Statistical Process Control

MSc: Statistics and Actuarial-Financial Mathematics

Konstantinos Bourazas
kbourazas@aegean.gr

Department of Statistics and Actuarial-Financial Mathematics

Course 3

March 5, 2026

$\bar{X}$ and s charts

- In the $\bar{X}$ and R charts we estimate σ from the average range $\bar{R}$
- The main reason was **computational convenience**: the range is just a subtraction, whereas the sample standard deviation requires squaring and a square root
- With modern software this advantage is negligible
- The sample standard deviation s_i is a **more efficient estimator** of σ than the range, especially when the subgroup size m is large
- The $\bar{X}$ **and** s **charts** replace the range with s_i as the basis for estimating σ and monitoring process variability
- They are **preferred over** the $\bar{X}$ and R charts, especially when $m > 6$

The sample standard deviation and its bias

- For each subgroup i we compute

$$s_i = \sqrt{\frac{1}{m-1}\sum_{j=1}^{m}(X_{ij} - \bar{X}_i)^2}, \qquad i = 1, 2, \ldots, n$$

- The sample variance s_i^2 is an **unbiased estimator** of σ^2, but s_i itself is a **biased estimator** of σ
  - Since $g(x) = \sqrt{x}$ is strictly concave, **Jensen's inequality** gives

$$E[s_i] = E\left[\sqrt{s_i^2}\right] < \sqrt{E[s_i^2]} = \sqrt{\sigma^2} = \sigma$$

- To correct the bias we define the constant

$$d_3(m) = E\left[\frac{s_i}{\sigma}\right] = \frac{\Gamma(m/2)}{\Gamma\big((m-1)/2\big)}\sqrt{\frac{2}{m-1}}$$

- An **unbiased estimator** of σ based on all n subgroups is

$$\hat{\sigma} = \frac{\bar{s}}{d_3(m)}, \qquad \bar{s} = \frac{1}{n}\sum_{i=1}^{n} s_i$$

Table of constants $d_3(m)$ (bias correction for s_i)

m	$d_3(m)$	m	$d_3(m)$
2	0.7979	14	0.9810
3	0.8862	15	0.9823
4	0.9213	16	0.9835
5	0.9400	17	0.9845
6	0.9515	18	0.9854
7	0.9594	19	0.9862
8	0.9650	20	0.9869
9	0.9693	21	0.9876
10	0.9727	22	0.9882
11	0.9754	23	0.9887
12	0.9776	24	0.9892
13	0.9794	25	0.9896

- As m grows, $d_3(m) \to 1$ and the bias vanishes
- For small m the correction matters as ignoring it would systematically underestimate σ and produce **control limits that are too narrow**

Control limits for the $\bar{X}$ chart using sample standard deviations

- We estimate the standard error of the subgroup mean as

$$\hat{\sigma}_{\bar{X}} = \frac{\bar{s}}{d_3(m)\sqrt{m}}$$

- The control limits follow the same $\bar{\bar{X}} \pm Z_{1-\alpha/2}\,\hat{\sigma}_{\bar{X}}$ structure as before

$$\text{UCL} = \bar{\bar{X}} + \frac{Z_{1-\alpha/2}}{d_3(m)\sqrt{m}}\,\bar{s}, \qquad \text{CL} = \bar{\bar{X}}, \qquad \text{LCL} = \bar{\bar{X}} - \frac{Z_{1-\alpha/2}}{d_3(m)\sqrt{m}}\,\bar{s}$$

- The only difference from the $\bar{X}$ and R chart is that $\bar{R}/d_1(m)$ has been replaced by $\bar{s}/d_3(m)$ as the estimator of σ
- As before, $Z_{1-\alpha/2} = 3$ is the standard choice

