

Introduction

Operations managers are responsible for developing and maintaining the production processes that deliver quality products and services.

Once the production process is operating, it is often necessary to constantly monitor the process to ensure that it functions the way it was designed.

Introduction

The statistical methods we introduce in this Chapter represent the most commonly-used applications of Statistics in Business, Marketing and Management – at any point in time, there are *thousands* of firms applying these methods.

In this Chapter, we deal with the subject of **Statistical Process Control**.

Process variation

All production processes result in variation; that is, no product is exactly the same as another.

For example, if you weigh two boxes of 500 gram breakfast cereal,

- it is unlikely that either of them will weight exactly 500 grams;
- it is unlikely that they will have the same weight.

Process variation

All products exhibit some degree of variation. This could be

- ① **Chance variation** – caused by a number of randomly occurring events that are part of the production process and that cannot be eliminated without changing the process;
- ② **Assignable variation** – caused by specific events or factors that are frequently temporary and that can usually be identified and eliminated.

Example: *Dulux paint*

Most tins of *Dulux* paint are 5 litres in volume.

The cans are filled by an automatic valve that regulates the amount of paint in each can.

The designers of the valve acknowledge that there will be some variation in the amount of paint, *even when the valve is working as it was designed to work*. This is **chance variation**.

Occasionally, the valve will malfunction, causing the variation in the amount delivered to each can to increase. This increase is the **assignable variation**.

Process variation

The best way to understand what is happening is to consider the volume of paint in each can as a *random variable*.

If the only sources of variation are caused by chance, then each can's volume is drawn from identical distributions – that is, each distribution has the same *shape*, *mean* and *standard deviation*.

Under such circumstances, the production process is said to be **under control**.

Process variation

In recognition of the fact that variation in output will occur even when the process *is* under control and operating properly, most processes are designed so that their products will fall within designated **specification limits** or “**specs**”.

For example, the process that fills the paint cans may be designed so that the cans contain between 4.98 and 5.02 litres.

Process variation

Inevitably, some event or combination of factors in a production process will cause the process distribution to change.

When it does, the process is said to be **out of control**.

Process variation

There are several ways for the process to go out of control:

- 1 Level shift
- 2 Instability
- 3 Trend
- 4 Cycle

Process variation

1. Level shift

This is a change in the *mean* of the process distribution.

- Assignable causes: machine breakdown, new machine and/or operator, or a change in the environment
- *Dulux* example: a temperature or humidity change may affect the density of the paint, resulting in less paint in each can.

Process variation

2. Instability

This is when the *standard deviation* increases.

- Assignable causes: a machine in need of repair, defective materials, poorly trained operator
- *Dulux* example: Valve worn down: amount of paint delivered becomes more erratic

Process variation

3. Trend

When there is a slow, steady shift (either up or down) in the process distribution mean.

- Assignable causes: less-than-regular maintenance, operator fatigue, residue or dirt build up
- *Dulux* example: If the paint control valve becomes increasingly clogged, we would expect to see a steady decrease in the amount of paint delivered

Process variation

4. Cycle

This is a repeated series of small observations followed by large observations.

- Assignable causes: environmental changes, worn parts, or operator fatigue
- *Dulux* example: Changes in the voltage in the electricity that runs the machines could result in a series of overfilled cans and a series of underfilled cans

Process variation

The key to quality is to detect when the process goes out of control so that we can correct the malfunction and restore control of the process.

The **control chart** is the statistical method that we use to detect such problems.

Control charts

A **control chart** is a plot of statistics over time.

For example, an $\bar{x}$ -**chart** plots a series of sample means taken over a period of time.

Control charts

Each control chart contains a *centreline* and *control limits*.

The control limit above the centreline is called the **upper control limit** (UCL) and that below the centreline is called the **lower control limit** (LCL).

If, when the sample statistics are plotted, all points are randomly distributed between the control limits, we conclude that the process is **under control**.

If the points are not randomly distributed between the control limits, we conclude that the process is **out of control**.

Example: Paint tins

Suppose that in the *Dulux* paint tins example described previously, we want to determine whether the central location (mean) of the distribution has changed from one period to another.

