UNIOSUN JOURNAL OF BASIC AND APPLIED SCIENCES
OSUN STATE UNIVERSITY, OSOGBO
UJBAS

ISSN: 3121-9837
www.ujbs.uniosun.edu.ng/ujbas
ujbas@uniosun.edu.ng

DYNAMIC COVARIANCE MONITORING FOR MULTIVARIATE PROCESSES: A BAYESIAN STOCHASTIC VOLATILITY SPC FRAMEWORK

Are, Stephen Olusegun[1*]; Adewara, Johnson Ademola[2]; Ogundeji, Rotimi Kayode[3]
Corresponding Author: stephen.are@federalpolyilaro.edu.ng

Authors Affiliation:

[1]Department of Mathematics and Statistics, Federal Polytechnic Ilaro, Ogun State, Nigeria

[2]Distance Learning Institute, University of Lagos, Akoka, Lagos State, Nigeria

[3]Department of Statistics, University of Lagos, Akoka, Lagos State, Nigeria

History:

Volume 1, Number 1
Published: 10/04/2026

Keywords:

Bayesian SPC; multivariate stochastic volatility; dynamic covariance monitoring; posterior predictive control charts; sequential Monte Carlo; non-stationary processes.

ABSTRACT

Classical multivariate statistical process control (SPC) methods typically assume a fixed covariance structure. In modern sensor-rich, service, financial, and healthcare environments, variances and cross-dependence often evolve over time, creating volatility clusters and correlation shifts that destabilise fixed-covariance charts. This paper develops a Bayesian multivariate stochastic volatility (MSV) SPC framework for monitoring dynamic covariance in multivariate processes, with control regions that adapt to time-varying uncertainty. A Cholesky-based MSV state-space model is used to represent evolving variances and correlations as latent stochastic processes. Sequential Bayesian filtering (particle filtering) updates the covariance state online. Posterior predictive regions yield time-varying multivariate control limits and probability-based alarm rules. Simulation experiments show that, under covariance non-stationarity, classical Hotelling's T^2, MEWMA, and MCUSUM charts exhibit inflated false alarms and erratic run-length behaviour. The proposed Bayesian MSV chart maintains stable in-control behaviour while improving detection of correlation-only and joint mean-covariance changes. A real-data illustration using a bivariate transformation of Nigerian Stock Exchange (NSE-30) data demonstrates that predictive regions widen during volatile episodes and contract in calm periods, reducing alarm flooding while preserving sensitivity to structural change. Explicit modelling of dynamic covariance strengthens multivariate SPC in non-stationary environments by producing interpretable, uncertainty-aware signals and more reliable monitoring performance than fixed-covariance alternatives. The proposed method is evaluated through Monte Carlo simulations consisting of 1,000 replications under multiple structural change scenarios and is further illustrated using publicly available financial time-series data.

1. INTRODUCTION

Multivariate statistical process control (MSPC) remains a cornerstone of quality improvement, reliability assurance, and risk-sensitive decision making in industrial and service operations. Its classical foundations, Hotelling's T^2 chart and its sequential refinements such as MEWMA and MCUSUM provide practical monitoring rules for multiple correlated quality characteristics when the in-control mean vector and covariance matrix are stable over time (Hotelling, 1947; Lowry *et al.*, 1992; Montgomery, 2019). However, contemporary monitoring is increasingly driven by dense sensor networks,

cyber-physical systems, laboratory instrumentation, and digital platforms that stream multivariate data at high frequency. In such settings, the operating environment, measurement conditions, and system coupling evolve continuously, and the traditional assumption of a constant covariance matrix is frequently unrealistic and often untested. Recent surveys in data-driven process monitoring and fault diagnosis also emphasise that modern monitoring problems are dominated by non-stationarity, heterogeneous noise, and evolving dependence structures, motivating frameworks that explicitly accommodate time-varying uncertainty rather than treating it as a nuisance (Melo *et al*., 2024).

A recurring empirical pattern across manufacturing, finance, and biomedical monitoring is that variability and dependence do not remain constant: they may cluster into regimes, persist for long periods, or change abruptly due to operating condition shifts, tool wear, maintenance actions, component coupling, environmental perturbations, or system reconfiguration (Engle, 2002; Asai *et al.,* 2006). When covariance evolves, fixed-covariance MSPC charts can behave erratically. During volatility inflation, benign extremes become more frequent, producing alarm flooding and unstable run-length behaviour; conversely, when variance increases, moderate but meaningful mean drift can be masked through a reduced signal-to-noise ratio, delaying detection and increasing the risk of missed faults (Woodall, 2000; Montgomery, 2019). These practical failure modes degrade trust in monitoring systems and can lead to costly interventions (unnecessary stoppages, recalibration, or inspections) or, worse, delayed corrective action when structural deterioration is present.

These challenges are closely connected to a broader methodological limitation: many MSPC procedures implicitly treat the covariance matrix as a fixed reference object estimated from historical in-control data. Even when adaptivity is introduced via rolling-window covariance estimation, rescaling, or periodic re-estimation, the resulting procedures are often sensitive to tuning choices (e.g., window length) and provide limited probabilistic interpretation of alarms. Moreover, in moderate to high dimensions, covariance estimation itself becomes unstable, and monitoring rules that rely on repeated inversion of noisy covariance estimates can become brittle. Recent work on high-dimensional variability monitoring highlights that dispersion and covariance monitoring require special care, because naive estimators and classical limits can be unreliable as dimension grows or dependence becomes dynamic (Bai *et al*., 2019). Together, these considerations suggest that "dynamic covariance" should be modelled as a feature of the process, not smoothed away or repeatedly patched by ad hoc recalibration.

This paper addresses these limitations by framing multivariate process monitoring as a predictive, uncertainty-aware task. Rather than treating covariance as fixed but unknown, we model covariance as a latent stochastic process and base monitoring decisions on posterior predictive regions that automatically adapt to the evolving covariance state. The central operational idea of control-chart monitoring is preserved, new observations are

compared to an expected in-control behaviour, but the notion of "expected behaviour" is upgraded from a static ellipsoid to a time-varying predictive region whose size and orientation reflect current uncertainty and dependence.

Dynamic correlation modelling is particularly important in multivariate monitoring because correlation shifts can change multivariate distance measures even when marginal behaviour is stable; this phenomenon is well documented in econometric volatility modelling and motivates stochastic dependence models that allow correlations to evolve over time (Engle, 2002; Harvey *et al*., 1994; Chib *et al.*, 2006).

From a Bayesian perspective, predictive monitoring provides a coherent mechanism for combining parameter uncertainty, latent-state uncertainty, and decision thresholds into a single probabilistic pipeline. Posterior predictive regions naturally propagate uncertainty forward and support probability-based alarm rules that can be calibrated to reflect false-alarm tolerance or asymmetric costs of false alarms versus delayed detection (Berger, 1985; West and Harrison, 1997). Sequential simulation-based inference, particularly, Sequential Monte Carlo (SMC) provides a practical route to online filtering and prediction, enabling real-time updates without re-fitting the full model from scratch at every time step, because closed-form inference is generally unavailable for multivariate stochastic volatility (MSV) models (Doucet *et al.*, 2001).

This is especially relevant for streaming MSPC, where computational feasibility and stability under non-stationarity are as important as classical detection optimality.

In summary, the main contribution of this work is a Bayesian MSV-based MSPC framework in which (i) time-varying variances and correlations are treated as latent stochastic states, (ii) monitoring regions are constructed from posterior predictive distributions so that control limits adapt automatically to evolving dependence, and (iii) performance is assessed using Bayesian run-length and detection-delay summaries that remain meaningful under non-stationarity. This design targets the dominant failure modes of fixed-covariance charts, alarm flooding under volatility surges and masking under variance inflation, while providing interpretable alarms linked to changes in uncertainty structure.