Control limits for the s chart

- We monitor s_i around its IC baseline $\bar{s}$
- The standard deviation of s_i under IC follows from $\mathrm{Var}(s_i) = E[s_i^2] - (E[s_i])^2 = \sigma^2 - d_3^2(m)\,\sigma^2$, which gives $\sigma_{s_i} = \sqrt{1 - d_3^2(m)}\;\sigma$. We estimate it by

$$\hat{\sigma}_{s_i} = \frac{\sqrt{1 - d_3^2(m)}}{d_3(m)}\;\bar{s}$$

- The s chart limits are therefore

$$\mathrm{UCL} = \left(1 + \frac{Z_{1-\alpha/2}\sqrt{1 - d_3^2(m)}}{d_3(m)}\right)\bar{s}, \quad \mathrm{CL} = \bar{s}, \quad \mathrm{LCL} = \left(1 - \frac{Z_{1-\alpha/2}\sqrt{1 - d_3^2(m)}}{d_3(m)}\right)\bar{s}$$

- As with the R chart, if the LCL turns out negative we set it to 0
- A point above the UCL signals **increased variability**, while a point below a positive LCL signals **decreased variability**

An alternative: the s^2 chart

- Instead of monitoring s_i, we can monitor the sample variance s_i^2 directly
- Under IC normality, $\dfrac{(m-1)\,s_i^2}{\sigma^2} \sim \chi^2_{m-1}$
- This leads to **exact chi-squared limits** rather than limits based on a normal approximation

$$\text{UCL} = \frac{\bar{s^2}\,\chi^2_{1-\alpha/2,\,m-1}}{m-1}, \qquad \text{CL} = \bar{s^2}, \qquad \text{LCL} = \frac{\bar{s^2}\,\chi^2_{\alpha/2,\,m-1}}{m-1}$$

  where $\bar{s^2} = \frac{1}{n}\sum_{i=1}^{n} s_i^2$
- Both the s chart and the s^2 chart require **normality** of the underlying observations. The difference is that the s chart approximates the distribution of s_i as normal when setting limits, whereas the s^2 chart uses the **exact sampling distribution** of s_i^2

Simulated example for the s^2 chart

- We generate $n = 25$ Phase I subgroups of size $m = 5$ from $N(50,\ 1^2)$
- We then generate 5 Phase II subgroups
  - The first three from the same IC distribution $N(50,\ 1^2)$
  - The last two from $N(50,\ 3^2)$, simulating a **doubling of** σ
- We build the s^2 chart using the Phase I data and the exact χ^2 quantiles with $\alpha = 0.0027$

$$\text{UCL} = \frac{\bar{s^2}\ \chi^2_{0.99865,\ 4}}{4}, \qquad \text{CL} = \bar{s^2}, \qquad \text{LCL} = \frac{\bar{s^2}\ \chi^2_{0.00135,\ 4}}{4}$$

- This example illustrates how the s^2 chart detects a **pure variability shift** that leaves the mean unchanged

Simulated example s^2 chart

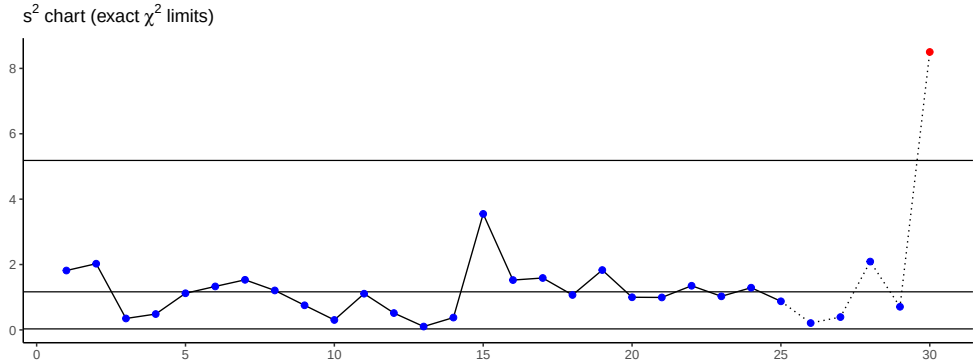

Phase I workflow for $\bar{X}$ and s charts

- The workflow is identical to the one for $\bar{X}$ and R charts
- Collect n rational subgroups of size m under conditions believed to be **IC**
- Compute $\bar{X}_i$ and s_i for each subgroup, then $\bar{\bar{X}}$ and $\bar{s}$
- Look up or compute $d_3(m)$ for the chosen subgroup size
- **First** check the s chart for stability of variability
  - If the s chart shows OOC points, the $\bar{X}$ chart limits based on $\bar{s}$ are not trustworthy
  - Only proceed to the $\bar{X}$ chart after variability is confirmed stable
- Investigate and remove any special causes, then rebuild the baseline