We will draw our conclusion from an $\bar{x}$ -chart.

For the moment, let us assume that we know the mean μ and the standard deviation σ of the process when it is under control; $\mu = 5$ litres and $\sigma = 0.15$ litres.

Example: Paint tins

Samples of 20 cans are taken every hour for a period of 24 hours; the resulting sample means (in litres) are shown in the table below.

Time	13:03	14:03	15:03	16:03	17:03	18:03	19:03	20:03
$\bar{x}$	5.01	4.94	4.96	5.02	5.05	4.88	4.87	5.03
Time	21:03	22:03	23:03	00:03	01:03	02:03	03:03	04:03
$\bar{x}$	5.00	5.02	4.98	4.91	5.09	4.95	4.95	5.08
Time	05:03	06:03	07:03	08:03	09:03	10:03	11:03	12:03
$\bar{x}$	4.90	5.00	5.03	5.07	4.86	4.80	4.85	4.82

Example: Paint tins

Recall the *Central Limit Theorem* from MAS1403.

This states that, no matter what the distribution of our random variable, the mean $\bar{x}$ of a random sample $x_1, x_2, \dots, x_n$ from this distribution (provided n is large) is approximately Normal with mean μ and variance σ^2/n , i.e.

$$\bar{x} \sim N\left(\mu, \frac{\sigma^2}{n}\right),$$

where μ and σ^2 are the mean and variance (respectively) of the distribution of our random variable.

Example: Paint tins

In the paint tins example, we are told that $\mu = 5$ litres and $\sigma = 0.15$ litres. Means are taken every hour from samples of size 20, and so we have

$$\bar{x} \sim N\left(5, \frac{0.15^2}{20}\right) \quad \text{i.e.}$$

$$\bar{x} \sim N(5, 0.001125)$$

Also recall from MAS1403 that

- $\bar{x} \pm 1$ standard deviation gives about 68% coverage,
- $\bar{x} \pm 2$ standard deviations gives about 95% coverage, and
- $\bar{x} \pm 3$ standard deviations gives about 99% coverage

for the Normal distribution.

Control charts

Thus, the lower and upper control limits for an $\bar{x}$ -chart are often defined as:

$$\text{Lower control limit} = \mu - 3 \text{ standard errors} = \mu - 3\sqrt{\frac{\sigma^2}{n}} = \mu - 3\frac{\sigma}{\sqrt{n}}$$

$$\text{Upper control limit} = \mu + 3 \text{ standard errors} = \mu + 3\sqrt{\frac{\sigma^2}{n}} = \mu + 3\frac{\sigma}{\sqrt{n}}$$

Control charts

For the *Dulux* paint example, this would give:

$$\text{Lower control limit} = 5 - 3 \times \frac{0.15}{\sqrt{20}} = 4.899$$

$$\text{Upper control limit} = 5 + 3 \times \frac{0.15}{\sqrt{20}} = 5.101$$

Control charts

In Figure 3.1 construct an $\bar{x}$ -chart for the *Dulux* paint example. Identify when the process is *out of control*.

2-sigma and 3-sigma control limits

The control limits used in the previous example are called “3-sigma” control limits – we used a distance of 3 standard errors either side of the mean as a tool to detect the process being out of control.

Why not use “2-sigma” control limits?

Sometimes these *are* used, although the following reasoning should help explain why 3-sigma control limits are often preferred.

2-sigma and 3-sigma control limits

Suppose we have

H_0 : The process is under control

H_1 : The process is out of control

In statistical process control, a **Type I error** occurs if we conclude that the process is out of control when in fact it is not (*i.e. reject H_0 when in fact it is true*).

Expensive error: the production process is usually stopped and the causes of the variation are found and repaired – stopping the production process is costly, especially if there was no need to do so!

2-sigma and 3-sigma control limits

So, we would like the probability of a Type I error to be as small as possible. For the paint tins example, we can work out the probability of a Type I error as follows.

