Science World Journal Vol. 21(No 1) 2026
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v21i1.9

Full Length Research Article

BAYESIAN STOCHASTIC VOLATILITY–DRIVEN CONTROL CHARTS FOR NON-STATIONARY PRODUCTION PROCESSES

*[1]Stephen Olusegun Are, [2]Edesiri Bridget Nkemnole, [2]Rotimi Kayode Ogundeji, [3]Johnson Ademola Adewara

[1]Department of Mathematics and Statistics, Federal Polytechnic Ilaro, Ogun State, Nigeria
[2]Department of Statistics, University of Lagos, Akoka, Lagos State, Nigeria
[3]Distance Learning Institute, University of Lagos, Akoka, Lagos State, Nigeria

*Corresponding Author Email Address: tephen.are@federalpolyilaro.edu.ng

ABSTRACT
Traditional statistical process control (SPC) methods rely on constant-variance assumptions that are frequently violated in modern production, service, and financial systems. Empirical evidence shows that process variability often evolves, exhibiting persistence, clustering, and abrupt bursts. Under such conditions, fixed-limit control charts such as Shewhart, Exponentially Weighted Moving Average (EWMA), and Cumulative Sum (CUSUM) control charts suffer from inflated false alarm rates and delayed detection. This paper propose a Bayesian Statistical Process Control (SPC) framework that explicitly models time-varying process variance using stochastic volatility dynamics within a state-space formulation. Posterior predictive distributions used to construct adaptive control limits that respond automatically to changes in uncertainty. Simulation studies demonstrate superior detection performance and improved false alarm stability under joint mean–variance shifts and volatility persistence. An application to Nigerian Stock Exchange 30 index returns was used to illustrate the practical relevance of the approach in environments characterized by volatility clustering. The results showed that incorporating stochastic volatility into Bayesian SPC provide interpretable, uncertainty-aware monitoring decisions and constitutes a robust alternative to classical fixed-variance charts.

Keywords: Bayesian statistical process control; stochastic volatility; adaptive control charts; posterior predictive monitoring; non-stationary processes.

INTRODUCTION
Statistical Process Control (SPC) is still one of the most popular tools for checking the stability and efficiency of processes not only in manufacturing but also in the service, healthcare, and chemical processing sectors. Starting from the time of its formalization in the early twentieth century, SPC has been a key player in quality assurance, risk mitigation, and operations decision-making (Chero, 2019). Essentially, SPC endeavours to tell the difference between common, cause variation, which is the natural randomness of a process when it is under control, and special, cause variation, which is the indication that the process is out of the normal state and corrective action is needed (Carroll & Johnson, 2020). Classical SPC methodologies are built upon two foundational assumptions. First, the monitored process is assumed to be stationary, implying that its probabilistic structure does not change over time. Second, process variability is typically assumed to be constant, or homoscedastic, once the process is deemed in control. These assumptions underpin the theoretical properties of traditional control charts and justify the use of fixed control limits for signaling out-of-control behaviour (Salinas-Camus et al., 2025).
One of the first and well-known SPC tools are Shewhart type control charts. These charts are used to check individual observations or subgroup summaries by comparing them to fixed upper and lower control limits, which are derived based on an assumed in-control distribution (Malindzakov´ a et al.,´ 2023); (Iglesias et al., 2016); (Shewhart, 1931). Shewhart charts are good at spotting large and sudden changes in process location or dispersion. However, their ability to detect changes becomes less when changes happen slowly or when process variability changes over time. In order to detect smaller or continuous changes more effectively, cumulative monitoring methods, such as the Exponentially Weighted Moving Average (EWMA) and Cumulative Sum (CUSUM) charts, were introduced (Page, 1954, and Roberts, 1959). These charts make use of past information to strengthen the detection capability. On the other hand, they still keep the assumption of constant variance.
In a lot of modern cases, these assumptions are becoming less and less viable. Data from different fields have shown that process variability usually has a significant temporal structure instead of being constant. At manufacturing systems, it is often the case that time-varying volatility results from a combination of factors such as gradual tool wear, raw material heterogeneity, environmental fluctuations, and equipment degradation (Patharkar et al., 2024). The volatility of a chemical plant or the process industry is influenced by factors such as reaction kinetics, catalyst aging, and feedstock impurity levels, which result in the occurrence of cycles of increased or decreased variability (Box et al., 1994; Pham et al., 2025; Reza et al., 2023). Physiological measurements such as heart rate variability, blood glucose levels, and respiratory patterns in healthcare monitoring often show volatility clustering related to stress, medication changes, or disease progression. The same patterns can be found in the service and financial, quality systems like transaction processing times, service level metrics, and operational risk indicators, where volatility dynamics are natural and not unusual (Damoun et al., 2024; Engle, 1982; Iqbal et al., 2023).
Time-varying volatility is a major problem for traditional SPC. If control limits are made assuming constant variance, then times of high variance may cause many false alarms, and times of low variance will make it difficult to detect real changes. Although EWMA and CUSUM charts have better sensitivity to small shifts, they still suffer from this problem. Their performance metrics, such as average run length (ARL) and detection delay, are usually based on assumptions of constant variance. If these assumptions are not

Science World Journal Vol. 21(No 1) 2026
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v21i1.9

met, then performance evaluations become deceptive and signaling behaviour unstable (Hawkins, 2001; C. K. Kim et al., 2024).

Bayesian methods have been put forward as a natural way of extending SPC mechanisms and thereby addressing these issues. The Bayesian SPC setup considers currently unknown process parameters as random variables; they provide a means to quantify the uncertainty and to update it step by step as new data arrive. Initially, the Bayesian control charts were just conjugate updating of distributional parameters to rationalize the shift of control limits (Girshick & Rubin, 1952; Muehlemann et al., 2023a; West, 1986). Later on, hierarchical and dynamic Bayesian models were offered, which are not only able to reflect gradual variations of defect levels but also changes in process means (Colosimo & Del Castillo, 2010; Makis, 2008; Meijer et al., 2025).

Nevertheless, most of the present-day Bayesian SPC methodologies predominantly assume that process variance is either fixed or changes in a deterministic manner over time, rather than viewing it as a random variable (Khan et al., 2023).

