Statistical Process Control

MSc: Statistics and Actuarial-Financial Mathematics

Konstantinos Bourazas
kbourazas@aegean.gr

Department of Statistics and Actuarial-Financial Mathematics

Course #11

May 7, 2026

Self-starting control charts for unknown IC parameters

- The CUSUM and EWMA charts we studied assume that the IC parameters are **known**. In practice they must be estimated, and the randomness in the estimates can severely distort the chart's performance
- A **self-starting** chart estimates the IC parameters **recursively**: calibration and testing happen **simultaneously**, with the estimates improving as more data arrive
- This makes self-starting charts ideal for **short production runs**, where there is not enough data for a separate Phase I exercise, and for processes that require **online monitoring from the very first observation**
- Today we cover three perspectives on self-starting monitoring
    - Frequentist parametric
    - Bayesian parametric
    - Nonparametric

The Q-chart[3]

- For individual observations $X_1, X_2, \ldots$ from $N(\mu_0, \sigma^2)$ with both parameters unknown, let $\bar{X}_{n-1}$ and s_{n-1} denote the mean and standard deviation of the first $n-1$ observations, and define

$$T_n = \frac{X_n - \bar{X}_{n-1}}{s_{n-1}}, \qquad n \geq 3$$

- Quesenberry showed that $\sqrt{(n-1)/n}\; T_n \sim t_{n-2}$, and the probability-integral transformation

$$Q_n = \Phi^{-1}\left[\Upsilon_{n-2}\left(\sqrt{\frac{n-1}{n}}\; T_n\right)\right]$$

  produces **i.i.d.** $N(0,1)$ values under IC, regardless of μ_0 and σ

- A Shewhart chart on Q_n signals when $|Q_n| > L$, with $L \approx 3.023$ for $\text{ARL}_0 = 370$
- The same transformation exists for Poisson[1] and Binomial[2] data

[1] Quesenberry, C. P. (1995). On properties of Poisson Q charts. *J. Qual. Technol.*, 27, 293–303.

[2] Quesenberry, C. P. (1995). On properties of binomial Q charts. *J. Qual. Technol.*, 27, 204–213.

[3] Quesenberry, C. P. (1991). SPC Q charts for start-up processes and short or long runs. *J. Qual. Technol.*, 23, 213–224.

Q-chart example: simulated Normal data

- We observe $n = 20$ individual observations with both μ_0 and σ unknown
- All observations are generated from the IC process $N(0, 1)$, except for a single large outlier
- A single **outlier** of $+4.5\sigma$ is planted at observation $n = 11$
- At each step $n \geq 3$, the running estimates $\bar{X}_{n-1}$ and s_{n-1} are updated and used to compute the Q-statistic
- The Q-statistic is available from $n = 3$ onward, so the first two positions in the Q-chart are empty
- The Shewhart chart signals when the Q-statistic exceeds the control limits, here using $L = 3.023$
- In this example, the outlier at $n = 11$ produces a signal in the Q-chart

Q-chart in action (Normal, both parameters unknown)

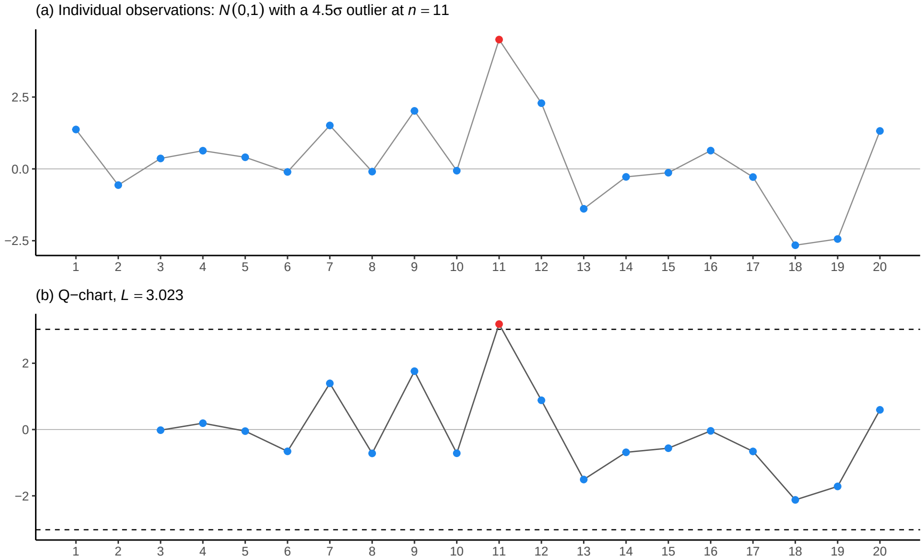

Self-starting CUSUM (SSC)[4]

