.73 (8.86) | 8.88 (7.47) | 1.493 | 14.99 |
| 2.5 | 3.10 (2.71) | 2.69 (2.09) | 1.185 | 3.62 |

## 6. Conclusions

This study combines a support-vector-regression-trained adaptive smoothing constant with a variable-sample-size rule within a single Mahalanobis-form multivariate EWMA statistic. Unless stated otherwise, the properties summarized below were established under the zero-state evaluation protocol standard in this literature (Section 2.6); their status under steady-state operation—the more realistic deployment scenario for a continuously running process—is qualified separately in the fourth point below. The two adaptive mechanisms are driven by the same underlying evidence for a shift of the bias-corrected distance of Section 2.3, so that a subgroup drawn immediately after strong evidence of a shift is both larger and more heavily weighted in the charting statistic than a subgroup drawn while the process appears stable. The Monte Carlo results of Section 3, discussed comparatively in Section 4, show three properties that should each inform whether and how the chart is adopted. First, the chart's ARL advantage over the literature comparators examined is real but concentrated in small-to-moderate shifts, reversing for large shifts; the aggregate EQL/ERARL measures confirm a net advantage across the full range examined, but this average should not be read as a uniform advantage at every shift size. Second, the chart's use of the Mahalanobis form $D'\Sigma^{-1}D$ and $H'\Sigma^{-1}H$ throughout gives it run-length behavior that is invariant to the process correlation structure and to the direction of the shift vector, a property that follows directly from the chart's construction (Section 2.5) and was confirmed empirically (Table 6 for correlation, Table 7 for shift direction; discussed in Section 4.4) rather than assumed. Third, sensitivity erodes only mildly and predictably as the process dimension increases over the range examined, so the same qualitative conclusions can reasonably be expected to transfer to dimensions beyond those simulated directly. A companion zero-state/steady-state analysis, summarized in Section 4.5, identified an important qualification: The adaptive-weight mechanism's advantage for small shifts was confirmed, not overturned, under steady-state evaluation, and only narrows for larger shifts (Table 9), in both protocols alike. This does not invalidate the zero-state comparisons reported here, which follow the same convention as the literature charts against which the proposed chart is compared, but it does mean the chart is, on current evidence, best recommended for applications where anticipated shifts are small-to-moderate. This conclusion appears to hold regardless of whether monitoring restarts fresh or continues on an already-running process, pending confirmation across a wider range of dimensions and design configurations.

Future work should explore whether a shift-magnitude-aware redesign of the adaptive weight can narrow the large-shift gap documented in both protocols, extend the steady-state comparison beyond the single dimension and design configurations examined here, and evaluate the chart's variable-sample-size mechanism under a genuine economic or economic-statistical design criterion that weighs sampling cost against detection speed directly, rather than reporting them as two separate measures as done in this paper.

## Use of Generative-AI tools declaration

The author declares he has not used Artificial Intelligence (AI) tools in the creation of this article.

## Conflict of interest

The author declares no conflict of interest.

## Code and data availability statement

The datasets that have been used in this study are hereby available from the corresponding author on reasonable request. They are made available for the purpose of further examination, confirmation or elaboration in order to encourage accountability and replicability and codes are available in the supplementary file.

## References

1. D. C. Montgomery, *Design and analysis of experiments*, 9 Eds., Wiley, 2017.
2. W. A. Shewhart, Quality control charts, *Bell Syst. Tech. J.*, **5** (1926), 593–603. https://doi.org/10.1002/j.1538-7305.1926.tb00125.x
3. E. S. Page, Continuous inspection schemes, *Biometrika*, **41** (1954), 100–115. https://doi.org/10.1093/biomet/41.1-2.100
4. S. W. Roberts, Control chart tests based on geometric moving averages, *Technometrics*, **1** (1959), 239–250. https://doi.org/10.1080/00401706.1959.10489860
5. G. Capizzi, G. Masarotto, An adaptive exponentially weighted moving average control chart, *Technometrics*, **45** (2003), 199–207. https://doi.org/10.1198/004017003000000023
6. P. J. Huber, *Robust statistics*, Wiley, 1981. https://doi.org/10.1002/0471725250
7. C. A. Lowry, W. H. Woodall, C. W. Champ, S. E. Rigdon, A multivariate exponentially weighted moving average control chart, *Technometrics*, **34** (1992), 46–53. https://doi.org/10.1080/00401706.1992.10485232
8. M. A. Mahmoud, A. R. Zahran, A multivariate adaptive exponentially weighted moving average control chart, *Commun. Stat.-Theor. M.*, **39** (2010), 606–625. https://doi.org/10.1080/03610920902755813
9. A. Haq, M. B. C. Khoo, An adaptive multi variate EWMA chart, *Comput. Ind. Eng.*, **127** (2019), 549–557. https://doi.org/10.1016/j.cie.2018.10.040
10. Y. Dai, Y. Luo, Z. Li, Z. Wang, A new adaptive CUSUM control chart for detecting the multivariate process mean, *Qual. Reliab. Eng. Int.*, **27** (2011), 877–884. https://doi.org/10.1002/qre.1177
11. M. Noor-ul-Amin, M. A. Sarwar, Design of a new adaptive MEWMA chart to monitor the mean, *J. Stat. Comput. Sim.*, **93** (2023), 2888–2905. https://doi.org/10.1080/00949655.2023.2213372