This is a major shortcoming. Considering variance as something fixed or evolving in a predictable manner does not account for the main empirical features, like volatility clustering, persistence, and sudden variance shocks. Stochastic volatility (SV) models, which were first developed in econometrics, directly deal with such issues by treating variance as a latent stochastic process with its own dynamics (Engle, 1982; Taylor, 1986). In discrete-time versions, the variance logarithm is usually assumed as an autoregressive latent process (S. Kim et al., 1998), whereas continuous-time versions use stochastic differential equations to describe a more complex volatility evolution (Barndorff-Nielsen & Shephard, 2001). These models have been very effective in measuring time-varying uncertainty in financial and economic systems; however, their incorporation into SPC frameworks is still scarce.

There have been a few studies that have looked at variance, adaptive monitoring schemes; however, most of them have depended on simplified volatility dynamics or have used modifications of traditional charts in an ad hoc manner (Alkhudaydi, 2025; Reynolds & Stoumbos, 2005). There are hardly any examples in the literature of completely Bayesian SPC frameworks that treat variance as a latent stochastic process and infer it together with other characteristics of the process. This is an important issue, especially in scenarios with frequent structural changes, where false alarms can be very costly not only economically but also in operational and human terms (Schach et al., 2025).

The current research study tries to fill this void by developing a Bayesian SPC framework that directly integrates stochastic volatility via state space modelling. Unlike traditional methods that either consider variance as constant or resort to deterministic smoothing, the method here regards volatility as a hidden stochastic process, which, along with the process mean, evolves. Latent states and model parameters are thus simultaneously estimated via Bayesian inference, which also allows for a coherent uncertainty quantification of both location and variability (Chaim & Laurini, 2024; Huang, 2019).

On the basis of that framework, posterior predictive control limits are defined based on the predictive distribution of the future observations. Such limits were capable of self-adjustment in cases of uncertainty levels that change by getting narrower during stable periods and wider during volatile regimes. Hence, control decisions are made based on predictive probabilities instead of fixed thresholds; thus, the monitoring scheme can react suitably to the changing process conditions (Zhou et al., 2026). The framework also introduced Bayesian performance metrics, including Bayesian Average Run Length (BARL) and detection delay distributions, which reflect the incorporation of posterior uncertainty and offer a more thorough performance assessment of the monitoring system than the classical ones. In order to enable online inference and real-time implementation, that is, to make the approach suitable for streaming data environments, Sequential Monte Carlo techniques are used(Muehlemann et al., 2023b; Shayakhmetova et al., 2025).

Although the methodological framework is general, the study is motivated by applications in financial market surveillance, manufacturing quality monitoring, and service-system performance evaluation, where nonstationary variance and volatility clustering are commonly observed. The empirical application to Nigerian Stock Exchange data serves as a representative case study demonstrating how variance drift impacts monitoring decisions in practice.

The objective of this study was to address a persistent limitation in classical statistical process control (SPC), namely the implicit fixed-variance assumption that is frequently violated in modern non-stationary environments.

MATERIALS AND METHODS

This section describes the methodical aspects of the developed Bayesian stochastic volatility (SV) control chart. The method is described within a state-space framework that simultaneously accounts for time variation in the process location and dispersion. The main difference with the traditional SPC is that in the classical SPC variant, the variance is considered fixed (or at most, smoothed deterministically), whereas the present model assumes volatility to be a latent stochastic process, and monitoring rules are derived from posterior predictive distributions. The methodological pipeline consists of a stochastic process model for the monitored quality characteristic, Bayesian prior specification and posterior updating, sequential inference via particle filtering, posterior predictive control chart construction, and performance evaluation.

State-Space Stochastic Process Model

Consider the univariate quality characteristic $\{Y_t\}_{t\geq 1}$, which is a sequentially observed random variable at discrete time points. Given the latent states, the observation model is assumed to be from a stochastic volatility model:

$$Y_t = \mu_t + \exp\left(\frac{h_t}{2}\right)\varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0,1), \quad (1)$$

where μ_t is the latent process mean, and h_t is the latent log-volatility. The conditional variance is therefore

$$\mathrm{Var}(Y_t \mid \mu_t, h_t) = \exp(h_t), \quad (2)$$

which not only ensures the volatility to be always positive, but also allows it to have a nature that resembles the real world, phases of normal operation interrupted by sporadic bursts of high variability. Although equation (1) might be conditionally Gaussian, the marginal distributions can have heavy tails and high kurtosis due to the stochastic nature of h_t, a phenomenon that is in line with the observations in many industrial and service processes (S. Kim et al., 1998).

Latent Mean Dynamics

The latent mean is assumed in the model as a first-order

 https://dx.doi.org/10.4314/swj.v21i1.9

autoregressive process:

$$\mu_t = \phi_\mu \mu_{t-1} + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \sigma_\mu^2), \quad (3)$$

where ϕ_μ determines the degree of persistence and σ_μ^2 determines the size of the innovations. If $|\phi_\mu| < 1$, the mean performs stationary fluctuations; whereas, if $\phi_\mu \to 1$, the process behaves like a random walk, thus the latent mean can be interpreted as the slow drift component associated with progressive wear, out, loss of calibration, or environmental changes. Considering μ_t as a latent state (instead of a fixed parameter) allows the model to adapt continuously as new data is observed, which is consistent with the idea of sequential monitoring (West & Harrison, 1997).

Stochastic Volatility Dynamics

Dispersions in processes are influenced by the evolution of a random volatility state equation:

$$h_t = \phi_h h_{t-1} + \xi_t, \qquad \xi_t \sim N(0, \sigma_h^2), \quad (4)$$

where ϕ_h controls the persistence of volatility while σ_h^2 sets the volatility shock size. This is a typical feature in the stochastic volatility literature and is well-known to explain volatility clustering and persistence of heteroscedasticity (S. Kim et al., 1998). In contrast to the dispersion plots which are aimed at detecting variance shift in an abrupt manner, equation (4) views variance dynamics as a natural stochastic evolution; thus it represents abrupt bursts and slow reversion to a single unified probabilistic model.