- Apply the standard one-sided CUSUM recursion to the Q-statistics

$$C_n^+ = \max\big(0,\ C_{n-1}^+ + Q_n - k\big), \qquad C_2^+ = 0$$

  and signal when $C_n^+ > h$
- Since the Q_n are i.i.d. $N(0,1)$ under IC, the **standard CUSUM design** can be used directly
- The reference value k controls the size of shift the chart is tuned to detect, while h controls the false alarm rate
- A key limitation is **shift absorption**: if the process shifts but no alarm is raised quickly, the OOC observations are gradually incorporated into the running estimates $\bar{X}_{n-1}$ and s_{n-1}
- As a result, the apparent signal can weaken over time, even though the process remains out of control
- SSC versions for Normal, Poisson, and Binomial data are all derived via the same Q-transformation principle

[4] Hawkins, D. M. and Olwell, D. H. (1998). *Cumulative Sum Charts and Charting for Quality Improvement*. Springer, New York.

SSC example: simulated Normal data

- We observe $n = 30$ individual observations with both μ_0 and σ unknown
- The first 10 observations are generated from the IC process $N(0, 1)$
- From observation $n = 11$ onward, the process mean shifts by 1σ, so the remaining observations are generated from $N(1, 1)$
- At each step $n \geq 3$, the running estimates $\bar{X}_{n-1}$ and s_{n-1} are updated and used to compute the Q-statistic
- The self-starting CUSUM is then applied to the sequence of Q-statistics
- For illustration, we use $k = 0.25$ and $h = 4.5$
- The conventional CUSUM with known $\mu_0 = 0$ and $\sigma = 1$ is also shown as a benchmark

SSC in action

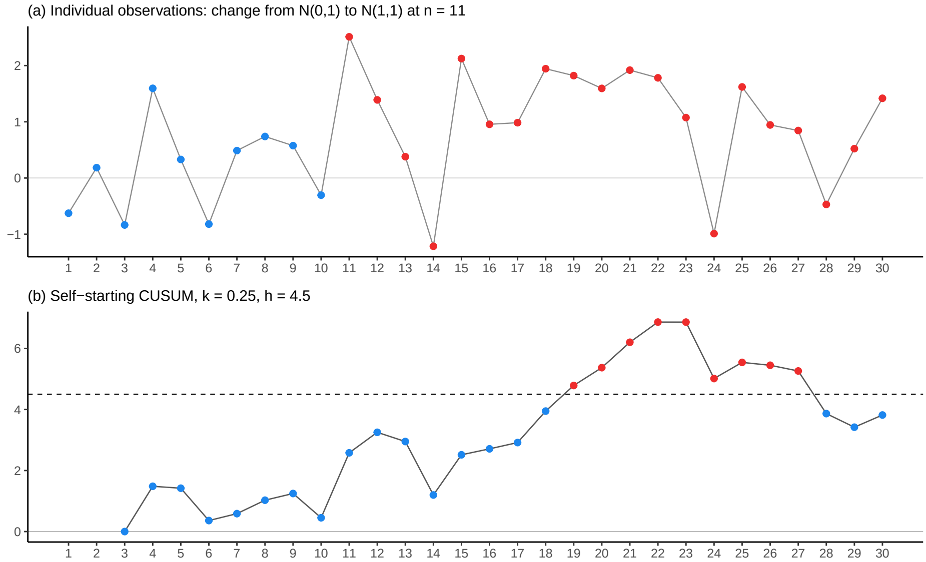

Bayesian self-starting charts: PCC and PRC

- Two Bayesian self-starting charts have been developed for any distribution in the **regular exponential family** (Normal, Poisson, Binomial, Gamma, ...): the **Predictive Control Chart (PCC)**[5] and the **Predictive Ratio CUSUM (PRC)**[6]