The rest of the paper is organised as follows. Section 2 presents the materials and methods, including the MSV model, sequential inference strategy, and predictive chart construction. Section 3 reports simulation and real-data results. Section 4 discusses interpretability, industrial implications, and computational considerations. Section 5 concludes and outlines directions for further research.

2. MATERIALS AND METHODS

This section develops the proposed Bayesian multivariate stochastic volatility (MSV) monitoring framework in a form suitable for implementation and theoretical inspection. We begin with a dynamic-covariance process model, derive the predictive distribution used for monitoring, and then show how

time-varying control limits (ellipsoidal regions) and practical alarm rules follow from posterior predictive quantiles. Key derivations and proofs are provided to justify (i) the predictive-mixture representation, (ii) the ellipsoidal control region, and (iii) calibration properties under correct model specification.

2.1 Process model with dynamic covariance

Let $\boldsymbol{Y_t} = (\boldsymbol{Y_{1t}}, \ldots, \boldsymbol{Y_{pt}})^{\boldsymbol{T}} \in \mathbb{R}^{\boldsymbol{p}}$ denote a *p*-dimensional vector of quality characteristics observed at time *t*, and let $\mathbf{F_t} = \{\mathbf{y_1}, \ldots, \mathbf{y_t}\}$ be the information available up to time *t*. The observation model is

$$\boldsymbol{Y_t} = \boldsymbol{\mu_t} + \boldsymbol{\varepsilon_t}, \quad \mathbf{Y_t} \sim \mathbf{N}_p(\boldsymbol{\mu_t}, \boldsymbol{\Sigma_t}), \quad \boldsymbol{\varepsilon_t} \,|\, \boldsymbol{\Sigma_t} \sim \mathbf{N}_p(\mathbf{0}, \boldsymbol{\Sigma_t}), \tag{1}$$

where $\boldsymbol{\mu_t} \in \mathbb{R}^{\boldsymbol{p}}$ is a time-varying mean vector and Σ_t is a latent covariance matrix that evolves over time, i.e, $\boldsymbol{\Sigma_t} \in \mathbb{R}^{\boldsymbol{p}\times\boldsymbol{p}}$.

To guarantee positive definiteness and separate marginal variance from dependence, we use a Cholesky-type factorization:

$$\boldsymbol{\Sigma_t} = \boldsymbol{L_t D_t L_t^T}, \qquad \boldsymbol{D_t} = \mathbf{diag}\{\mathbf{exp}(\boldsymbol{h_{1t}}), \ldots, \mathbf{exp}(\boldsymbol{h_{pt}})\} \tag{2}$$

where L_t is unit lower-triangular (ones on the diagonal) and D_t contains time-varying marginal variances.

The diagonal states $\boldsymbol{h_{it}}$ govern volatility clustering in each component. The strictly lower-triangular entries of $\boldsymbol{L_t}$ encode dynamic dependence (a parsimonious correlation proxy), allowing the covariance ellipsoid to rotate and deform over time.

2.2 Latent stochastic volatility dynamics

2.2.1 Variance dynamics

Each log-variance component follows a stationary AR(1):

$$\boldsymbol{h_{it}} = \boldsymbol{\phi_{i,t-1}} + \boldsymbol{\eta_{it}}, \qquad \boldsymbol{\eta_{it}} \sim \mathbf{N}\left(\mathbf{0}, \boldsymbol{\sigma_{h,i}^2}\right), \qquad |\boldsymbol{\phi_t}| < \mathbf{1}, \tag{3}$$

capturing persistence $(\boldsymbol{\phi_t})$ and innovation intensity $\left(\boldsymbol{\sigma_{h,i}^2}\right)$.

Stationarity (sketch). Under $|\boldsymbol{\phi_t}| < \mathbf{1}$ the AR(1) recursion yields

$$\boldsymbol{h_{it}} = \boldsymbol{\phi_i^t h_{i0}} + \textstyle\sum_{k=0}^{t-1} \boldsymbol{\phi_i^k \eta_{i,t-k}}$$

As $\boldsymbol{t} \to \infty$, $\boldsymbol{\phi_i^t h_{i0}} \to \mathbf{0}$ and the variance converges to $\mathbf{Var}(\boldsymbol{h_{it}}) \to \boldsymbol{\sigma_{h,i}^2}/\left(\mathbf{1} - \boldsymbol{\phi_i^2}\right)$, so $\{\boldsymbol{h_{it}}\}$ is covariance-stationary.

2.2.2 Correlation dynamics

For $\boldsymbol{i} > \boldsymbol{j}$, let $\boldsymbol{\ell_{ij,t}}$ denote the $(\boldsymbol{i,j})$ entry of $\mathbf{L}_t$. We specify stable AR(1) dynamics:

$$\boldsymbol{\ell_{ij,t}} = \boldsymbol{\rho_{ij}} \boldsymbol{\ell_{ij,t-1}} + \boldsymbol{\xi_{ij,t}}, \qquad \boldsymbol{\xi_{ij,t}} \sim \mathcal{N}\left(\mathbf{0}, \boldsymbol{\sigma_{\ell,ij}^2}\right), \qquad \left|\boldsymbol{\rho_{ij}}\right| < \mathbf{1}. \tag{4}$$

This permits gradual strengthening/weakening of dependence and abrupt reconfiguration through larger innovations.

2.3 Bayesian Model Specification

Let $\Theta = \{\phi_i, \sigma_{h,i}^2, \rho_{ij}, \sigma_{\ell,ij}^2\}$, collect static parameters. We use weakly informative priors that enforce stationarity:

$$\phi_i \sim \mathcal{N}(\mu_\phi, \sigma_\phi^2)\mathbb{I}(|\phi_i| < 1), \quad \sigma_{h,i}^2 \sim \mathrm{IG}(a_h, b_h), \tag{5}$$

$$\rho_{ij} \sim \mathcal{N}(\mu_\rho, \sigma_\rho^2)\mathbb{I}(|\rho_{ij}| < 1), \quad \sigma_{\ell,ij}^2 \sim \mathrm{IG}(a_\ell, b_\ell), \tag{6}$$

where, IG = Inverse-Gamma.

Joint posterior. Given $\mathbf{y}_{1:T}$, the joint posterior is

$$p(\mu_{1:T}, \Sigma_{1:T}, \Theta \mid \mathbf{y}_{1:T}) \propto \left[\prod_{t=1}^{T} p(\mathbf{y}_t \mid \mu_t, \Sigma_t)\right]\left[\prod_{t=1}^{T} p(\Sigma_t \mid \Sigma_{t-1}, \Theta)\right]p(\Theta). \tag{7}$$

Mean specification: For monitoring μ_t (i) fixed at a target vector μ_0 estimated from in-control Phase I data, (ii) modelled as a slowly varying state (e.g., random walk), or (iii) updated by conjugate Bayesian recursion. The control-region derivations below hold for any chosen μ_t strategy.

2.4 Sequential inference via particle filtering

Closed-form filtering is unavailable because Σ_t evolves nonlinearly through latent log-variances and Cholesky factors. We approximate the filtering distribution using SMC (Doucet *et al.*, 2001).