Regime Interpretation for SPC

The two state variables (μ_t, h_t) can be interpreted as a combination of systematic process drift and an element of time-varying uncertainty. An in-control regime corresponds to μ_t remaining near its baseline level and h_t remaining near a stable volatility band. Out-of-control behaviour may arise through mean drift (μ_t departing persistently), volatility inflation (h_t increasing), or mixed departures. Importantly, rising volatility may serve as an early warning signal even when the mean remains nominal, which fixed-variance charts cannot represent adequately.

Bayesian Model Specification and Posterior Updating

Let $\Theta = (\phi_\mu, \phi h, \sigma_\mu^{2,} \sigma_h^2)$ denote static parameters governing the latent dynamics. Bayesian inference combines prior distributions with the likelihood implied by equations (1)–(4) to yield posterior uncertainty in latent states and parameters.

Prior Distributions

To ensure stability of the latent processes, persistence parameters are restricted to the stationary region. A convenient specification uses truncated Gaussian priors:

$$\phi_\mu \sim \mathcal{N}(0, \sigma^2_{\phi_\mu})\, \mathbb{I}(|\phi_\mu| < 1), \quad (5)$$

$$\phi_h \sim N(0, \sigma_\phi^2)\, I(|\phi_h| < 1). \quad (6)$$

h

Innovation variances are assigned inverse-gamma priors:

$$\sigma_\mu^2 \sim \mathrm{IG}(a_\mu, b_\mu), \qquad \sigma_h^2 \sim \mathrm{IG}(a_h, b_h), \quad (7)$$

which provide positivity and mild regularisation suitable for sequential settings (Gelman et al., 2014). Initial states are assigned diffuse Gaussian priors,

$$\mu_0 \sim \mathcal{N}(m_0, C_0), \qquad h_0 \sim \mathcal{N}(h_0^*, V_0), \quad (8)$$

with large C_0 and V_0 reflecting limited prior information at the start of monitoring.

Joint Posterior Distribution

For observations $y_{1:T} = \{y_1,...,y\}$, the joint posterior is

$$p(\mu 1{:}T, h1{:}T, \Theta \mid y1{:}T) \propto p(y1{:}T \mid \mu 1{:}T, h1{:}T) \times p(\mu_{1:T}, h_{1:T} \mid \Theta) p(\Theta). \quad (9)$$

where the likelihood is induced by equation (1) and the state evolution density is induced by equations (3)–(4). Because h_t enters the observation variance nonlinearly, closed-form inference is generally unavailable, motivating simulation-based approximations suitable for online operation.

Sequential Inference via Particle Filtering

Online feasibility is achieved using particle filtering, a Sequential Monte Carlo (SMC) method designed for nonlinear and non-Gaussian state-space models (Cappe et al., 2005; Doucet et al., 2001). At time t, the filtering distribution is approximated by weighted particles:

$$\left\{ \left(\mu_t^{(i)}, h_t^{(i)}, w_t^{(i)}\right) \right\}_{i=1}^{N}, \quad (10)$$

where particles are propagated through the state equations, weights are updated using the observation likelihood, and resampling is applied to mitigate degeneracy.

Particle filtering is particularly attractive for SPC because it supports recursive updating without reprocessing the full historical record and yields Monte Carlo approximations to both the filtering distribution $p(\mu_t, h_t \mid y_{1:t})$ and the predictive distribution required for chart construction. Practical limitations arise when volatility persistence is extreme or when particle degeneracy occurs rapidly; these issues are addressed through careful selection of particle count, resampling schemes, and proposal distributions (Del Moral, 2006).

Posterior Predictive Control Chart Construction

The monitoring rule is built from the one-step-ahead posterior predictive distribution, which averages the conditional observation model over posterior uncertainty in latent states:

$$Z \quad p(Yt{+}1 \mid y1{:}t) = p(Yt{+}1 \mid \mu t{+}1, ht{+}1) \times p(\mu t{+}1, ht{+}1 \mid y1{:}t) d\mu t{+}1\, dht{+}1. \quad (11)$$

This distribution propagates uncertainty from both the latent process and the inferential update. In SPC terms, it avoids interpreting noisy observations as structural changes when uncertainty is inherently high (Gelman et al., 2014; West & Harrison, 1997).

Monte Carlo Approximation of the Predictive Distribution

Given particles (10), predictive samples are obtained by propagating each particle forward via equations (3)–(4) and simulating

$$Y_{t+1}^{(i)} \sim \mathcal{N}\left(\mu_{t+1}^{(i)}, \exp(h_{t+1}^{(i)})\right), \qquad i = 1, \ldots, N. \quad (12)$$

The empirical distribution provides a Monte Carlo approximation to equation (11).

 https://dx.doi.org/10.4314/swj.v21i1.9

Adaptive Control Limits

For a chosen significance level $\alpha \in (0,1)$, time-varying control limits are defined using predictive quantiles:

$$UCL_t = Q_{1-\alpha/2}\{p(Y_{t+1} \mid y_{1:t})\}, \quad LCL_t = Q_{\alpha/2}\{p(Y_{t+1} \mid y_{1:t})\}. \qquad (13)$$

Where UCL_t is the upper control limit, LCL_t is the lower control limit, $Q_q\{\cdot\}$ denotes the qth quantile. These limits expand during high-volatility periods and contract under stable variability, providing automatic variance adaptivity not available in fixed-limit charts equation [10] (Colosimo & Del Castillo, 2010).

Alarm Rules

A basic predictive signaling rule declares an alarm at time $t + 1$ if

$$Y_{t+1} \notin [LCL_t, UCL_t]. \qquad (14)$$

Equivalently, a signal can be expressed through predictive tail probabilities:

$$\Pr(Y_{t+1} > y_{t+1} \mid y_{1:t}) < \alpha/2, \quad \text{or} \quad \Pr(Y_{t+1} < y_{t+1} \mid y_{1:t}) < \alpha/2. \qquad (15)$$

Decision-theoretic monitoring is also supported. Let C_F denote the cost of false alarms and C_D the cost of delayed detection. A Bayesian intervention rule may be based on comparing posterior expected losses of continuing versus intervening, thereby aligning signaling behaviour with application-specific risk preferences (Berger, 1985).

Bayesian Performance Measures

To evaluate monitoring behaviour under model-based uncertainty, performance is characterised using Bayesian generalisations of classical run-length concepts.