Example 3.2 Injection molding process revisited

- We return to the same injection molding data from Example 3.1, but now build $\bar{X}$ and s charts instead of $\bar{X}$ and R charts
- From the $n = 20$ Phase I subgroups with $m = 5$ we compute

$$\bar{\bar{X}} = 79.533, \qquad \bar{s} = 3.575, \qquad d_3(5) = 0.940$$

- With $Z_{1-\alpha/2} = 3$ the $\bar{X}$ chart limits become

$$\text{UCL} = 79.533 + \frac{3}{0.940 \times \sqrt{5}} \times 3.575 = 84.636, \qquad \text{LCL} = 74.430$$

- The s chart limits become

$$\text{UCL} = \left(1 + \frac{3\sqrt{1 - 0.940^2}}{0.940}\right) \times 3.575 = 7.468, \qquad \text{LCL} = 0$$

since the raw LCL of -0.318 is negative and we set it to zero

Example 3.2 $\bar{X}$ and s charts

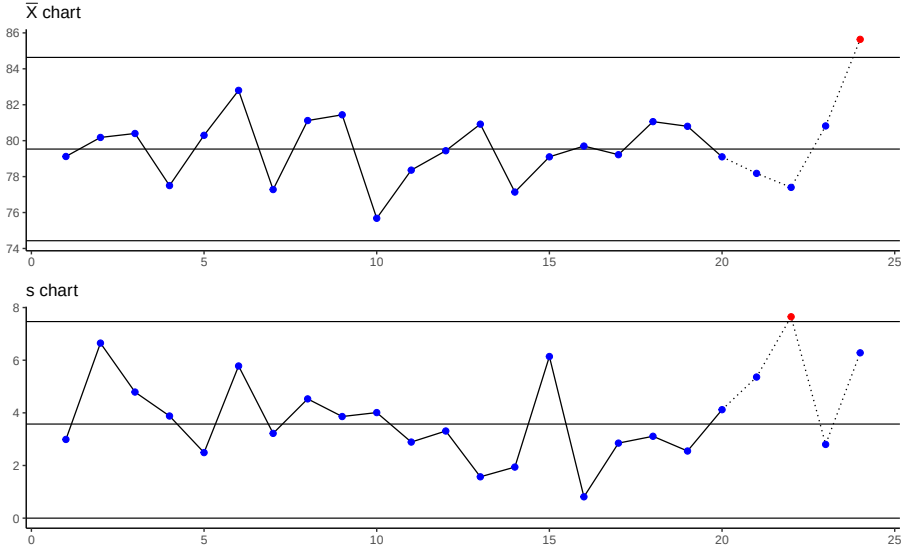

Example 3.2 Interpretation

- During Phase I (subgroups 1 to 20), both charts show all points within limits, confirming that the process is **IC**
- In Phase II the four new subgroups are plotted against the same limits
- The s chart gives a signal at the **22nd time point**, where $s_{22} = 7.65$ exceeds the UCL of 7.468
- This variability alarm arrives **two subgroups earlier** than the mean alarm at time 24 that the $\bar{X}$ and R chart gave in Example 3.1
- This illustrates the advantage of the s chart: it can be **more sensitive** to variability shifts than the R chart, particularly when m is not very small

When to prefer the $\bar{X}$ and s charts

- For $m \leq 6$ the range and the sample SD carry nearly the same information about σ, so either pair of charts is fine
- For $m > 6$ the relative efficiency of the range drops noticeably and the $\bar{X}$ and s charts should be used instead
- The s chart also generalises more naturally
  - It connects directly to the chi-squared distribution through the s^2 chart
  - It extends to the multivariate setting where the sample covariance matrix replaces s_i
- In modern software, both chart types are computed with equal ease, so the choice is driven purely by **statistical efficiency**

X and MR charts for individual observations

- In some applications we collect **only one observation per time point**
  - Chemical processes where each batch yields a single measurement
  - Expensive or destructive tests
  - Continuous production where readings are taken at regular intervals
- With subgroup size $m = 1$, the within-subgroup range and SD are both **zero**, so the standard $\bar{X}$ and R charts cannot be used
- We need a different strategy for estimating σ
- The standard approach is based on **moving ranges**, which create artificial subgroups from consecutive observations