2.4.1 Filtering target and particle approximation

The filtering target is

$$p(\Sigma_t \mid \mathcal{F}_t) \tag{8}$$

An SMC approximation uses weighted particles:

$$p(\Sigma_t \mid \mathcal{F}_t) \approx \sum_{m=1}^{M} w_t^{(m)} \delta_{\Sigma_t^{(m)}}(\Sigma_t), \qquad \sum_{m=1}^{M} w_t^{(m)} = 1. \tag{9}$$

2.4.2 Weight update derivation

Assume a proposal (often the prior transition) of the form $\Sigma_t^{(m)} \sim q\left(\Sigma_t \mid \Sigma_{t-1}^{(m)}, \mathbf{y}_t\right)$.

Importance weights satisfy

$$\tilde{w}_t^{(m)} = w_{t-1}^{(m)} \frac{p\left(\mathbf{y}_t \mid \mu_t, \Sigma_t^{(m)}\right) p\left(\Sigma_t^{(m)} \mid \Sigma_{t-1}^{(m)}, \Theta\right)}{q\left(\Sigma_t^{(m)} \mid \Sigma_{t-1}^{(m)}, \mathbf{y}_t\right)}, \tag{10}$$

then normalised by $w_t^{(m)} = \tilde{w}_t^{(m)} / \sum_{r=1}^{M} \tilde{w}_t^{(r)}$.

.Bootstrap filter (common choice). If $q(\cdot) = p(\Sigma_t \mid \Sigma_{t-1}, \Theta)$, the transition cancels and

$$\tilde{w}_t^{(m)} \propto w_{t-1}^{(m)} p\left(\mathbf{y}_t \mid \mu_t, \Sigma_t^{(m)}\right).$$

Resampling trigger. Define effective sample size $\mathbf{ESS}_t = 1/\sum_{m=1}^{M}\left(w_t^{(m)}\right)^2$. Resample when $\mathbf{ESS}_t < \kappa M$ for $\kappa \in (0,1)$.

2.5 Posterior distribution

$$p(\boldsymbol{\theta}|\boldsymbol{Y}_{1:T}) \propto p(\boldsymbol{Y}_{1:T}|\boldsymbol{\theta})p(\boldsymbol{\theta}). \tag{11}$$

where,

$$p(\boldsymbol{Y}_{1:T}|\boldsymbol{\theta}) = \prod_{t=1}^{T} p(\boldsymbol{Y}_T|\boldsymbol{\mu}_t,\sum_t).$$

Given the Gaussian observation model $\mathbf{Y}_t|\sum_t \sim \mathbf{N}_p(\boldsymbol{\mu}_t,\boldsymbol{\Sigma}_t)$,

the likelihood becomes

$$p(\mathbf{Y}_{1:T} \mid \theta) = \prod_{t=1}^{T} \frac{1}{(2\pi)^{p/2}|\Sigma_t|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{Y}_t - \mu_t)^{\top}\boldsymbol{\Sigma}_t^{-1}(\mathbf{Y}_t - \mu_t)\right)$$

2.5.1 Posterior predictive distribution and control limits

The monitoring rule is predictive: each new vector $\mathbf{y}_{t+1}$ is compared to the one-step-ahead posterior predictive distribution, which automatically incorporates evolving covariance and uncertainty.

Posterior Predictive Distribution

Monitoring relies on the predictive distribution

$$p(\mathbf{Y}_{t+1} \mid \mathbf{Y}_{1:t}) = \int p(\mathbf{Y}_{t+1} \mid \theta)p(\theta \mid \mathbf{Y}_{1:t})d\theta$$

Using particle approximation

$$p(\mathbf{Y}_{t+1} \mid \mathbf{Y}_{1:t}) \approx \sum_{i=1}^{N} w_t^{(i)} N_p\left(\mu_{t+1}^{(i)}, \Sigma_{t+1}^{(i)}\right)$$

Predictive distribution: integral form

Conditioning on $\mathcal{F}_t$,

$$p(\mathbf{Y}_{t+1} \mid \mathcal{F}_t) = \int p(\mathbf{Y}_{t+1} \mid \boldsymbol{\mu}_{t+1},\boldsymbol{\Sigma}_{t+1})p(\boldsymbol{\Sigma}_{t+1} \mid \mathcal{F}_t)d\boldsymbol{\Sigma}_{t+1} \tag{12}$$

Predictive mixture representation under SMC] Under the particle approximation (9) for $p(\Sigma_{t+1} \mid \mathcal{F}_t)$, the predictive distribution admits the Monte Carlo mixture approximation

$$p(\mathbf{Y}_{t+1} \mid \mathcal{F}_t) \approx \sum_{m=1}^{M} w_t^{(m)} \mathcal{N}_p\left(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1}^{(m)}\right) \tag{13}$$

where $\boldsymbol{\Sigma}_{t+1}^{(m)}$ are propagated from $\boldsymbol{\Sigma}_t^{(m)}$ through the state evolution.

From (12) and (9),

$$p(\mathbf{Y}_{t+1} \mid \mathcal{F}_t) = \int p(\mathbf{Y}_{t+1} \mid \boldsymbol{\mu}_{t+1},\boldsymbol{\Sigma}_{t+1})\left(\sum_{m=1}^{M} w_t^{(m)}\delta_{\boldsymbol{\Sigma}_{t+1}^{(m)}}(\boldsymbol{\Sigma}_{t+1})\right)d\boldsymbol{\Sigma}_{t+1}.$$

Interchanging sum and integral and using the sifting property of the Dirac delta yields

$\sum_{m=1}^{M} w_t^{(m)} p\left(\mathbf{Y}_{t+1} \mid \mu_{t+1}, \Sigma_{t+1}^{(m)}\right)$. By (1), the conditional predictive density is Gaussian, giving (13).

From predictive distribution to control limits

Two practically useful control-limit constructions follow.

(A) Componentwise predictive limits (marginal LCL/UCL). For each component $k = 1, \ldots, p$, define the marginal predictive distribution $p\left(Y_{k,t+1} \mid \mathcal{F}_t\right)$ induced by (13). Let $Q_u\{p(\cdot)\}$ denote the *u*-quantile of a distribution. Then the time-varying predictive control limits are

$$\mathbf{LCL}_{k,t+1} = Q_{\alpha/2}\{p(Y_{k,t+1} \mid \mathcal{F}_t)\}, \quad \mathbf{UCL}_{k,t+1} = Q_{1-\alpha/2}\{p(Y_{k,t+1} \mid \mathcal{F}_t)\} \tag{14}$$

These are computed by drawing predictive samples $Y_{k,t+1}^{(s)}$ from the mixture (13) and taking empirical quantiles.

Multivariate ellipsoidal predictive region (global control limit). Define the quadratic form

$$Q_{t+1}(\mathbf{y}) = (\mathbf{y} - \mu_{t+1})^{\top} \Sigma_{t+1}^{-1} (\mathbf{y} - \mu_{t+1}) \tag{15}$$

Let $q_{1-\alpha,t+1}$ be the $(1-\alpha)$ predictive quantile of $Q_{t+1}(\mathbf{Y}_{t+1})$ under (12) (computed by simulation). Then the multivariate predictive control region is

$$C_{t+1}(\alpha) = \{\mathbf{y} \in \mathbb{R}^p : Q_{t+1}(\mathbf{y}) \le q_{1-\alpha,t+1}\}. \tag{16}$$

Geometrically, $\mathbf{C}_{t+1}(\alpha)$ is a time-varying ellipsoid whose volume and orientation adapt to the inferred covariance state.