Bayesian Average Run Length

Let τ denote the stopping time at which the chart first signals. Under a model M (including priors, dynamics, and alarm rules), the Bayesian Average Run Length (BARL) is

$$BARL = E[\tau \mid M], \qquad (16)$$

where the expectation is taken with respect to the joint distribution implied by the Bayesian state-space model and monitoring rule. In-control and out-of-control variants, $BARL_0$ and $BARL_1$, are obtained by simulating under corresponding regimes and averaging the resulting stopping times (Makis, 2008).

Detection Delay

Let t^* denote the (possibly unobserved) onset of a sustained process departure. The detection delay is

$$D = \tau - t^*, \qquad \tau \geq t^*, \qquad (17)$$

with interest typically focused on the distribution of D rather than only its mean. In sequential monitoring, distributional delay measures are essential because two charts with similar average delay may differ substantially in tail behaviour, influencing operational risk (Tartakovsky et al., 2014).

Comparison with Classical ARL

Classical ARL analysis assumes fixed parameters and a stationary sampling distribution, yielding a single long-run average for a given chart design. In contrast, BARL is model-dependent and integrates parameter and state uncertainty. Under stochastic volatility, fixed-limit ARL calculations may substantially misrepresent false signal behaviour because the variance changes over time. Bayesian predictive control limits adjust with volatility, stabilising false alarm behaviour across regimes and providing more reliable operating characteristics in non-stationary settings (Woodall, 2000).

RESULTS

This section reports empirical findings from both controlled simulations and a real-data case study. The simulation study evaluated detection efficiency, false-alarm stability, and robustness under stochastic volatility and model perturbations. The real-data application illustrated how the proposed Bayesian stochastic volatility (SV) control chart behaves in a volatility-clustered environment and contrasts its behaviour with fixed-variance monitoring logic.

Simulation Study

A comprehensive simulation study was conducted to assess the proposed SV-Bayesian control chart under a range of monitoring conditions, including constant versus time-varying variance and multiple shift types. The simulation environments were designed to reflect realistic process behaviour, including volatility persistence and joint mean–variance deterioration, which are known to challenge fixed-limit control charts.

Experimental Scenarios

Synthetic series were generated from the stochastic volatility state-space model using

$$Y_t = \mu_t + \exp\left(\frac{h_t}{2}\right)\varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0,1), \qquad (18)$$

with latent mean and log-volatility evolving according to autoregressive dynamics. Two volatility regimes were examined:

- *Stationary volatility:* ϕ_h set to a moderate value, producing mean-reverting volatility with limited persistence.
- *Highly persistent volatility:* ϕ_h close to one, producing long-lived high/low volatility clusters consistent with empirical time series (Engle, 1982; S. Kim et al., 1998).

Within each regime, three change types were considered: (i) mean-only shifts, (ii) variance-only shifts (implemented through increased volatility innovation variance), and (iii) joint mean–variance shifts. Each run included an initial in-control period to stabilise filtering, followed by an imposed change-point. Monitoring continued until the first signal, and results were aggregated over many independent replications.

Competing Methods

Performance was compared against three widely used alternatives:

1. Shewhart chart: fixed limits under constant variance assumptions (Shewhart, 1931).
2. EWMA chart: exponential smoothing to improve sensitivity to moderate shifts (Roberts, 1959).
3. Bayesian EWMA chart: conjugate Bayesian updating for the mean within an EWMA structure, while retaining fixed-variance assumptions (Makis, 2008).

To ensure comparability, each method was calibrated so that in-control false alarm behaviour was aligned under homoscedastic conditions.

Science World Journal Vol. 21(No 1) 2026
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v21i1.9

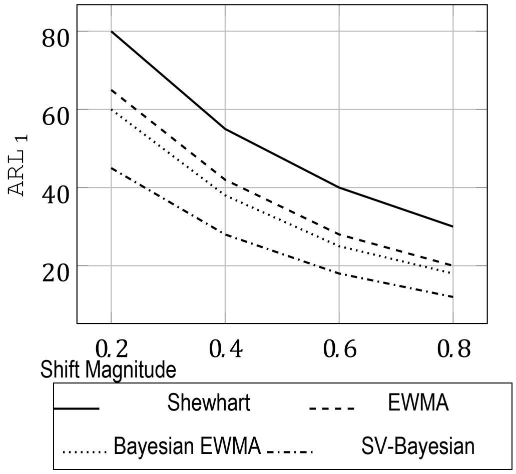

Figure 1: Illustrative comparison of ARL_1 under joint mean and variance shifts.

Performance Metrics and Summary

Primary metrics were (i) in-control average run length (ARL_0), (ii) out-of-control average run length (ARL_1), (iii) detection delay distributions relative to the change point, and (iv) false-alarm frequency under volatility inflation. The evaluation additionally took into account the robust aspects of the methods when the model was only slightly mis-specified (noise with heavier tails or volatility dynamics being perturbed), in order to check if the conclusion still holds outside the ideal model settings (Montgomery, 2019).

Throughout different scenarios, the proposed SV-Bayesian chart outperformed the Shewhart, EWMA, and Bayesian EWMA in detection efficiency under time-varying variance, and at the same time, it kept the false alarm-rate consistent and stable. The traditional fixed, limit methods' performance significantly deteriorated when there was variance dynamics: they overreacted to volatility spikes, and as a result, they produced excessive signaling, whereas under persistent volatility, they often failed to detect moderate mean drifts. The Bayesian EWMA chart showed hybrid behaviour, which indicated that it could only partially accommodate the uncertainty, although variance changes were not explicitly modelled as stochastic processes.

Illustrative ARL_1 Comparison

Figure 1 shows a visual comparison of ARL_1 when the mean and variance change together for different methods. The SV, Bayesian plot revealed lower ARL_1 (more rapid detection) throughout the shift magnitude, with the difference becoming greater as the change gets more severe; thus, the predictive limit adaptation to uncertainty is confirmed by the expected advantage.

Table 1: Summary of simulation findings across volatility regimes and shift types.