- We illustrate with $X_i \mid \theta \sim N(\theta, \sigma^2)$ (known σ^2, unknown mean) with conjugate prior $\theta \sim N(\mu_0, \sigma_0^2)$. The posterior after $n-1$ observations is $\theta \mid X_{n-1} \sim N(\hat{\mu}_{n-1}, \hat{\sigma}^2_{n-1})$, where

$$\hat{\sigma}^2_{n-1} = \left(\frac{n-1}{\sigma^2} + \frac{1}{\sigma_0^2}\right)^{-1}, \qquad \hat{\mu}_{n-1} = \hat{\sigma}^2_{n-1}\left(\frac{(n-1)\bar{X}_{n-1}}{\sigma^2} + \frac{\mu_0}{\sigma_0^2}\right)$$

- The **predictive distribution** of X_n given X_{n-1} integrates out θ

$$X_n \mid X_{n-1} \sim N\big(\hat{\mu}_{n-1},\ \hat{\sigma}^2_{n-1} + \sigma^2\big)$$

- The predictive variance combines **parameter uncertainty** ($\hat{\sigma}^2_{n-1}$, shrinks with data) and **observation noise** (σ^2, constant). With a flat prior $p_0(\theta) \propto 1$ it coincides with the Q-chart

[5] Bourazas, K., Kiagias, D., and Tsiamyrtzis, P. (2022). Predictive Control Charts (PCC). *J. Qual. Technol.*, 54, 367–391.
[6] Bourazas, K., Sobas, F., and Tsiamyrtzis, P. (2023). Predictive ratio CUSUM (PRC). *J. Qual. Technol.*, 55, 391–403.

Predictive Control Chart (PCC)

- PCC is a **Shewhart-type** Bayesian chart that signals when a new observation falls outside the $100(1-\alpha)\%$ **Highest Predictive Density** region

$$R_n = \hat{\mu}_{n-1} \ \pm \ z_{1-\alpha/2} \sqrt{\hat{\sigma}^2_{n-1} + \sigma^2}$$

- The probability level α can be selected in different ways, depending on the monitoring objective
- For a fixed monitoring horizon, α may be chosen to control the overall false alarm probability
- Alternatively, α can be calibrated by simulation to attain a target ARL_0, as in standard control-chart design
- The predictive limits are **not constant**: they update after each observation and typically narrow as posterior uncertainty about the IC parameter decreases
- PCC is designed for detecting **large transient shifts** or outliers, just like a Shewhart chart
- With a moderately informative prior, PCC can improve early detection, especially when only a few observations are available

PCC example: simulated Normal data

- We use the same 20 individual observations as in the Q-chart example
- All observations are generated from the IC process $N(0, 1)$, except for a single large outlier
- A single **outlier** of $+4.5\sigma$ is planted at observation $n = 11$
- The variance is treated as known, while the IC mean is estimated sequentially using the Bayesian posterior distribution
- Before each new observation, PCC computes the current predictive distribution and its corresponding predictive control limits
- The limits are updated at every step and are therefore **time-varying**, not horizontal
- In this example, the outlier at $n = 11$ falls outside the predictive limits and produces a PCC signal

PCC in action (Normal, known variance)

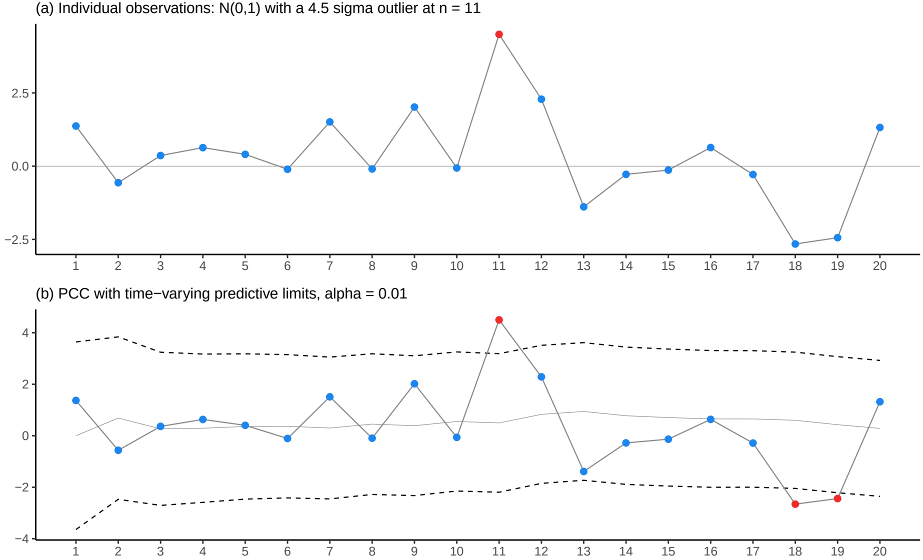

Predictive Ratio CUSUM (PRC)