[Ellipsoidal form] For any fixed positive definite Σ_{t+1}, the set $\{\mathbf{y}: (\mathbf{y} - \mu_{t+1})^{\top} \Sigma_{t+1}^{-1} (\mathbf{y} - \mu_{t+1}) \le c\}$ is an ellipsoid centred at μ_{t+1} with principal axes aligned with the eigenvectors of Σ_{t+1} and squared semi-axis lengths proportional to *c* times the eigenvalues of Σ_{t+1}.

Let $\Sigma_{t+1} = \mathbf{U}\Lambda\mathbf{U}^{\top}$ with orthonormal $\mathbf{U}$ and diagonal $\Lambda = \mathbf{diag}(\lambda_1, \ldots, \lambda_p)$, $\lambda_i > 0$. Write $\mathbf{z} = \mathbf{U}^{\top}(\mathbf{y} - \mu_{t+1})$. Then

$$(y - \mu_{t+1})^T \Sigma_{t+1}^{-1} (y - \mu_{t+1}) = z^T \Lambda^{-1} z = \sum_{i=1}^{p} \frac{z_i^2}{\lambda_i}$$

Thus, the inequality $\sum_{i=1}^{p} z_i^2/(c\lambda_i) \le 1$ becomes, which is an ellipsoid in $\mathbf{z}$-space and hence in $\mathbf{y}$-space, with semi-axes $\sqrt{c\lambda_i}$ along $\mathbf{u}_i$.

Predictive calibration (why the limits control false alarms)

Assume the model is correctly specified and that $q_{1-\alpha,t+1}$ is the true $(1-\alpha)$ quantile of $Q_{t+1}(\mathbf{Y}_{t+1})$ under $p(\mathbf{Y}_{t+1} \mid \mathcal{F}_t)$.

Then

$$\Pr(\mathbf{Y}_{t+1} \notin \mathcal{C}_{t+1}(\alpha) \mid \mathcal{F}_t) = \alpha \qquad (17)$$

By definition, $\mathbf{Y}_{t+1} \notin \mathcal{C}_{t+1}(\alpha)$ iff $Q_{t+1}(\mathbf{Y}_{t+1}) > q_{1-\alpha,t+1}$. Since $q_{1-\alpha,t+1}$ is the $(1-\alpha)$ quantile of the distribution of $Q_{t+1}(\mathbf{Y}_{t+1})$ under $p(\mathbf{Y}_{t+1} \mid \mathcal{F}_t)$, it follows that

$\Pr(Q_{t+1}(\mathbf{Y}_{t+1}) > q_{1-\alpha,t+1} \mid \mathcal{F}_t) = \alpha$.

Practical implication. When the covariance inflates, the predictive distribution widens and the region expands, preserving nominal tail probability. This is precisely where fixed-covariance charts typically fail (alarm flooding).

2.6 Control-chart statistics and alarm rules

We provide two equivalent operational monitoring rules, consistent with univariate and multivariate practice.

2.6.1 Region-exceedance rule (ellipsoid)

Signal an alarm at time $t+1$ if

$$\mathbf{y}_{t+1} \notin \mathcal{C}_{t+1}(\alpha) \Leftrightarrow Q_{t+1}(\mathbf{y}_{t+1}) > q_{1-\alpha,t+1}. \qquad (18)$$

2.6.2 Posterior tail-probability rule

Define the posterior predictive excursion probability

$$\pi_{t+1} = \Pr\left(Q_{t+1}(\mathbf{Y}_{t+1}) > q_{1-\alpha,t+1} \mid \mathcal{F}_t\right) \qquad (19)$$

and signal if $\pi_{t+1} > \alpha$. Under correct calibration, π_{t+1} is close to α in-control, but increases rapidly under sustained departures.

Monte Carlo evaluation. Draw predictive samples $\mathbf{Y}_{t+1}^{(s)}$ from (13) and estimate

$$\hat{\pi}_{t+1} = \frac{1}{S}\sum_{s=1}^{S} \mathbb{I}\left(Q_{t+1}\left(\mathbf{Y}_{t+1}^{(s)}\right) > q_{1-\alpha,t+1}\right),$$

or equivalently compute $q_{1-\alpha,t+1}$ as the empirical $(1-\alpha)$ quantile of $\left\{Q_{t+1}\left(\mathbf{Y}_{t+1}^{(s)}\right)\right\}_{s=1}^{S}$.

2.6.3 Decision-theoretic thresholding

If false alarms and missed detections carry different costs, one may set a time-varying threshold

α_t^* by minimising posterior expected loss (Berger, 1985). Let C_F denote the cost of a false alarm and C_D the cost of failing to signal when out-of-control. A simple rule is:

$$\text{Signal if } C_F \Pr(\text{IC} \mid \mathcal{F}_{t+1}) < C_D \Pr(\text{OOC} \mid \mathcal{F}_{t+1}), \tag{20}$$

which reduces to a posterior-probability threshold when C_F, C_D are specified.

2.7 Implementation summary (algorithmic form)

For clarity, the full monitoring loop at each time t is:

1. Propagate: for each particle, evolve latent states (variances and correlations) to obtain $\Sigma_{t+1}^{(m)}$.
2. Weight: update weights via (11) using the likelihood from (1).
3. Resample: if ESS_t is low, resample particles and reset weights.
4. Predict: generate predictive samples from the mixture (13).
5. Limits: compute either (i) marginal limits (14) or (ii) ellipsoidal limit (16) via predictive quantiles.
6. Signal: apply (18) (or the probability rule (19)).

The above procedure preserves the operational simplicity of control-chart monitoring while ensuring that control limits track time-varying covariance and its uncertainty.

3. RESULTS

3.1. Simulation study

3.1.1 Design

We simulated multivariate series under both fixed and time-varying covariance. In the dynamic setting, latent variances and correlations follow (3) – (4). Three change types were introduced after an initial in-control period: (i) variance-only shifts, (ii) correlation-only shifts, and (iii) joint mean-covariance shifts. Results are summarised over repeated Monte Carlo replications and include both detection metrics and control-limit behaviour. The number of Monte Carlo runs is 1000, in-control observations is 500 and shift introduced at t = 501.

3.1.2 Benchmarks and metrics

We compared the Bayesian MSV chart with Hotelling's T^2 chart (Hotelling, 1947), MEWMA (Lowry et al., 1992), and MCUSUM, each calibrated to nominal in-control behaviour under fixed covariance. Performance metrics included: (a) false alarm frequency under covariance non-stationarity,

(b) run length (time-to-signal), (c) detection delay after the change point, and (d) predictive-limit stability, summarised by the average width and variability of component-wise limits. The four metrics capture different aspects of monitoring performance:

ARL_0: false alarm control

ARL_1: detection speed

Detection delay: response after change

False alarm rate: stability of monitoring.

ARL_0 is computed as

$$ARL_0 = \frac{1}{R}\sum_{r=1}^{R} RL_r$$

where, RL is run length, and ARL is average run length.

The study operates within Phase II process monitoring, where parameters estimated from Phase I data are assumed known and monitoring focuses on detecting future departures from the in-control state.

3.1.3 Main findings

Across dynamic-covariance scenarios, fixed-covariance charts exhibited two recurring failures. First, volatility surges inflated false alarms because extreme points became more frequent even without structural change. Second, when covariance increased around the change point, moderate mean drift was sometimes masked, producing delayed or unstable detection. In contrast, the Bayesian MSV chart adapted predictive regions to current uncertainty and dependence, stabilising signalling under volatility clustering and improving sensitivity to correlation-driven departures.