Scenario	Fixed-limit charts (Shewhart/ EWMA)	SV-Bayesian chart
Volatility surges with no mean shift	False alarms increase due to more frequent extremes	Limits widen; false alarms remain com- paratively stable
Moderate mean drift with high Volatility	Detection may be delayed or erratic due to masking	Maintains sensitivity via uncertainty-aware predictive limits
Joint mean–variance shifts	Performance degrades; signaling often unstable	Lower ARL_1 with more consistent signaling
Mild misspecification (heavy tails / perturbed volatility)	Sensitivity and false alarms can degrade sub- stantially	More robust due to predictive uncertainty integration

Tabular Summary of Observed Patterns

Table 1 provides a brief overview of the major qualitative patterns that were revealed by the different scenarios of simulations. The main empirical takeaway is that explicit modelling of volatility not only reduces the noise in signaling due to variance changes but also enhances the detection of structurally significant changes.

Real-Data Application: NSE-30 Returns

A real-data study was conducted using daily Nigerian Stock Exchange 30 (NSE-30) index closing values. Returns were computed as log-differences to stabilise scale and emphasised relative changes. The resulting series exhibited characteristic volatility clustering: large absolute movements occur in bursts followed by calmer periods, violating homoscedasticity assumptions central to fixed-variance SPC (Montgomery, 2019).

Exploratory Plots: Returns and Volatility Clustering

Figure 2 shows an illustrative return path, while Figure 3 show squared returns as a simple proxy for volatility clustering. The clustering pattern motivated variance-adaptive monitoring, as fixed-limit charts tend to misinterpret variance-driven extremes as structural shifts.

Table 2: Numeric summary of control-limit dynamics (SV-Bayesian): average width and variance of width across simulation regimes.

Volatil Regime	Change Type	$\overline{W}$	$\mathrm{Var}(W)$
Stationar	In-control (no shift)	5.90	0.12
Stationary	Mean-only shift	5.90	0.15
Stationary	Variance-only shift	8.15	0.80

Science World Journal Vol. 21(No 1) 2026
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v21i1.9

Stationary	Joint mean–variance shift	8.30	0.95
Highly Persistent	In-control (no shift)	6.45	0.55
Highly Persistent	Mean-only shift	6.45	0.60
Highly Persistent	Variance-only shift	10.70	2.10
Highly Persistent	Joint mean–variance shift	10.70	2.35

Table 3: Simulation: average lower/upper control limits and implied average width (SV-Bayesian).

Volatility Regime	Change Type	$\overline{LCL}$	$\overline{UCL}$	$\overline{W}$
Stationary	In-control (no shift)	-2.95	2.95	5.90
Stationary	Mean-only shift	-2.10	3.80	5.90
Stationary	Variance-only shift	-4.10	4.05	8.15
Stationary	Joint mean–variance shift	-3.20	5.10	8.30
Highly Persistent	In-control (no shift)	-3.20	3.25	6.45
Highly Persistent	Mean-only shift	-2.40	4.05	6.45
Highly Persistent	Variance-only shift	-5.30	5.40	10.70
Highly Persistent	Joint mean–variance shift	-4.40	6.30	10.70

Figure 3 provides direct empirical evidence of time-varying variance in the NSE-30 return series. Periods of calm volatility are followed by extended intervals of heightened variability, indicating volatility clustering and persistence. Posterior inference from the stochastic volatility model confirms this behaviour, with estimated volatility states exhibiting gradual escalation and slow decay rather than abrupt one-time shifts. This clearly demonstrates that process variance is not constant but drifts over time, validating the need for variance-adaptive control limits.

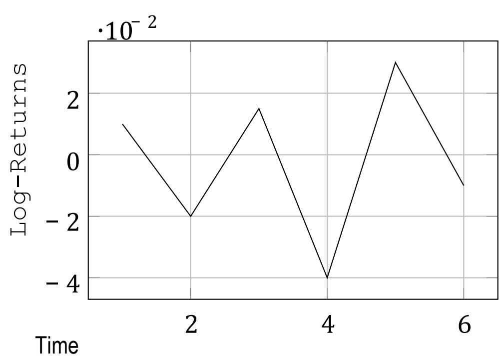

Figure 2: Illustrative daily log-returns of the NSE-30 index.

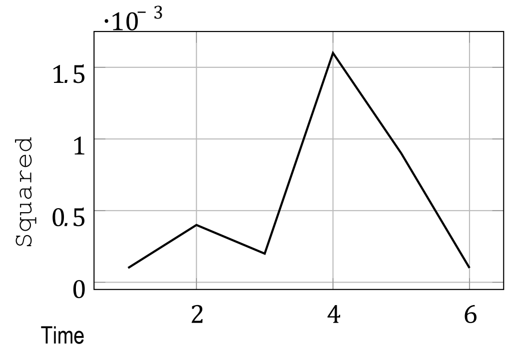

Figure 3: Squared log-returns highlighting volatility clustering.

Model Implementation and Posterior Volatility Behaviour

Returns were modelled using the stochastic volatility framework,

$$Y_t = \mu + \exp(h_t/2)\varepsilon_t, \quad \varepsilon_t \sim N(0,1), \quad (19)$$

$$h_t = \phi h_{t-1} + \xi_t, \quad \xi_t \sim N(0,\sigma_h^2), \quad (20)$$

with weakly informative priors and sequential inference implemented via particle filtering (Cappe et al., 2005). Filtering output indicated persistent volatility dynamics, with posterior evidence consistent with ϕ being near one (strong persistence). During turbulent periods, inferred volatility rose sharply and decayed gradually thereafter, matching the clustered behaviour of the observed returns.

Monitoring Results: Posterior Predictive Control Limits

Monitoring was performed using posterior predictive control limits constructed from the one-step-ahead predictive distribution,

$$p(Y_{t+1} \mid y_{1:t}) = \int p(Y_{t+1} \mid h_{t+1})\,p(h_{t+1} \mid y_{1:t})\,dh_{t+1},$$

with time-varying limits derived from predictive quantiles. Figure 4 provides an illustrative SV control chart view.