- PRC is a **CUSUM-type** Bayesian chart for detecting **small persistent shifts**
- It is based on the ratio of two posterior predictive distributions: the usual IC predictive distribution and an OOC predictive distribution obtained by applying a specified intervention to the posterior parameters
- In the Normal mean example, this intervention corresponds to a target mean shift of $k \cdot \sigma$
- Let $z_n = (X_n - \hat{\mu}_{n-1})/\sqrt{\hat{\sigma}^2_{n-1} + \sigma^2}$ be the standardised predictive residual and $\kappa_n = k\sigma/\sqrt{\hat{\sigma}^2_{n-1} + \sigma^2}$. The log predictive ratio is

$$\log L_n = \kappa_n\left(z_n - \frac{\kappa_n}{2}\right)$$

  and the CUSUM recursion is $S_n = \max(0,\ S_{n-1} + \log L_n)$, $S_0 = 0$
- The form mirrors the classical CUSUM log-ratio, but the predictive variance changes over time as the posterior distribution is updated
- The threshold h is calibrated by simulation to achieve the desired false alarm performance

PRC example: simulated Normal data

- We use the same data structure as in the self-starting CUSUM example, with $n = 30$ individual observations
- The first 10 observations are generated from the IC process $N(0, 1)$
- From observation $n = 11$ onward, the process mean shifts by 1σ, so the remaining observations are generated from $N(1, 1)$
- The variance is treated as known, while the IC mean is assigned the prior $\theta \sim N(0, 1)$
- At each step, the posterior predictive distribution is updated using the observations already seen
- PRC accumulates the log predictive ratios between the shifted and IC predictive distributions
- For illustration, we use $k = 0.5$ and $h = 4.5$

PRC in action (Normal, known variance)

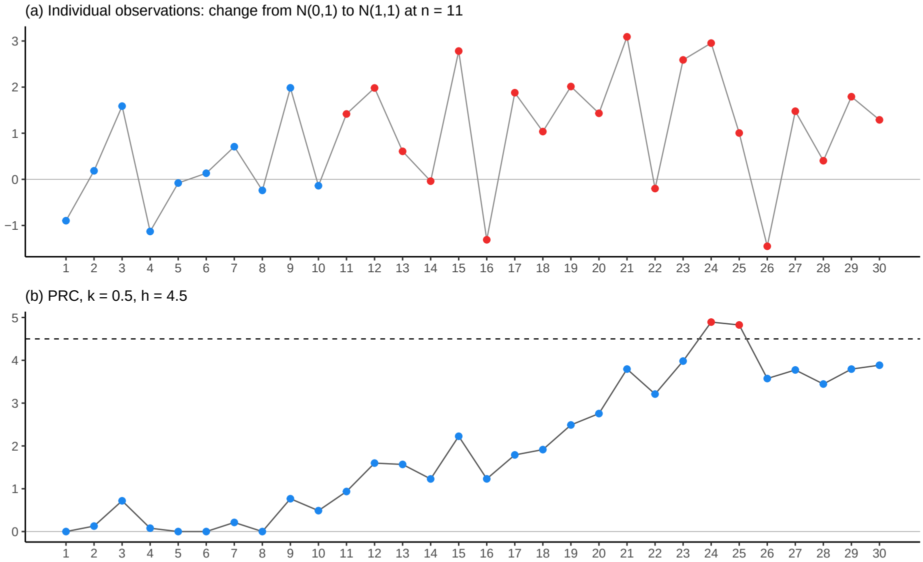

Shift absorption and practical recommendations

- All self-starting charts (frequentist and Bayesian) face the **shift absorption** problem: undetected OOC observations are incorporated into the running parameter estimates, which gradually absorb the shift and reduce the charting statistic back toward its IC level
- Absorption is most severe **early in the process**, when few IC observations are available and the estimates are most sensitive to incoming data
- This creates a **window of opportunity** for detection. React to the first signal quickly
- When an alarm is raised: stop the process, investigate, and if corrective action is taken, **restart the chart** using the previous IC data
- In the Bayesian framework, an **informative prior** anchors the parameter estimates and can substantially **mitigate absorption**: the posterior is less easily pulled by OOC observations, keeping the charting statistic elevated for longer
- PCC and PRC are **complementary**: run both simultaneously to cover large transient shifts (PCC) and small persistent shifts (PRC)