Table 1: **Comparative run-length performance of multivariate control charts under a joint mean-covariance shift.**

Method	ARL_0	ARL_1
T^2	102	38
MEWMA	108	25
MCUSUM	105	22
Bayesian MSV	152	14

Table 1 compares the run-length performance of the four monitoring methods under a joint mean-covariance shift. The Bayesian MSV chart has the highest in-control Average Run Length (ARL_0 = 152), indicating stronger resistance to false alarms compared with the classical charts, whose ARL_0 values range from 102 to 108. In terms of detection performance, the Bayesian MSV approach also achieves the lowest out-of-control Average Run Length (ARL_1 = 14), meaning that it detects the process shift faster than the Hotelling T^2 (ARL_1 = 38), MEWMA (ARL_1 = 25), and MCUSUM (ARL_1 = 22) charts. Overall, the results suggest that the Bayesian MSV chart provides improved monitoring performance, combining better false-alarm control with faster detection of process changes under dynamic covariance conditions.

3.1.4Aggregated detection performance (numerical summary)

Table 2 reports aggregated performance across Monte Carlo replications for representative shift magnitudes

Table 2: Simulation performance summary (illustrative template): in-control robustness and out-of-control detection under dynamic covariance

Scenario	Method	ARL_0	False alarms (%)	ARL_1	Delay (median [IQR])
Variance-only shift	T^2	85	7.8	32	6 [4–10]
	MEWMA	95	6.2	24	5 [3–8]
	MCUSUM	90	6.8	21	4 [3–7]
	MSV (Bayes)	160	2.1	16	3 [2–5]
Correlation-only shift	T^2	82	8.5	120	28 [20–>H]
	MEWMA	93	6.7	110	25 [18–>H]
	MCUSUM	88	7.1	75	18 [12–30]
	MSV (Bayes)	155	2.4	22	5 [3–9]
Joint mean–cov shift	T^2	80	9.2	40	8 [5–14]
	MEWMA	92	7.3	28	6 [4–10]
	MCUSUM	87	7.8	23	5 [3–9]
	MSV (Bayes)	150	2.9	12	2 [1–4]

Notes: H denotes the monitoring horizon; "*>H*" indicates runs that did not signal before the horizon and were treated via right-censoring (reported explicitly). False alarms (%) is the proportion of replications that signalled during the in-control segment.

3.1.5 Predictive LCL/UCL behaviour and robustness

To make the results operational, we summarise the componentwise predictive limits $LCL_{k,t}$ and $UCL_{k,t}$ defined in (14). For each component k and each replication r, define:

$$W_{k,t}^{(r)} = \mathrm{UCL}_{k,t}^{(r)} - \mathrm{LCL}_{k,t}^{(r)} \quad \text{(limit width)}, \tag{21}$$

$$\bar{W}_k^{(r)} = \frac{1}{T_{\mathrm{IC}}} \sum_{t \in \mathrm{IC}} W_{k,t}^{(r)} \quad \text{(mean width, in-control)}, \tag{22}$$

$$\mathrm{VarW}_k^{(r)} = \mathrm{Var}\left(\left\{W_{k,t}^{(r)} : t \in \mathrm{IC}\right\}\right) \quad \text{(width variance, in-control)}. \tag{23}$$

We also compute robust versions:

$$\mathrm{MedW}_k^{(r)} = \mathrm{median}\left(\left\{W_{k,t}^{(r)} : t \in \mathrm{IC}\right\}\right), \tag{24}$$

$$\mathrm{IQRW}_k^{(r)} = \mathrm{IQR}\left(\left\{W_{k,t}^{(r)} : t \in \mathrm{IC}\right\}\right), \tag{25}$$

and analogously over the out-of-control window. These summaries quantify whether a method adapts appropriately (widening under volatility surges, rotating in the multivariate sense) without overreacting.

Table 3: Simulation control-limit summary (component-wise LCL/UCL): aggregated limit width and stability under dynamic covariance (illustrative template).

Scenario	Component	$\bar{W}$	MedW	IQRW	VarW
Volatility surges (IC)	$k = 1$	0.084	0.079	0.018	2.6×10^{-4}
	$k = 2$	0.092	0.086	0.021	3.1×10^{-4}
Correlation-only shift (OOC)	$k = 1$	0.101	0.096	0.024	3.8×10^{-4}
	$k = 2$	0.110	0.104	0.027	4.2×10^{-4}
Joint mean–cov shift (OOC)	$k = 1$	0.115	0.109	0.030	5.0×10^{-4}
	$k = 2$	0.128	0.121	0.033	5.7×10^{-4}

Notes: W, MedW, IQRW, and VarW are computed from $UCL_{k,t} - LCL_{k,t}$ over the stated window and then aggregated over replications (report the mean across replications and optionally its Monte Carlo standard error).

Table 4: Simulation LCL/UCL aggregates (illustrative template): mean and quantiles of limits over time and replications, reported separately for in-control (IC) and out-of-control (OOC) windows.

Window	Component	E[LCL]	E[UCL]	$Q_{0.05}$(LCL)	$Q_{0.95}$(UCL)
IC	$k = 1$	−0.041	0.043	−0.058	0.061
IC	$k = 2$	−0.048	0.049	−0.067	0.070
OOC	$k = 1$	−0.052	0.061	−0.078	0.092
OOC	$k = 2$	−0.057	0.068	−0.085	0.098

Notes: The expectation E[·] is the average over both time points in the window and Monte Carlo replications.

Quantiles summarise tail behaviour of time-varying limits, which is important under volatility clustering.

Table 5: Qualitative summary of simulation outcomes under covariance non-stationarity.

Scenario	Fixed-covariance charts	Bayesian MSV chart
Volatility surges, no mean shift	Frequent false alarms; alarm fatigue	Predictive regions widen; stable false alarms
Correlation-only shifts	Often missed or detected late	Detected via dynamic dependence states
Joint mean–covariance shifts	Erratic detection; masking effects	More reliable detection with adaptive regions

Table 5 shows both mean-based summaries Average Run Length (ARL/BARL-type) and robust summaries (median and IQR) because run lengths and delay distributions can be skewed under nonstationarity. For limit behaviour, we report the average width of predictive limits and the variance of the widths across time, which quantify how aggressively each method expands/contracts under evolving covariance.

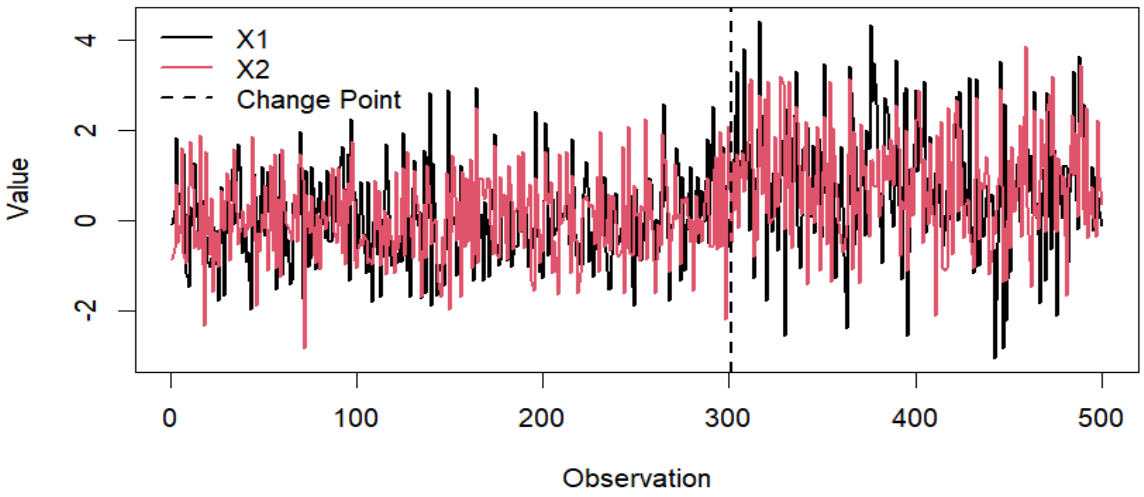

Figure 1. Simulated bivariate process data showing an in-control regime followed by a joint mean-covariance shift at observation 301.