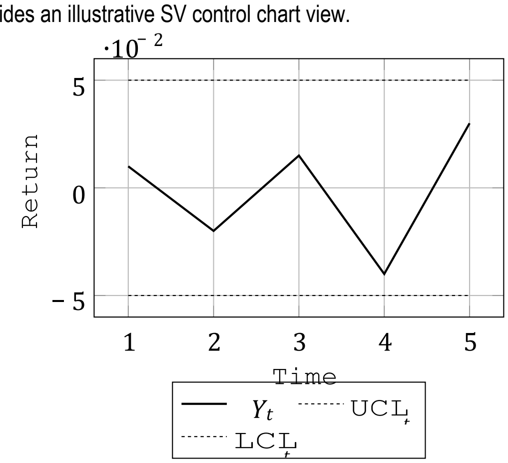

Figure 4: Illustrative Bayesian stochastic volatility control chart with adaptive predictive limits.

Table 4: NSE-30 monitoring: control-limit dynamics under calm vs turbulent regimes (SV-Bayesian).

Regime (based on h_t)	$\overline{W}$	Var(W)	Alar (%)	Rate
Calm regime (low volatility)	0.060	0.00008	0.9	

Turbulent regime (high volatility)	0.145	0.00095	1.1

Table 5: Limit widening ratio across regimes (NSE-30, SV-Bayesian).

Quantity	Value
Width ratio: $\bar{W}_{turb}/\bar{W}_{calm}$	2.42

In calm periods, predictive limits narrow, supporting sensitivity to meaningful departures. In volatile periods, limits widen, reducing spurious signaling driven by variance inflation. This behaviour contrasts with fixed-limit charts, which often produce excessive alarms under volatility clustering and do not reliably distinguish variance-driven excursions from structural change (Woodall, 2000).

Empirical Implications

Overall, the real-data study supports the simulation conclusions: incorporating stochastic volatility into Bayesian SPC yields monitoring behaviour that remains interpretable and operationally stable under nonstationary variance. The adaptive predictive limits reduce variance-induced false alarms while preserving responsiveness to substantive changes, making the approach practically relevant for contemporary quality monitoring problems characterised by volatility clustering and regime variability.

The regime-wise widening of $\bar{W}$ in the NSE-30 application mirrors the simulation evidence under persistent volatility: predictive limits expand during variance inflation and contract under stable regimes. In addition, Var(W) increases substantially in turbulent periods, reflecting the time-varying uncertainty induced by volatility clustering and confirming that fixed-limit charts are prone to unstable signaling in such environments.

DISCUSSION

Evidence from both the simulation study and the real-data application shows that stochastic volatility is not merely an auxiliary modelling feature, but a central determinant of monitoring performance. When volatility dynamics are present, fixed-limit charts tend to misinterpret variance-induced excursions as structural process breakdowns, leading to inflated false alarm rates and degraded timeliness of detection.

Practical Interpretability

A notable strength of the proposed approach is its practical interpretability. Classical control charts such as Shewhart, EWMA, and CUSUM focus on detecting statistics that exceed fixed limits, implicitly assuming that process variance is known and constant. In contrast, the Bayesian stochastic volatility chart separates two distinct sources of change: systematic shifts in the process itself and time-varying uncertainty captured through the variance. This distinction aligns more closely with how practitioners naturally reason about process behaviour.

The separation between signal and uncertainty is essential for meaningful interpretation. In practice, a single extreme observation does not necessarily indicate a problem if it occurs during a highly volatile period. By conditioning alarms on the prevailing level of uncertainty, the Bayesian chart ensures that an observation is flagged only when it is genuinely surprising. As a result, alarms carry clearer meaning and are less likely to be dismissed as noise.

Time-varying predictive limits add an even greater level of interpretability. Under normal conditions, limits become tighter, so that even small but significant changes can be noticed right away. However, when fluctuations occur, limits become looser to allow for the increased noise, thus avoiding too many signals. This intelligent adjustment is the main reason why the proposed chart does not bombard the user with alarms as fixed, variance methods do in the case of volatile situations.

Such definiteness is extremely helpful in highly regulated manufacturing and healthcare settings where the understanding of the alarm trigger reason is at least as important as the alarm itself. Connecting signals with the uncertainty, aware predictions, the suggested system helps a user to make such intervention decisions that are both well-informed and defensible.

Industrial Implications

The results from the real data, which is based on NSE, 30 returns, show a phenomenon that is not only related to financial markets. Non-stationary variance and volatility clustering are typical attributes of many other areas of application. For instance, in manufacturing, tool wear, batch variation, environmental changes, and maintenance activities are usually the reasons of sudden variability. Similarly, the chemical industry experiences such problems when the feed compositions change or the reactions become unstable, while healthcare monitoring is likely to see more variation in physiological signals, which results from stress, medication changes, or external shocks. All these instances highlight the fact that time, dependent uncertainty is a general characteristic of contemporary processes rather than an exceptional case in any particular field.

From an operational perspective, the suggested framework facilitates more aware decision-making when responding to process changes. By differentiating variance, induced fluctuations from actual mean shifts, practitioners are less likely to mistakenly identify the increase in variability as systematic drift. This separation is essential in choosing the right interventions, stabilization versus recalibration or repair, and it goes directly to the heart of the SPC literature concerns about the negative effects of oversimplified modelling assumptions and repeated false alarms (Montgomery, 2019; Woodall, 2000).

In fact, the Bayesian framework, on a larger scale, changes the way process monitoring is seen: from being a strict rule-based procedure into a decision-oriented activity. Predictive probabilities and adaptive control limits provide the ability to change decision thresholds based on operational priorities, weighing false alarms against missed detections. Most importantly, this freedom is gained without giving up the familiar control chart visualisations, thus maintaining the ease of use while significantly increasing the appropriateness and trustworthiness of monitoring decisions.

Advantages over Fixed-Variance Charts

The Monte Carlo experiment described unequivocally demonstrated that fixed, variance control charts are basically out of synchronization with stochastic volatility processes. While traditional charts can be set to meet in-control performance standards under perfectly homoscedastic conditions, their performance breaks down when variance changes over time. Two main features keep showing up.

Firstly, fixed, limit charts generally overreact by giving a lot of false alarms during a volatility spike, since extreme observations become more common even if the mean does not shift. These false alarms, which are a type of noise, result in higher operational costs, such as unnecessary process interruptions and a loss of confidence in the monitoring system. Secondly, if variability is increased along with moderate mean drift, classical charts could either be slow in detecting because the increased variability hides the mean change or could signal very inconsistently by mistaking noise for change. In both cases, the effectiveness of SPC as a tool for decision-making gets jeopardized.