2-sigma and 3-sigma control limits

Using 3-sigma control limits:

$$\begin{aligned}\Pr(\bar{x} < \text{lower limit}) &= \Pr(\bar{x} < 4.899) \\ &= \Pr\left(z < \frac{4.899 - 5}{\sqrt{0.001125}}\right) \\ &= \Pr(z < -3) \\ &= 0.0013\end{aligned}$$

2-sigma and 3-sigma control limits

Also

$$\begin{aligned}\Pr(\bar{x} > \text{upper limit}) &= \Pr(\bar{x} > 5.101) \\ &= 1 - \Pr(\bar{x} < 5.101) \\ &= 1 - \Pr\left(z < \frac{5.101 - 5}{\sqrt{0.001125}}\right) \\ &= 1 - \Pr(z < 3) \\ &= 1 - 0.9987 \\ &= 0.0013.\end{aligned}$$

2-sigma and 3-sigma control limits

Thus

$$\begin{aligned}\Pr(\text{Type I error}) &= 0.0013 + 0.0013 \\ &= 0.0026\end{aligned}$$

2-sigma and 3-sigma control limits

Of course, we could have turned to the standard Normal distribution immediately and, owing to the symmetry of the Normal distribution:

$$\begin{aligned}\Pr(\text{Type I error}) &= 2 \times \Pr(z < -3) \\ &= 2 \times 0.0013 \\ &= 0.0026.\end{aligned}$$

2-sigma and 3-sigma control limits

Similarly, if we use 2-sigma control limits, we would have

$$\begin{aligned}\Pr(\text{Type I error}) &= 2 \times \Pr(z < -2) \\ &= 2 \times 0.0228 \\ &= 0.046.\end{aligned}$$

Thus, using 3-sigma control limits will give a **false alarm** in about 0.26% of samples.

Using 2-sigma control limits increases this to about 4.6% of samples.

2-sigma and 3-sigma control limits

A **Type II error** occurs when we do not detect when the process is out of control (*i.e. we retain H_0 when in fact it is false*).

Average run length

The **Average Run Length** (ARL) is the expected number of samples that must be taken before the control chart indicates that the process has gone out of control.

The ARL is determined by

$$ARL = \frac{1}{P}$$

where P is the probability that a sample mean falls outside the control limits.

Average run length

Assuming that 3-sigma control limits are used, this would give

$$ARL = \frac{1}{0.0026} = 385.$$

This means that when the process is under control, the $\bar{x}$ -chart will erroneously conclude that it is *out* of control once every 385 samples, on average.

If the sampling plan calls for samples to be taken once every hour (as in the paint tins example), on average there will be a *false alarm* once every 385 hours.

Examples

1. Find the upper and lower 2-sigma control limits for the *Dulux* paint example. Also find the *ARL* for such a control chart.

Recall that we have

$$\mu = 5 \text{ litres}$$

$$\sigma = 0.15 \text{ litres}$$

$$n = 20.$$

Thus, we have

$$\text{Lower control limit} = 5 - 2 \times \frac{0.15}{\sqrt{20}} = 4.9329$$

and

$$\text{Upper control limit} = 5 + 2 \times \frac{0.15}{\sqrt{20}} = 5.0671$$

Examples

The *ARL* is given by

$$ARL = \frac{1}{P}$$

where, from previous calculations, $P = 0.046$. Thus,

$$ARL = \frac{1}{0.046} = 22;$$

this means that when the process is under control, there will be a false alarm once every 22 samples – compared to once every 385 samples when using 3-sigma control limits!

Examples

2. If the control limits of an $\bar{x}$ -chart are set at 2.5 standard errors from the centreline, what is the probability that on any sample the control chart will give a false alarm?

Setting the control limits at 2.5 standard errors gives

$$\begin{aligned}\Pr(\text{Type I error}) &= 2 \times \Pr(z < -2.5) \\ &= 2 \times 0.0062 \\ &= 0.0124.\end{aligned}$$

Examples

3. A firm manufactures notebook computers. For one component, the company draws samples of size 10 every half an hour. The company makes 4000 of these components every hour. The control limits of the $\bar{x}$ -chart are set at 3 standard errors from the mean. When the process goes out of control, it usually shifts the mean by about 0.75 standard deviations.
- (a) On average, how many units will be produced until the control chart gives a false alarm?