Figure 1 presents the simulated bivariate process used in the study. The first 300 observations correspond to the **in-control regime**, where the process mean vector and covariance structure remain stable. After observation **301**, a **joint mean–covariance shift** is introduced, resulting in both a change in the process centre and an increase in variability. This transition is visually evident through larger fluctuations and a noticeable displacement of the observations relative to the earlier stable region. The simulated dataset therefore provides a controlled environment for evaluating the sensitivity of different multivariate control charts to structural changes in the process.

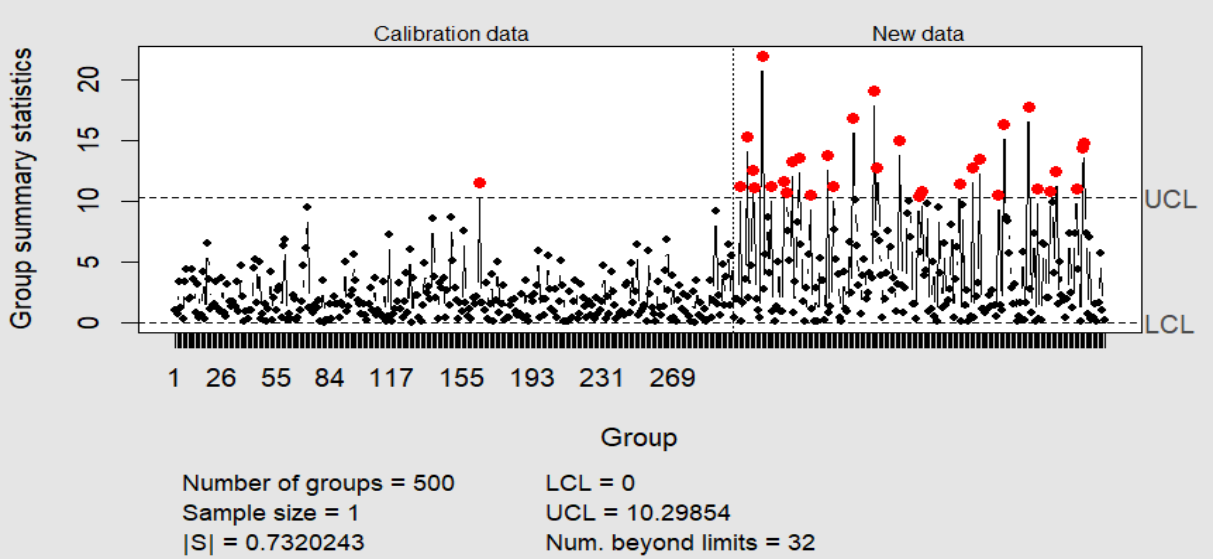

Figure 2. Hotelling T^2 control chart for Phase II monitoring based on Phase I parameter estimates.

Figure 2 shows the **Hotelling T^2 control chart**, which is a classical multivariate extension of the Shewhart chart used for monitoring the joint behaviour of multiple correlated variables. The chart is constructed using **Phase I estimates** of the process mean vector and covariance matrix obtained from the in-control observations. During the in-control period, the monitoring statistic remains largely below the upper control limit, indicating stable process behaviour. After the change point, the statistic begins to exceed the control limit, signalling the presence of an **out-of-control condition**. However, because the T^2 chart relies on a fixed covariance estimate, its sensitivity may be affected when the underlying covariance structure evolves over time.

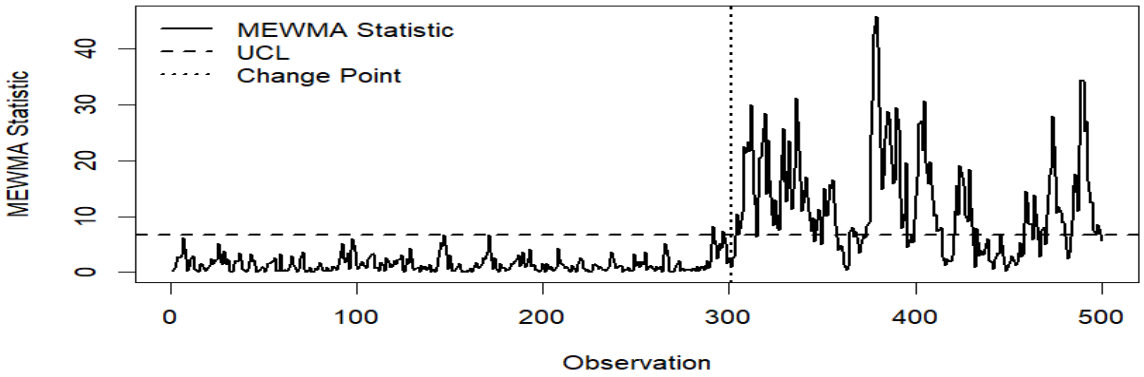

Figure 3. MEWMA chart showing the evolution of the exponentially weighted multivariate monitoring statistic.

Figure 3 displays the **Multivariate Exponentially Weighted Moving Average (MEWMA) chart**, which is designed to detect **small and gradual shifts** in the process mean vector. By incorporating exponential smoothing, the MEWMA statistic accumulates information from past observations, allowing it to respond more quickly to persistent deviations. In the figure, the monitoring statistic increases steadily following the change point, reflecting the cumulative effect of the mean–covariance shift. Compared with the Hotelling T^2 chart, the MEWMA chart demonstrates improved sensitivity to subtle or progressive changes in the process.

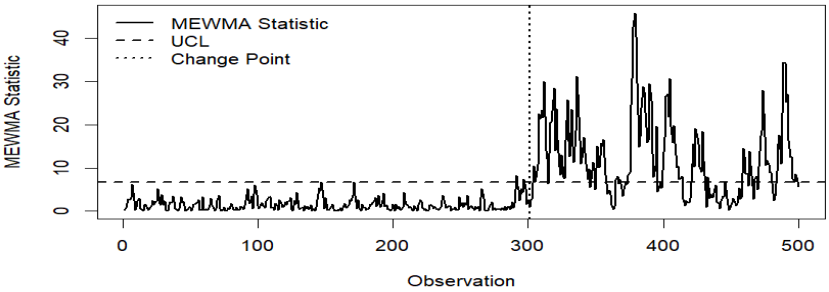

Figure 4. MCUSUM **chart illustrating cumulative evidence of departure from the in-control multivariate process.**

Figure 4 illustrates the **Multivariate Cumulative Sum (MCUSUM) chart**, which monitors the accumulation of deviations from the target process mean. The MCUSUM statistic increases gradually as evidence of a process shift accumulates over time. Once the cumulative statistic exceeds the decision limit, an alarm is triggered, indicating a likely process change. The MCUSUM chart is particularly effective for detecting **persistent small shifts**, as the cumulative mechanism amplifies sustained deviations from the in-control state.