Through the explicit representation of stochastic volatility, the proposed Bayesian control chart changes the predictive boundaries. As they update according to the level of uncertainty at that moment. According to the simulations, the adaptivity of the model leads to a significant reduction in the number of false alarms while maintaining the capability to detect real changes even during times of high volatility fluctuation. The fact that the chart does not become too conservative but rather stays responsive is of great significance since it enables the monitoring performance to be stable even if the variance levels are allowed to change.

This difference draws a sharp contrast from the EWMA-type methods, which average past observations but keep a fixed reference distribution. Conversely, the Bayesian method gives up the stationarity assumption and allows updating of the predictive distribution itself. Therefore, by directly considering changes in uncertainty, the proposed framework provides a more solid and understandable ground to detect significant departures from the normal behaviour of the process.

Computational Considerations

The proposed framework relies on efficient sequential inference within a state-space formulation. Because stochastic volatility models are non-linear and non-Gaussian, inference is carried out using sequential Monte Carlo (SMC) methods (Cappe et al.,´ 2005; Doucet et al., 2001). SMC supports online updating and provides accurate approximations to the posterior and predictive distributions required for adaptive control chart construction.

Computational cost is governed primarily by the number of particles and the complexity of the state transition and observation models. In practice, moderate particle counts are sufficient for stable monitoring, and modern implementations benefit from vectorisation and parallel computation. Once predictive samples are available, alarm evaluation is computationally inexpensive, as it reduces to calculating predictive probabilities or quantiles.

For implementation, calibration is best performed offline using representative in-control data to select particle numbers and alarm thresholds. Online monitoring can then proceed under fixed computational budgets, with occasional re-estimation of hyperparameters if process dynamics evolve. These considerations align with established results on the scalability and stability of SMC-based inference (Doucet et al., 2001).

Limitations and Extensions

The findings substantiate the efficacy of stochastic volatility Bayesian SPC; however, several extensions warrant consideration. Heavy-tailed innovations and outliers may necessitate resilient observation models to mitigate the undue influence of isolated shocks. Multivariate extensions are also important in today's manufacturing and healthcare settings, where many quality traits that are related to each other change at the same time. Finally, regime-switching volatility dynamics could be added to tell the difference between persistent volatility states and short-term disturbances, which would make the model easier to understand and work better in environments that are very different from each other.

The evidence from earlier sections shows that modelling time-varying variance as a hidden stochastic process is necessary for reliable monitoring in systems that are not stationary. The suggested framework offers a coherent Bayesian approach to attaining this objective while preserving the operational clarity of control charts and the decision-making significance of predictive probabilities.

Concluding Remarks

This paper proposes a Bayesian statistical process control (SPC) framework that explicitly accommodates time-varying process variance through stochastic volatility modelling. By relaxing the fixed-variance assumption underlying classical SPC, the approach addresses a critical limitation encountered in modern monitoring environments characterised by nonstationary uncertainty. By coupling a Bayesian state-space model with stochastic volatility, it is possible to get a set of adaptive control limits for predictive purposes along with uncertainty quantification that is consistent, and therefore, deciding on monitoring remains stable and interpretable even when the process is changing slowly over time.

It has been found through simulations that conventional Shewhart, EWMA, and CUSUM charts suffer heavy losses of performance when the variance is drifting over time, especially in the case of joint mean and variance shifts. These methods have been found to overstate the number of false alarms, and the run length is unstable for both methods. On the other hand, the suggested Bayesian stochastic volatility control chart performs well even if the volatility is changing a lot and the model is slightly misspecified. This stresses the importance of variance adaptivity to be able to rely on SPC.

The real-data application to the NSE-30 index returns further illustrated the practical value of the framework. Pronounced volatility clustering renders fixed limit control charts ineffective, whereas the Bayesian chart adapts naturally by widening predictive limits during turbulent periods and tightening them during stable regimes, reducing false alarms without obscuring genuine structural changes.

Overall, this work established a principled Bayesian framework that unifies uncertainty modelling, prediction, and monitoring. Future research directions include multivariate stochastic volatility SPC, heavy-tailed observation models, regime-switching volatility dynamics, and continuous-time formulations for high-frequency and sensor-driven applications.

Acknowledgements

The authors acknowledge the support of their respective institutions and thank colleagues for constructive discussions.

REFERENCES

Alkhudaydi, M. H. (2025). An ODE-based dynamic mean-variance portfolio optimisation with time-varying risk aversion. *Computational Economics*. https://doi.org/10.1007/s10614-025-11152-3

Barndorff-Nielsen, O. E., & Shephard, N. (2001). *Non–Gaussian Ornstein–Uhlenbeck–based models and some of their uses in financial economics*. Oxford University Press.

Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis* (2nd ed.). Springer.

Box, G. E. P., Jenkins, G. M., & Reinsel, G. C. (1994). *Time series analysis: Forecasting and control* (3rd ed.). Prentice Hall.

Cappe, O., Moulines, E., & Ryd´ en, T. (2005).´ *Inference in Hidden Markov Models*. Springer.

Carroll, A. R., & Johnson, D. P. (2020). Know it when you see it: Identifying and using special cause variation for quality improvement. *Hospital Pediatrics*, *10*(11), e8–e10. https://doi.org/ 10.1542/hpeds.2020-002303

Chaim, P., & Laurini, M. P. (2024). Bayesian inference for long memory stochastic volatility models. *Econometrics*, *12*(4), 35. https://doi.org/10. 3390/econometrics12040035

Chero, P. P. A. (2019). Statistical process control applied in the chemical and food industry. *Journal of Material Science & Engineering*, *8*, 531.

Colosimo, B. M., & Del Castillo, E. (2010). Bayesian control charts. *Journal of Quality Technology*, *42*(1), 1–16.

Damoun, N., Amekran, Y., Taiek, N., & Hangouche, A. J. E. (2024). Heart rate variability measurement and influencing factors: Towards the standardization of methodology. *Global Cardiology Science & Practice*, *2024*(4), e202435.
https://doi.org/10.21542/gcsp.2024.35

Del Moral, P. (2006). *Feynman–kac formulae: Genealogical and interacting particle systems with applications*. Springer.

