Monitoring the Process Mean with the EWMA Control Chart: A Brief Simulation Study

Student Name
MSc in Statistics and Actuarial-Financial Mathematics

Spring 2026

1 Introduction

A fundamental objective of statistical process control is the rapid detection of persistent shifts in the mean of a quality characteristic. The classical Shewhart $\bar{X}$ chart (Shewhart, 1931) bases each decision on the most recent subgroup alone and is therefore highly effective for large shifts but relatively slow to react when the displacement is small. Two families of charts have been developed to address this limitation: the cumulative sum (CUSUM) chart, introduced by Page (1954), and the exponentially weighted moving average (EWMA) chart, proposed by Roberts (1959).

In this report we focus on the EWMA chart applied to a univariate normal process. We present the method, implement it in a simulation study that mimics Phase I parameter estimation followed by Phase II online monitoring, and compare its detection speed for two choices of the smoothing parameter λ. Throughout, the process variable is assumed to be normally distributed with known (or Phase I–estimated) parameters; shift detection is evaluated on a single realisation for illustration, although in practice the average run length (ARL) would be studied via repeated simulation (Lucas and Saccucci, 1990).

2 Methodology

Let $X_1, X_2, \ldots$ be sequentially observed values from a process with in-control (IC) mean μ_0 and standard deviation σ. The EWMA statistic is defined recursively as

$$Z_i = \lambda X_i + (1 - \lambda) Z_{i-1}, \qquad Z_0 = \mu_0, \tag{1}$$

where $\lambda \in (0, 1]$ is the smoothing parameter. Expanding the recursion shows that Z_i is a weighted average of the entire observation history, with weights that decay geometrically at rate $(1 - \lambda)$ per step.

Because the X_i are independent with common variance σ^2, the variance of Z_i is

$$\operatorname{Var}(Z_i) = \frac{\lambda}{2 - \lambda} \left[1 - (1 - \lambda)^{2i}\right] \sigma^2,$$

which converges to the steady-state value $\sigma_Z^2 = \lambda \sigma^2/(2 - \lambda)$ as $i \to \infty$. Using the steady-state variance, the control limits are

$$\text{UCL} = \mu_0 + L \sigma \sqrt{\frac{\lambda}{2 - \lambda}}, \qquad \text{LCL} = \mu_0 - L \sigma \sqrt{\frac{\lambda}{2 - \lambda}}, \tag{2}$$

where L is typically set to 3.

The parameter λ controls the trade-off between sensitivity and smoothness. A small λ (e.g. 0.05–0.10) assigns most of the weight to the past, producing a smooth statistic that is very sensitive to small, persistent shifts. A large λ (approaching 1) concentrates the weight on the most recent observation and the EWMA reduces to the standard Shewhart chart. Values in the range $0.05 \le \lambda \le 0.25$ are recommended in practice (Lucas and Saccucci, 1990).

In practice μ_0 and σ are unknown and must be estimated from a Phase I reference sample collected under IC conditions. In our simulation we use the Phase I sample mean $\hat{\mu}$ and sample standard deviation $\hat{\sigma}$ in place of μ_0 and σ in (1)–(2).

3 Simulation Study

We generate $n_{\mathrm{I}} = 50$ Phase I observations from $N(0, 1)$ and compute $\hat{\mu}$ and $\hat{\sigma}$. In Phase II we generate $n_{\mathrm{II}} = 50$ observations: the first 29 from the same IC distribution $N(\hat{\mu}, \hat{\sigma}^2)$, and observations 30–50 from $N(\hat{\mu} + 1 \cdot \hat{\sigma}, \hat{\sigma}^2)$, simulating a persistent 1σ shift in the mean at $\tau = 30$. We apply the EWMA chart with $L = 3$ for two choices of the smoothing parameter, $\lambda = 0.05$ and $\lambda = 0.25$.