Creating subgroups with moving windows

- Let $X_1, X_2, \ldots, X_n$ be individual observations
- Choose a window size $\tilde{m} > 1$ and form overlapping groups

$$\text{Group 1: } X_1, X_2, \ldots, X_{\tilde{m}}, \quad \text{Group 2: } X_2, X_3, \ldots, X_{\tilde{m}+1}, \quad \ldots$$

- The range of each group is the **moving range**

$$MR_i = \max(X_i, \ldots, X_{i+\tilde{m}-1}) - \min(X_i, \ldots, X_{i+\tilde{m}-1}), \quad i = 1, \ldots, n - \tilde{m} + 1$$

- The most common choice is $\tilde{m} = 2$, which gives

$$MR_i = |X_{i+1} - X_i|, \quad i = 1, \ldots, n-1$$

- From the average moving range $\overline{MR}$ we estimate $\hat{\sigma} = \overline{MR} \,/\, d_1(\tilde{m})$, exactly as in the grouped-data case

Control limits for the X chart (individual observations)

- The centre line is the overall mean $\bar{X}$, and the limits use $\hat{\sigma} = \overline{MR} \,/\, d_1(\tilde{m})$

$$\text{UCL} = \bar{X} + \frac{Z_{1-\alpha/2}}{d_1(\tilde{m})}\,\overline{MR}, \qquad \text{CL} = \bar{X}, \qquad \text{LCL} = \bar{X} - \frac{Z_{1-\alpha/2}}{d_1(\tilde{m})}\,\overline{MR}$$

- Note that **there is no** $\sqrt{m}$ in the denominator because each "subgroup" that is plotted consists of a single observation
- For $\tilde{m} = 2$ and $Z_{1-\alpha/2} = 3$ the limits simplify to

$$\bar{X} \pm \frac{3}{1.128}\,\overline{MR} = \bar{X} \pm 2.66\,\overline{MR}$$

- We declare the process mean OOC at time i when X_i falls outside the limits

Control limits for the MR chart (individual observations)

- The moving ranges MR_i are monitored around their average $\overline{MR}$
- The limits follow the same structure as the R chart, with $\tilde{m}$ in place of m

$$\text{UCL} = \left(1 + \frac{Z_{1-\alpha/2}\, d_2(\tilde{m})}{d_1(\tilde{m})}\right)\overline{MR}, \quad \text{CL} = \overline{MR}, \quad \text{LCL} = \left(1 - \frac{Z_{1-\alpha/2}\, d_2(\tilde{m})}{d_1(\tilde{m})}\right)\overline{MR}$$

- For $\tilde{m} = 2$ with $Z_{1-\alpha/2} = 3$

$$\text{UCL} = \left(1 + \frac{3 \times 0.853}{1.128}\right)\overline{MR} = 3.267\,\overline{MR}, \qquad \text{LCL} = 0$$

since the raw LCL is negative

- If MR_i exceeds the UCL, we should investigate the process at time points i through $i + \tilde{m} - 1$ for special causes

Important remarks on the individual-observation charts

- With $m = 1$ we **lose the CLT protection** that makes subgroup means approximately normal, so normality of X_i should be checked with a histogram or normal probability plot before using these charts
- The standard choice is $\tilde{m} = 2$ because it maximises locality and minimises dependence between consecutive MR_i. Larger $\tilde{m}$ stabilises the estimate of σ but smooths out local changes
- There is **no separate s chart**: the MR chart plays the role of the variability chart. A moving SD would be just $MR_i/\sqrt{2}$ for $\tilde{m} = 2$, so it adds nothing
- Consecutive moving ranges **share observations** and are therefore dependent; run-rule patterns on the MR chart should be interpreted with caution
- Under **positive autocorrelation**, neighbouring observations are similar, so MR_i becomes artificially small, σ is underestimated, and the limits become too narrow, causing excessive false alarms