Examples

$$\begin{aligned}\Pr(\text{False alarm}) &= 2 \times \Pr(z < -3) \\ &= 2 \times 0.0013 \\ &= 0.0026.\end{aligned}$$

The *ARL* is

$$ARL = \frac{1}{0.0026} = 385,$$

so we would get a false alarm once every 385 samples.

Since samples are taken every half an hour, this would give a false alarm once every 192.5 hours – in this time, the firm would have produced

$$192.5 \times 4000 = 770,000 \text{ units.}$$

Examples

- (b) Find the probability that the $\bar{x}$ -chart does not detect a shift of 0.75 standard deviations on the first sample after the shift occurs.

Recall that the standard normal distribution is $Z \sim N(0, 1)$. If a shift in the mean by 0.75 standard deviations occurs, that gives us a new random variable, say Z_1 , where $Z_1 \sim N(0.75, 1)$.

Examples

If the $\bar{x}$ -chart does not detect this, then the values plotted remain inside the control limits – in this case, within the range $-3 \leq Z_1 \leq 3$. So we want

$$\begin{aligned}\Pr(-3 \leq Z_1 \leq 3) &= \Pr(Z_1 \leq 3) - \Pr(Z_1 \leq -3) \\&= \Pr\left(Z \leq \frac{3 - 0.75}{1}\right) - \Pr\left(Z \leq \frac{-3 - 0.75}{1}\right) \\&= \Pr(Z \leq 2.25) - \Pr(Z \leq -3.75) \\&= 0.9878 - 0 \\&= 0.9878.\end{aligned}$$

Control charts

In the examples presented so far, we have assumed that the process parameters (μ and σ) were known.

When the parameters are *unknown*, we estimate their values from the sample data.

In the next section, we discuss how to construct and use control charts in more realistic situations.

We will re-visit $\bar{x}$ -charts, as well as S -charts, and we call these *control charts for variables*.

In section 3.5 we will also consider *control charts for attributes*.

Control charts for variables

There are several ways to determine whether a change in the process distribution has occurred.

To determine whether the distribution means have changed, we employ the $\bar{x}$ -chart.

To determine whether the process distribution standard deviation has changed, we can use the S -chart or the R -chart (S for standard deviation, R for range).

$\bar{x}$ -chart

In the previous section we determined the centreline and the control limits of an $\bar{x}$ -chart using the mean and standard deviation of the process distribution.

However, it is unrealistic to believe that the mean (μ) and standard deviation (σ) of the process distribution will be known.

Thus, to construct the $\bar{x}$ -chart, we need to *estimate* the relevant parameters from the data.

$\bar{\bar{x}}$ -chart

We begin by drawing samples when we have determined that the process is under control.

For each sample, we compute the mean and standard deviation.

The estimator of the mean of the distribution is the mean of the sample means (denoted $\bar{\bar{x}}$ and pronounced “x bar bar” or “x double bar”), given by

$$\bar{\bar{x}} = \frac{\sum_{j=1}^k \bar{x}_j}{k},$$

where $\bar{x}_j$ is the mean of the j th sample and there are k samples.

$\bar{x}$ -chart

To estimate the standard deviation of the process distribution, we calculate the “pooled standard deviation”:

$$S = \sqrt{\frac{\sum_{j=1}^k s_j^2}{k}}.$$

Previously, we used:

$$\text{centreline} = \mu$$

$$\text{Lower control limit} = \mu - 3\frac{\sigma}{\sqrt{n}}$$

$$\text{Upper control limit} = \mu + 3\frac{\sigma}{\sqrt{n}}$$

$\bar{x}$ -chart

Because the values of μ and σ are now unknown, we *now* use:

- $\bar{\bar{x}}$ in place of μ
- S in place of σ

Therefore, the centreline and control limits are now:

$$\begin{aligned}\text{centreline} &= \bar{\bar{x}} \\ \text{Lower control limit} &= \bar{\bar{x}} - 3 \frac{S}{\sqrt{n}} \\ \text{Upper control limit} &= \bar{\bar{x}} + 3 \frac{S}{\sqrt{n}}\end{aligned}$$

Example: Car Giant car seats

Car Giant manufactures seats for Mazda and Lexus cars. One of the components of a front seat cushion is a wire spring, produced from 4mm steel wire.