Figure 1 shows the resulting charts. The vertical dotted line marks the true change point. With $\lambda = 0.05$ the EWMA statistic is smoother and signals the shift shortly after the change point. With $\lambda = 0.25$ the statistic is noisier and less smooth, and the first out-of-control signal may arrive slightly later because the wider steady-state limits partially offset the faster response. Both charts detect the 1σ shift well within Phase II, confirming the EWMA's superiority over the standard Shewhart chart for shifts of this magnitude.

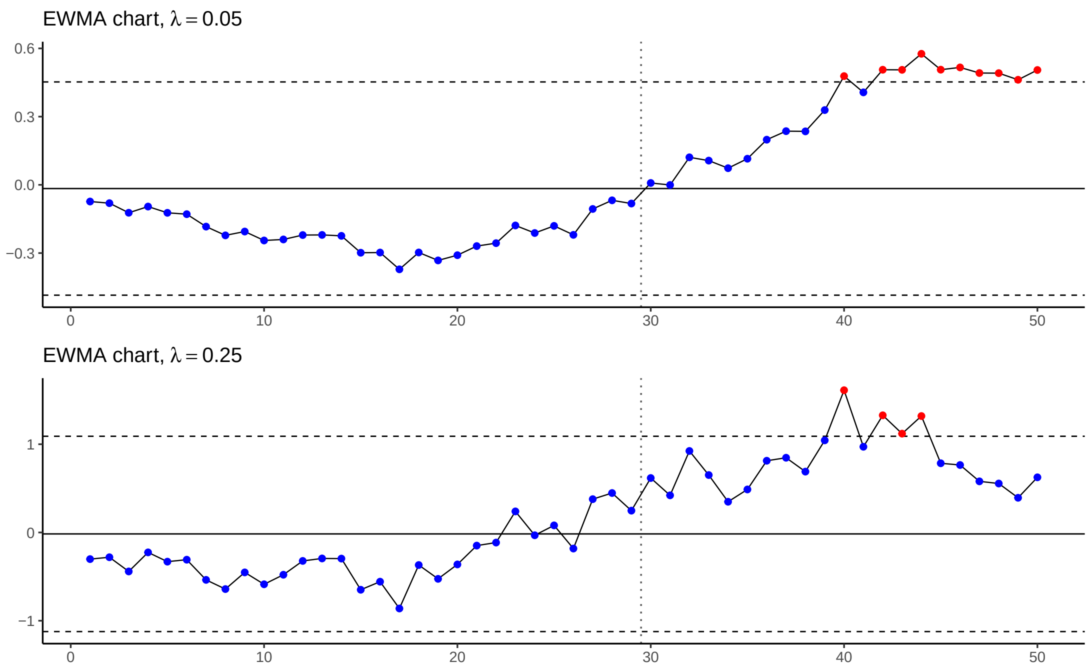

Figure 1: EWMA charts for a simulated 1σ shift at $\tau = 30$. Top panel: $\lambda = 0.05$. Bottom panel: $\lambda = 0.25$. Red points denote out-of-control signals.

4 Conclusion

The simulation confirms that the EWMA chart successfully detects a persistent 1σ shift in the process mean for both values of λ considered. The smaller smoothing parameter ($\lambda = 0.10$) produces a smoother statistic and narrower control limits, which makes it particularly effective for detecting small to moderate shifts. The larger value ($\lambda = 0.30$) offers a compromise between Shewhart-like responsiveness to large shifts and EWMA-type sensitivity to smaller ones. These results are consistent with the findings of Roberts (1959) and the more extensive ARL analysis of Lucas and Saccucci (1990), and illustrate why the EWMA chart has become a standard tool for Phase II process monitoring.

References

Lucas, J. M. and Saccucci, M. S. (1990). Exponentially weighted moving average control schemes: properties and enhancements. *Technometrics*, 32(1):1–12.

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

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

Shewhart, W. A. (1931). *Economic Control of Quality of Manufactured Product*. Van Nostrand, New York.