Nonparametric self-starting charts

- The Q-chart, PCC, and PRC all require a **known distributional family**. With individual observations and no separate Phase I, there is no opportunity to estimate or validate the distributional form, so a misspecification directly compromises IC performance
- Nonparametric alternatives use **ranks** or **empirical distributions** instead, so their IC properties hold for **any continuous distribution**
- We distinguish charts by their target
  - **Location shifts**: SNS CUSUM
  - **Scale shifts**: Squared Ranks chart
  - **Arbitrary distributional changes**: NLR EWMA and CPM
- All require an IC **reference sample** (warm-up) of moderate size ($t_0 \geq 50$)

Sequential Normal Score (SNS) CUSUM for location shifts[7]

- Let $X_1, \ldots, X_{t_0}$ be the IC warm-up sample of size t_0 and $X_{t_0+1}, X_{t_0+2}, \ldots$ the Phase II observations. At step $j \geq 1$, let R_j denote the rank of X_{t_0+j} among all $t_0 + j$ observations seen so far
- Convert the rank to a **sequential normal score**

$$Z_j = \Phi^{-1}\left(\frac{R_j}{t_0 + j + 1}\right)$$

  Under IC, the Z_j are approximately i.i.d. $N(0,1)$ for any continuous IC distribution
- Apply the standard one-sided CUSUM recursion directly to the Z_j

$$C_j^+ = \max(0,\ C_{j-1}^+ + Z_j - k), \qquad C_0^+ = 0$$

  where $k > 0$ is the reference value (typically $k = 0.5$ for a 1σ shift) and the chart signals when $C_j^+ > h$
- The chart is **distribution-free**: the ranks do not depend on the shape of the IC distribution
- Effective for **location shifts** across symmetric and skewed distributions alike

[7] Conover, W. J., Guerrero-Serrano, A. J., and Tercero-Gómez, V. G. (2017). Sequential normal scores in nonparametric SPC. *J. Qual. Technol.*, 49, 1–15.

Squared Ranks (SR) change-point chart for scale shifts[8]

- The SR chart is a **nonparametric change-point chart** for detecting sustained changes in process variance
- At each time N, centre all observations around the overall mean, compute $U_i = |X_i - \bar{X}|$, and rank the U_i's
- For each possible change point k, compute the squared-ranks statistic and standardise it:

$$T_k^* = \sum_{i=1}^{k} [R(U_i)]^2, \qquad T_k = \frac{T_k^* - E_0(T_k^*)}{\sqrt{\mathrm{Var}_0(T_k^*)}}, \qquad 2 \le k \le N-2$$

- The monitored statistic is

$$C_N = \max_{2 \le k \le N-2} |T_k|$$

  so the chart searches for the strongest evidence of a variance change over all possible split points
- The chart signals when $C_N > h_{N,\alpha}$, where the control limit depends on the current sample size N and is calibrated by simulation

[8] Villanueva-Guerra, E. C. et al. (2017). A control chart for monitoring scale based on squared ranks. *J. Stat. Comp. Simul.*, 87, 3243–3268.

Charts for arbitrary distributional changes

- When the type of change (location, scale, or shape) is **not known in advance**, omnibus methods are needed
- The **Nonparametric Likelihood Ratio EWMA (NLR EWMA)**[9] replaces parametric densities with exponentially weighted empirical cdfs. It detects any shift in the cdf of the observations
- The **change-point model (CPM)**[10] tests at each time whether a distributional change has occurred, using statistics like Cramér-von Mises (for location) or Kolmogorov-Smirnov (for general changes)
- Both are pay a price in **detection delay**: they are less powerful than targeted charts when the shift type is known

[9] Zou, C. and Tsung, F. (2010). Likelihood ratio-based distribution-free EWMA control charts. *J. Qual. Technol.*, 42, 174–196.

[10] Ross, G. J. (2012). Parametric and nonparametric sequential change detection in R: the `cpm` package. *J. Stat. Softw.*, 66(3), 1–20.