A machine is used to bend the wire so that the spring's length is 500mm.

If the springs are longer than 500mm they will loosen and eventually fall out. If they are too short, they won't easily fit into position.

To determine whether the process is under control, random samples of four springs are taken every hour. The last 25 samples are shown in the table in your notes.

Example: Car Giant car seats

Complete the table by finding the mean and standard deviation for samples 21–25.

You should use the Statistics mode on your calculator for this!

Example: Car Giant car seats

The mean of the means – $\bar{\bar{x}}$ – is given by

$$\begin{aligned}\bar{\bar{x}} &= \frac{502.027 + 499.015 + \dots}{25} \\ &= 500.296.\end{aligned}$$

Also, the pooled standard deviation S is found as

$$\begin{aligned}S &= \sqrt{\frac{1.56689^2 + 0.83309^2 + \dots}{25}} \\ &= \sqrt{\frac{97.125}{25}} \\ &= 1.971\end{aligned}$$

Example: Car Giant car seats


Thus, the centreline and control limits are

$$\text{Centreline} = \bar{\bar{x}} = 500.296$$

$$\text{Lower control limit} = \bar{\bar{x}} - 3 \frac{S}{\sqrt{n}} = 500.296 - 3 \times \frac{1.971}{\sqrt{4}} = 497.34$$

$$\text{Upper control limit} = \bar{\bar{x}} + 3 \frac{S}{\sqrt{n}} = 500.296 + 3 \times \frac{1.971}{\sqrt{4}} = 503.253$$

Example: Car Giant car seats



S-charts

The S -chart graphs sample standard deviations to determine whether the process distribution standard deviation has changed.

The format is similar to that of the $\bar{x}$ chart – the S chart displays a centreline and control limits.

However, the formulae for the centreline and control limits are more complicated than those for the $\bar{x}$ -chart – consequently, we will let Minitab do the work here!

Monitoring the production process

In this section, we have introduced $\bar{x}$ -charts and S -charts as separate procedures.

In actual practice, the two charts must be drawn and assessed together.

The reason for this is that the $\bar{x}$ -chart uses S to calculate the control limits; consequently, if the S -chart indicates that the process is out of control, the value of S will not lead to an accurate estimate of the standard deviation of the process distribution.

Monitoring the production process

The usual procedure is:

- 1 Draw the S -chart first;
- 2 If this indicates that the process is under control, then draw the $\bar{x}$ -chart;
- 3 If the $\bar{x}$ -chart also indicates that the process is under control, we can then use both charts to maintain control;
- 4 If either chart shows that the process was out of control at some time during the creation of the charts, we can detect and fix the problem and then re-draw the charts with new data.

Back to Car Giant car seats

We have determined that the process is under control – none of the points in the S -chart nor the $\bar{x}$ -chart fell outside the control limits.

We now use the statistics generated in the creation of these charts to monitor the production process.

Back to Car Giant car seats

Recall that the sampling plan calls for samples of size 4 every hour.

The following table lists the lengths of the springs taken during the *next* six hours.

Sample	Spring 1	Spring 2	Spring 3	Spring 4	$\bar{x}_j$	s_j
26	502.653	498.354	502.209	500.080		
27	501.212	494.454	500.918	501.855		
28	500.086	500.826	496.426	503.591		
29	502.994	500.481	502.996	503.113		
30	506.549	508.708	502.480	504.442		
31	500.441	508.887	505.005	503.254		

Back to Car Giant car seats

1. Complete the last two columns of this table.
2. Update the S -chart and $\bar{x}$ -chart shown in figures 3.5 and 3.6 below (respectively).
3. Comment.