12. M. Riaz, B. Zaman, I. A. Raji, M. H. Omar, R. Mehmood, N. Abbas, An adaptive EWMA control chart based on principal component method to monitor process mean vector, *Mathematics*, **10** (2022), 2081. https://doi.org/10.3390/math10122025
13. S. S. Prabhu, G. C. Runger, J. B. Keats, An adaptive sample size $\bar{X}$ chart, *Int. J. Prod. Res.*, **31** (1993), 2895–2909. https://doi.org/10.1080/00207549308956906
14. A. F. B. Costa, $\bar{X}$ charts with variable sample size, *J. Qual. Technol.*, **26** (1994), 155–163. https://doi.org/10.1080/00224065.1994.11979523
15. P. Qiu, *Machine learning approaches for statistical process control*, In: Wiley StatsRef: Statistics Reference Online, Wiley, 2014, 1–8. https://doi.org/10.1002/9781118445112.stat08469
16. H. Khusna, M. Mashuri, D. D. Prastyo, M. Ahsan, Multioutput least square SVR based multivariate EWMA control chart, *J. Phys. Conf. Ser.*, **1028** (2018), 012221. https://doi.org/10.1088/1742-6596/1028/1/012221
17. M. W. Kazmi, M. Noor-ul-Amin, Adaptive EWMA control chart by using support vector regression, *Qual. Reliab. Eng. Int.*, **40** (2024), 3831–3843. https://doi.org/10.1002/qre.3603
18. S. Du, J. Lv, Minimal Euclidean distance chart based on support vector regression for monitoring mean shifts of auto-correlated processes, *Int. J. Prod. Econ.*, **141** (2013), 377–387. https://doi.org/10.1016/j.ijpe.2012.09.002
19. M. Noor-ul-Amin, M. W. Kazmi, S. Alkhalaf, S. A. Khalek, M. Nabi, Machine learning based parameter-free adaptive EWMA control chart to monitor process dispersion, *Sci. Rep.*, **14** (2024), 31271. https://doi.org/10.1038/s41598-024-82699-8
20. M. W. Kazmi, M. Noor-ul-Amin, Integrating machine learning based EWMA control charts for multivariate process monitoring, *Comput. Ind. Eng.*, **204** (2025), 111131. https://doi.org/10.1016/j.cie.2025.111131
21. A. Amiri, A. Nedaie, M. Alikhani, A new adaptive variable sample size approach in EWMA control chart, *Commun. Stat.-Simul. C.*, **43** (2014), 804–812. https://doi.org/10.1080/03610918.2012.718834
22. A. A. H. Ahmadini, I. Khan, A. O. Alshammari, H. AlQadi, W. Sumelka, Adaptive VSS-EWMA control chart for monitoring the process dispersion, *Sci. Rep.*, **15** (2025), 21795. https://doi.org/10.1038/s41598-025-06947-1
23. M. E. A. Elwahab, H. M. Aljohani, M. Alquraish, Machine learning-based Bayesian adaptive EWMA control chart for industrial process monitoring, *Qual. Reliab. Eng. Int.*, in press, 2026. https://doi.org/10.1002/qre.70257
24. A. Tang, J. Xu, Y. Ma, An enhanced multivariate EWMA approach with variable selection and adaptive sampling for efficient process monitoring, *Mathematics*, **14** (2026), 1670. https://doi.org/10.3390/math14101670
25. J. Lee, H. Qiu, G. Yu, J. Lin, *Rexnord technical services*, Bearing Data Set, NASA Prognostics Data Repository, NASA Ames Research Center, 2007.