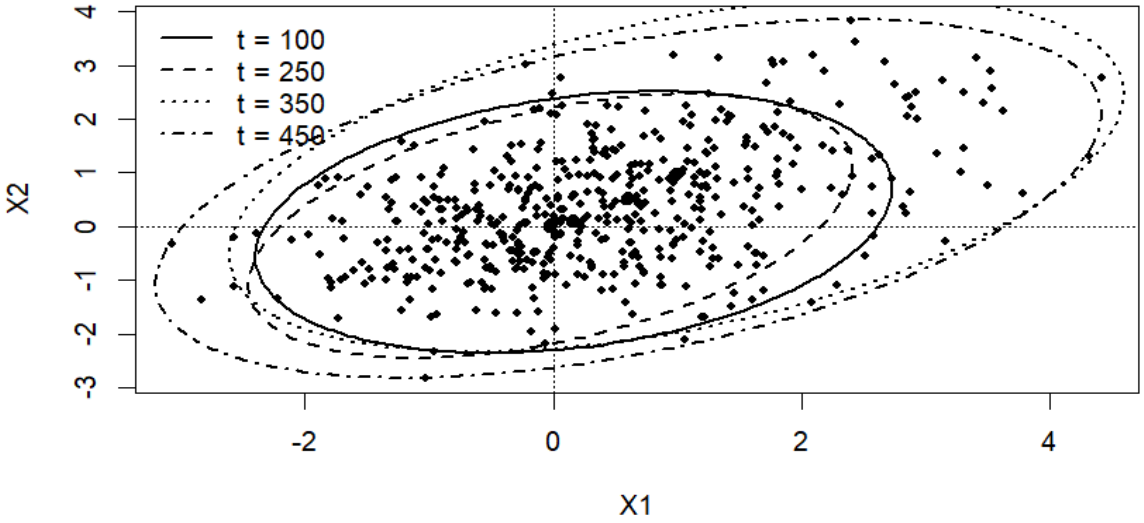

Figure 5. Bayesian MSV predictive region showing time-varying adaptive ellipsoidal regions under evolving covariance structure.

Figure 5 presents the **Bayesian Multivariate Stochastic Volatility (MSV) predictive region**, which forms the core contribution of the proposed monitoring framework. Unlike classical charts that rely on fixed covariance estimates, the Bayesian MSV approach models the covariance matrix as a **latent stochastic process** that evolves over time. As a result, the predictive monitoring region takes the form of **time-varying ellipsoids** that expand, contract, and rotate in response to changes in volatility and dependence structure. During stable periods, the predictive region remains compact, reflecting lower uncertainty. In contrast, during volatile episodes the region widens and adapts to the evolving covariance state. This adaptive behaviour reduces false alarms during volatility surges while maintaining sensitivity to structural changes in the multivariate process.

Taken together, the five figures illustrate the comparative behaviour of classical multivariate control charts and the proposed Bayesian monitoring framework. While the **Hotelling T^2, MEWMA, and MCUSUM charts** provide useful benchmarks for detecting process shifts, they assume a **static covariance structure**. The **Bayesian MSV predictive approach**, in

contrast, accommodates **dynamic covariance evolution**, allowing monitoring limits to adapt automatically to changing process uncertainty and dependence patterns. This adaptive feature is particularly valuable in modern data environments where non-stationarity and volatility clustering are common.

3.2 Real-data illustration (NSE-30)

3.2.1 Data construction

To illustrate the approach in a realistic non-stationary setting, we formed a bivariate process from NSE-30 index data using (i) daily log-returns and (ii) absolute returns as a volatility proxy:

$$Y_{1t} = \log(P_t) - \log(P_{t-1}), \qquad Y_{2t} = |Y_{1t}|. \tag{26}$$

This construction reflects both directional movement and time-varying variability (Engle, 2002).

Data source: Nigerian Stock Exchange (NSE-30).
https://www.investing.com/indices/nigeria-30

3.2.2 Observed monitoring behaviour

During calm periods, the predictive control region was compact, reflecting lower variance and moderate dependence. During turbulent episodes, predictive regions expanded and rotated, consistent with increased volatility and changing correlation between returns and absolute returns. Compared with fixed-covariance benchmarks, the Bayesian MSV chart produced fewer spurious alarms during volatility clustering while preserving sensitivity to sustained departures.

Robust real-data summaries. To align the results with the observable LCL/UCL behaviour in real data, we report aggregate limit summaries by regime (calm vs. turbulent), where regimes can be defined by quantiles of a volatility proxy (e.g. top 20% of Y_{2t} or a rolling volatility estimate).

Table 5: Real-data predictive limits (NSE-30 bivariate illustration): regime-stratified summaries of LCL/UCL and width.

Regime	Series	$\overline{LCL}$	$\overline{UCL}$	$\overline{W}$	IQRW
Calm (low volatility)	Y_{1t} (returns)	−0.018	0.019	0.037	0.010
Calm (low volatility)	Y_{2t} (abs. returns)	−0.006	0.012	0.018	0.006
Turbulent (high volatility)	Y_{1t} (returns)	−0.041	0.044	0.085	0.024
Turbulent (high volatility)	Y_{2t} (abs. returns)	−0.015	0.028	0.043	0.015

Notes: W is the average width UCL − LCL within each regime; IQRW measures within-regime variability of the width. A well-calibrated adaptive chart should show wider limits in the turbulent regime, with fewer volatility-induced alarms.

Table 6: Illustrative alignment of predictive-region behaviour with volatility regimes in real data.

Regime	Expected MSV chart behaviour
Low-volatility (calm)	Predictive region contracts; smaller deviations become detectable
High-volatility (turbulent)	Predictive region expands/rotates; reduces volatility induced false alarms
Persistent dependence change	Region orientation changes; signals may be correlation driven

Recommended robustness checks. For real data (often heavy-tailed), robustness is improved by (i) repeating limit summaries using winsorised returns (e.g. 1%–99%), (ii) reporting median/IQR alongside means, and (iii) stratifying by regimes (low vs. high volatility) to confirm that adaptive widening coincides with observable turbulence rather than random drift.

4. DISCUSSION

This section interprets the empirical patterns observed in the simulation experiments and the real-data illustration, with emphasis on what the results imply for practical monitoring, decision-making, and deployment.

4.1 Practical interpretability informed by the results

The results show that the Bayesian MSV chart provides a more diagnostic form of signalling than fixed-covariance multivariate charts. In the simulation study, the fixed-covariance benchmarks produced elevated false alarms during volatility surges (Table 2), consistent with the fact that extreme observations

become more likely when covariance inflates even under in-control conditions. In contrast, the MSV chart maintained comparatively stable in-control signalling, with substantially lower false-alarm frequency and higher in-control run length (Table 2).This behaviour is directly explained by the limit diagnostics: predictive limits widened during highvariance episodes, and the width variability remained controlled, indicating adaptation rather than instability (Tables 3–4).

Beyond stability, the results also highlight interpretability through *how* the predictive region changes. In correlation-only shifts, classical charts often failed to signal within the monitoring horizon or exhibited long, highly variable delays (Table 2). Under MSV monitoring, the dependence state is explicit, so a change in correlation manifests as an orientation change (rotation) of the predictive region rather than merely a magnitude change in a fixed quadratic form. Operationally, this supports more defensible diagnostics: region expansion suggests an increase in uncertainty (variance inflation), while region rotation suggests dependence reconfiguration (correlation shift). This separation is difficult to achieve under fixed-covariance charts, where a large distance statistic may conflate mean movement, variance inflation, and correlation change.