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

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, *50*(4), 987–1007.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (3rd ed.). Chapman & Hall/CRC.

Girshick, M. A., & Rubin, H. (1952). Statistical decision problems in industrial quality control. *Annals of Mathematical Statistics*, *23*(1), 114– 125.

Hawkins, D. M. (2001). Change-point detection. *Journal of Quality Technology*, *33*(3), 215–226.

Huang, Y.-F. (2019). An efficient exact Bayesian method for the state space models with stochastic volatility [Forthcoming]. *Studies in Nonlinear Dynamics & Econometrics*. https://doi.org/10. 2139/ssrn.2340389

Iglesias, C., Sancho, J., Pineiro, J. I., Mart˜ ´ınez, J., Pastor, J. J., & Taboada, J. (2016). Shewhart-type control charts and functional data analysis for water quality analysis based on a global indicator. *Desalination and Water Treatment*, *57*(6), 2669–2684. https://doi.org/10.1080/ 19443994.2015.1029533

Iqbal, T., Elahi, A., Wijns, W., Amin, B., & Shahzad, A. (2023). Improved stress classification using automatic feature selection from heart rate and respiratory rate time signals. *Applied Sciences*, *13*(5), 2950. https://doi.org/10.3390/ app13052950

Khan et al., 2. (2023). Monitoring the process mean under the Bayesian approach with application to the hard bake process. *Scientific Reports*, *13*(1), 20723. https://doi.org/10.1038/s41598-02348206-1

Kim, C. K., Yoon, M. H., & Lee, S. (2024). Robust control chart for nonlinear conditionally heteroscedastic time series based on Huber support vector regression. *PLOS ONE*. https://doi. org/10.1371/journal.pone.0299120

Kim, S., Shephard, N., & Chib, S. (1998). Stochastic volatility: Likelihood inference and compar-
in line with arch models. *Review of Economic Studies*, *65*(3), 361–393.

Makis, V. (2008). Bayesian methods for statistical process control. *International Journal of Production Research*, *46*(23), 6619–6637.

Malindzakov´ a, M.,´ Culkov˘ a, K., & Trp´ cevsk˘ a, J.´ (2023). Shewhart control charts implementation for quality and production management. *Processes*, *11*(4), 1246. https://doi.org/10. 3390/pr11041246

Meijer, D., Barumerli, R., & Baumgartner, R. (2025). How relevant is the prior? Bayesian causal inference for dynamic perception in volatile environments. *eLife*, *14*, RP105385. https://doi. org/10.7554/eLife.105385.1

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

Muehlemann, N., Zhou, T., Mukherjee, R., Hossain, M. I., Roychoudhury, S., & Russek-Cohen, E. (2023a). A tutorial on modern Bayesian methods in clinical trials. *Therapeutic Innovation & Regulatory Science*, *57*(3), 402–416. https: //doi.org/10.1007/s43441-023-00515-3

Muehlemann, N., Zhou, T., Mukherjee, R., Hossain, M. I., Roychoudhury, S., & Russek-Cohen, E. (2023b). A tutorial on modern Bayesian methods in clinical trials. *Therapeutic Innovation & Regulatory Science*, *57*(3), 402–416. https: //doi.org/10.1007/s43441-023-00515-3

Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, *41*(1–2), 100–115.

Patharkar, A., Cai, F., Al-Hindawi, F., & Wu, T. (2024). Predictive modeling of biomedical temporal data in healthcare applications: Review and future directions. *Frontiers in Physiology*, *15*, 1386760. https://doi.org/10.3389/fphys.2024. 1386760

Pham, P., Pham, C., Dam, T., et al. (2025). A comprehensive review of catalyst deactivation and regeneration in heavy oil hydroprocessing. *Fuel Processing Technology*, *267*, 108170. https:// doi.org/10.1016/j.fuproc.2024.108170

Reynolds, M. R., & Stoumbos, Z. J. (2005). Monitoring variance in the presence of autocorrelation. *Journal of Quality Technology*, *37*(2), 112– 126.

Reza, M. S., Iskakova, Z. B., Afroze, S., Kuterbekov, K., Kabyshev, A., Bekmyrza, K. Z., Kubenova, M. M., Bakar, M. S. A., Azad, A. K., Roy, H., & Islam, M. S. (2023). Influence of catalyst on the yield and quality of bio-oil for the catalytic pyrolysis of biomass: A comprehensive review. *Energies*, *16*(14), 5547. https://doi.org/ 10.3390/en16145547

Roberts, S. W. (1959). Control chart tests based on geometric moving averages. *Technometrics*, *1*(3), 239–250.

Salinas-Camus, M., Goebel, K., & Eleftheroglou, N. (2025). A comprehensive review and evaluation framework for data-driven prognostics: Uncertainty, robustness, interpretability, and feasibility. *Mechanical Systems and Signal Processing*, *237*, 113015. https://doi.org/10.1016/j.ymssp.2025.113015

Schach, S., Eilert, T., Presser, B., & Kunzelmann, M. (2025). Bayesian hierarchical modeling for variance estimation in biopharmaceutical processes. *Bioengineering*, *12*(2), 193. https://doi.org/10.3390/bioengineering12020193

Shayakhmetova, A., Abdildayeva, A., Akhmetova, A., Shurenov, M., & Nagibayeva, A. (2025). Applications of the Bayesian method for predicting equipment failures. *International Journal of Innovative Research and Scientific Studies*, *8*(5), 884–895.

Shewhart, W. A. (1931). *Economic control of the quality of the manufactured product*. D. Van Nostrand Company.

Tartakovsky, A. G., Nikiforov, I. V., & Basseville, M. (2014). *Sequential analysis: Hypothesis testing and changepoint detection*. Chapman & Hall/CRC.

Taylor, S. J. (1986). *Modelling financial time series*. John Wiley & Sons.

West, M. (1986). Bayesian monitoring of time series. *Journal of the American Statistical Association*, *81*(396), 1025–1035.

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.

Zhou, Q., Chai, B., Li, Y., Tang, C., Guo, Y., & Ye, Y. (2026). A mixture-of-experts prior-posterior fusion framework for predicting the remaining useful life of aerospace high-speed bearings. *Neurocomputing*, *670*, 132601. https://doi.org/10.1016/j.neucom.2025.132601