Example 3.3 Liquid cleaner concentration

- A chemical process produces a liquid cleaner whose quality characteristic is the **concentration (in g/l)** of the active ingredient
- We have $n = 20$ individual observations
- We build the X and MR charts using a moving window of size $\tilde{m} = 2$
- From the data we compute

$$\bar{X} = 72.38, \qquad \overline{MR} = 8.72$$

- With $d_1(2) = 1.128$, $d_2(2) = 0.853$, and $Z_{1-\alpha/2} = 3$ the limits become

$$X \text{ chart:} \quad \text{UCL} = 95.57, \quad \text{LCL} = 49.19$$

$$MR \text{ chart:} \quad \text{UCL} = 28.50, \quad \text{LCL} = 0$$

Example 3.3 Data and moving ranges

i	X_i	MR_i	i	X_i	MR_i
1	78.7	—	11	78.4	10.3
2	72.8	5.9	12	78.2	0.2
3	78.4	5.6	13	60.0	18.2
4	79.6	1.2	14	74.7	14.7
5	60.4	19.2	15	75.8	1.1
6	69.5	9.1	16	76.6	0.8
7	64.9	4.6	17	68.4	8.2
8	75.5	10.6	18	83.1	14.7
9	70.4	5.1	19	61.1	22.0
10	68.1	2.3	20	73.0	11.9
$\bar{X} = 72.38$			$\overline{MR} = 8.72$		

Example 3.3 X and MR charts

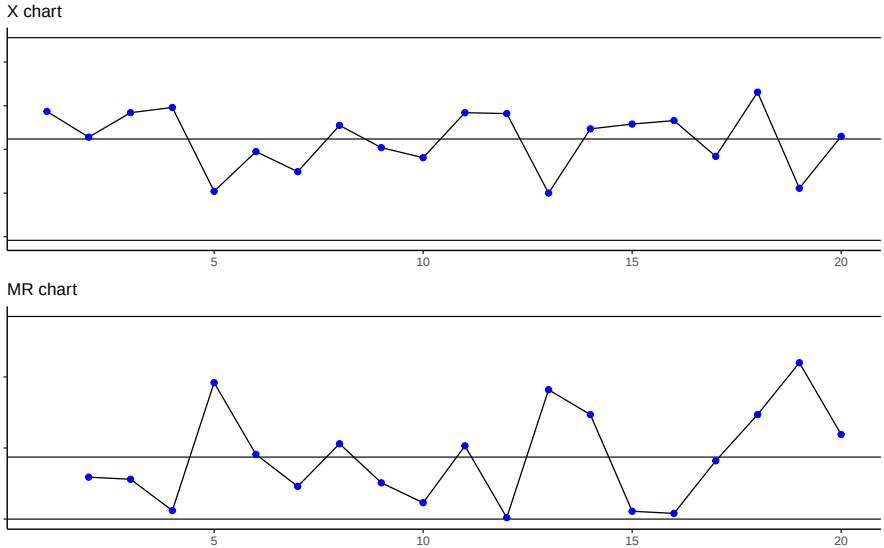

Example 3.3 Interpretation

- Both the X chart and the MR chart show **all points within limits**
- There is no signal of a shift in either the process mean or the process variability
- The limits are relatively **wide** compared to the data spread
  - This is typical of individual-observation charts: with $m = 1$ the standard error of the plotted statistic equals σ itself, rather than $\sigma/\sqrt{m}$
  - Shewhart charts for individual observations are therefore **less sensitive** to moderate shifts
- If faster detection of small shifts is needed, CUSUM or EWMA charts are better alternatives

Summary of Shewhart charts for numerical variables

Chart pair	Data type	Subgroup size	Variability estimator
$\bar{X}$ and R	grouped	$m \leq 6$	$\hat{\sigma} = \bar{R}/d_1(m)$
$\bar{X}$ and s	grouped	any m (prefer $m > 6$)	$\hat{\sigma} = \bar{s}/d_3(m)$
X and MR	individual	$m = 1$	$\hat{\sigma} = \overline{MR}/d_1(\tilde{m})$

- All three share the same structure: centre line $\pm$ $Z_{1-\alpha/2}$ times an estimated standard error
- The choice depends on the subgroup size and the type of data available
- In all cases: check the variability chart **first**, then the location chart