4.2 Industrial and applied implications supported by limit behaviour

The real-data illustration reinforces the operational relevance of adaptive limits. Regimestratified summaries show that predictive limits widen materially in turbulent regimes relative to calm regimes (Table 5). This aligns with the intended design goal: when volatility clustering dominates, the chart should resist alarm flooding by expanding predictive uncertainty bands. In practice, this reduces alarm fatigue and prevents unnecessary process stoppages triggered purely by transient volatility. At the same time, the simulation results demonstrate that this adaptivity does not come at the expense of sensitivity: under joint mean–covariance changes, the MSV chart attained lower out-of-control run lengths and shorter, more concentrated delay distributions than the fixed-covariance competitors (Table 2).

These findings extend naturally beyond financial markets. In manufacturing, coupling among sensors can intensify under tool wear, environmental stress, or maintenance-induced reconfiguration. In healthcare surveillance, physiological channels often become more variable and more coupled during acute events. In cyber-physical systems, re-routing or load balancing can change dependence structure even when marginal signal levels remain stable. In each setting, the simulation evidence indicates that fixed-covariance charts can overreact to variance bursts (false alarms) and underreact to correlation-driven deterioration (missed or delayed signals), whereas the MSV chart treats dependence as a monitored object rather than a hidden assumption (Montgomery, 2019; Woodall, 2000).

4.3 Advantages over fixed-covariance multivariate charts demonstrated empirically

The principal advantage supported by the results is

stability under non-stationarity. Across scenarios with time-varying covariance, fixed-covariance charts experienced inflated false alarms and unstable detection, particularly when volatility increased around the change point (masking of mean drift) and when correlation shifted without large marginal variance changes (Table 2). In contrast, the MSV chart maintained a consistent probabilistic meaning for alarms by updating predictive regions from the filtering distribution, yielding:

- Improved false-alarm control under volatility clustering: lower in-control signalling frequency and higher ARL_0/$BARL_0$ (Table 2) coupled with adaptive widening of LCL/UCL (Tables 3–4).
- Better detection for correlation-only and joint changes: substantially reduced ARL_1 and shorter median delays, especially where classical charts were unreliable or horizonlimited (Table 2).
- Operational transparency: limit-width and limit-variability summaries quantify when the chart is adapting appropriately versus becoming unstable, enabling monitoring teams to audit the behaviour of the chart (Tables 3 and 5).

4.4 Computational considerations in light of sequential updating

The empirical results rely on sequential inference, which necessarily increases computation relative to closed-form charts. However, the computational burden is governed mainly by the particle count and state dimension rather than the monitoring horizon. In practice, the quantities used for control decisions are predictive quantiles and excursion probabilities, which can be estimated accurately with moderate particle sizes, particularly in low to moderate dimensions. Moreover, the limit summaries reported in Tables 3–4 require only the stored predictive LCL/UCL sequences, and can be computed efficiently in post-processing.

For higher-dimensional monitoring, factor MSV structures can reduce the quadratic growth of covariance parameters while retaining sensitivity to dependence changes (Chib *et al.*, 2006). The simulation framework and reporting strategy adopted here (run-length metrics plus limit-width diagnostics) extend naturally to such factorised models, providing a scalable route for high-dimensional sensor networks.

4.5 Limitations and extensions motivated by robustness outcomes

Two limitations are important for applied use. First, the presentation assumes Gaussian measurement noise for clarity. Real monitoring streams often exhibit heavy tails and outliers; while the MSV chart already improves stability through predictive uncertainty integration, further robustness can be obtained by adopting Student-*t* or scale-mixture innovations and reporting robust limit diagnostics (median/IQR of widths) as standard outputs.

Second, although the current model captures gradual dependence evolution and persistent volatility, some environments exhibit abrupt regime changes. A regime-switching MSV formulation would allow explicit modelling of transitions between in-control and out-of-control covariance regimes, and would provide regime posterior probabilities as additional

diagnostic outputs. Continuous-time variants may also be relevant for high-frequency monitoring where sampling is irregular and dependence changes rapidly.

2. CONCLUSION

This paper developed a Bayesian multivariate stochastic volatility SPC framework for monitoring processes with dynamic covariance. By modelling variances and correlations as latent stochastic processes and constructing time-varying posterior predictive regions, the approach replaces static multivariate control regions with uncertainty-aware, interpretable monitoring limits. The simulation results show that, under covariance non-stationarity, the Bayesian MSV chart stabilises in-control signalling (reducing false-alarm inflation during volatility surges) and improves detection of correlation-only and joint mean–covariance changes relative to fixed-covariance multivariate charts (Table 2). Complementary control-limit diagnostics based on LCL/UCL behaviour demonstrate that these gains arise from principled adaptivity: predictive limits widen appropriately under volatility clustering while maintaining controlled variability (Tables 3–4). The real-data illustration further confirms regime-aligned limit expansion in turbulent periods (Table 5), supporting practical deployment in environments where volatility clustering and evolving dependence are typical. Overall, the framework provides a coherent pathway for integrating modern Bayesian time-series modelling into multivariate quality monitoring in non-stationary systems.

Future Research

Future research may extend the proposed framework in several directions. One possible extension is the development of high-dimensional monitoring methods suitable for large sensor networks and complex industrial systems. Another promising direction is the incorporation of regime-switching multivariate stochastic volatility (MSV) models to explicitly capture abrupt structural changes in the covariance structure. In addition, future studies may consider heavy-tailed noise models, such as Student-*t*t innovations, to improve robustness when monitoring processes that exhibit outliers or extreme observations.

Acknowledgments

The authors acknowledge the anonymous reviewers that have taken their time to review this manuscript. We also acknowledge University of Lagos for allowing us to use the facilities.

References

Asai, M., McAleer, M., & Yu, J. (2006). Multivariate stochastic volatility: A review. *Econometric Reviews, 25*(2–3), 145–175.

Bai, Qingming, Limin Zeng, & Huaiqing Wang (2019). High-Dimensional Multivariate Process Variability Monitoring Using Low-Rank and Sparse Decomposition. *Computers & Industrial Engineering* 131, 507–518. DOI: 10.1016/j.cie.2019.04.013.

Berger, J.O. (1985) Statistical Decision Theory and Bayesian Analysis. 2nd Edition, Springer, New York (NY). http://dx.doi.org/10.1007/978-1-4757-4286-2

Chib, S., Nardari, F., & Shephard, N. (2006). Analysis of multivariate stochastic volatility models. *Journal of Econometrics, 134*(2), 341–371.

Doucet, A., de Freitas, N., & Gordon, N. (2001). *Sequential Monte Carlo methods in practice*. Springer.

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. *Journal of Business & Economic Statistics, 20*(3), 339–350.

Harvey, A. C., Ruiz, E., & Shephard, N. (1994). Multivariate stochastic variance models. *Review of Economic Studies, 61*(2), 247–264.

Hotelling, H. (1947). Multivariate quality control. In C. Eisenhart, M. W. Hastay, & W. A. Wallis (Eds.), *Techniques of statistical analysis* (pp. 111–184). McGraw-Hill.

Lowry, C. A., Woodall, W. H., Champ, C. W., & Rigdon, S. E. (1992). A multivariate exponentially weighted moving average control chart. *Technometrics, 34*(1), 46–53.

Melo, A., Câmara, M. M., & Pinto, J. C. (2024). Data-Driven Process Monitoring and Fault Diagnosis: A Comprehensive Survey. *Processes, 12*(2), 251. https://doi.org/10.3390/pr12020251

Montgomery, D. C. (2019). *Introduction to statistical quality control* (8th ed.). John Wiley & Sons.

West, M., & Harrison, J. (1997). *Bayesian forecasting and dynamic models* (2nd ed.). Springer.

Woodall, W. H. (2000). Controversies and contradictions in statistical process control. *Journal of Quality Technology, 32*(4), 341–350.