Quality Control Management

1. Introduction and motivation

We associate the idea of 'quality' with almost everything. Quality has different meanings under different contexts. For example:

- ***Performance***: A product that performs better than others designed for the same purpose can be said to have better quality; for example, sound quality from two MP3 players playing the same song can be quantitatively compared. Another aspect of performance may be in terms of the number of features provided by a product; e.g a mobile phone with a camera and MP3 player compared to one that can only be used as a phone.
- ***Reliability:*** A product that often breaks down unexpectedly, and therefore requires frequent repair or maintenance is said to be unreliable (i.e. poor in quality).
- ***Durability:*** A longer expected service life implies higher durability (i.e. better quality).
- ***Aesthetics:*** Visual appeal, including factors such as style, color shape etc. are related to the product's aesthetics. For example, the popular iPod MP3 player from Apple is popularly believed to have good aesthetics.

Every company wishes to **control the quality** of the products it makes or the services it provides. Before we see how to do so, here is an example of why it is important to do so.

Example 1: This example is based on a real case study performed in ~1980 for a car manufacturer in USA. For a particular model, it used transmissions (gear-box) manufactured by two factories, one in Japan and one in USA. The study revealed that the transmissions produced by the USA factory had almost 4-times the total warranty repair costs than the ones made in Japan.

The smooth operation of transmission depends on proper fitting between several mechanical parts -- loose fitting parts tend to vibrate more during use and possibly cause earlier break-down. The company dis-assembled several transmissions from both factories, and measured the part dimensions that were critical for good performance. They plotted the variations of size of these critical dimensions, and in several cases, the graph looked as follows.

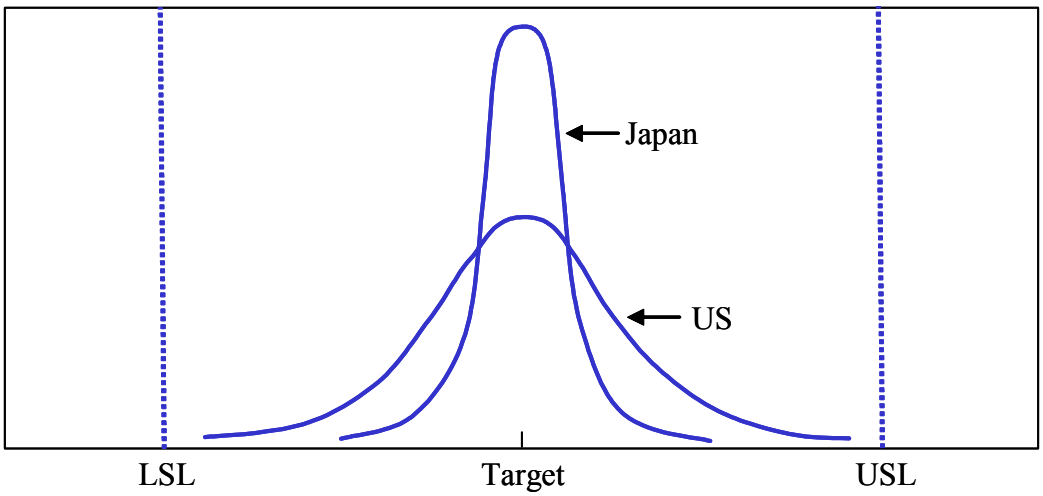

Figure 1. Distribution of critical dimensions from transmissions [source: Montgomery]

Notice that most parts produced by either factory are acceptable (i.e. the measured dimension falls between the range between the **LSL (lower specified limit)** and the **USL**

(upper specified limit). However, the parts from the Japan factory used only about 25% of the allowed tolerance band, while the US factory parts were spread over ~75% of it. Further experiments verified that the lower variability was the reason for lower break-down rate of the Japanese transmissions.

At first, it appears that both factories are producing parts whose dimensions vary within the specified limits, so both are of acceptable quality. But it is clear that the Japanese parts were somehow performing better. This leads to the following definitions.

Definitions:
Quality is inversely proportional to variability

Quality improvement is the reduction in variability of products/services.

But how can we reduce variability? We shall look at three methods used together or separately by modern industry. All of these methods use tools from statistics. In this topic, I will introduce some basic methods for improvement or control of quality, and also some fundamental ideas from statistics that form the basis of these methods.

2. Techniques for Quality Management

1. Acceptance **sampling**
2. Statistical Process Control (**SPC**)
3. Robust Design through Design of Experiments (**DOE**)

2.1 Sampling and some basic statistics

In most cases, we will associate the idea of quality with some property/properties of a product/service. We will denote the best possible quality with a target value, and any deviation from this value (i.e. variation) will indicate the level of quality (larger deviation implies lower quality). Of course, this measure is applied to each product we make, and it is tedious to look at the quality level of individual items -- therefore we shall try to get an estimate of the overall quality level by looking at a summary of the data corresponding to all the products -- namely a statistical summary of the data. In particular, two statistical estimates of interest are the **mean (denoted as μ)** and the **standard deviation (denoted as stdev, or σ)**.

Example 2. *Average value (mean) and spread (standard deviation)*
Given a list of n numbers, $a_1, \ldots, a_n$; say: 19, 21, 18, 20, 20, 21, 20, 20.

The mean = $\mu = \Sigma a_i / n = (19+21+18+20+20+21+20+20) / 8 = 19.875$

The variance $\sigma^2 = \dfrac{\sum (a_i - \mu)^2}{n} \approx 0.8594$

The standard deviation = $\sigma = \sqrt{\dfrac{\sum (a_i - \mu)^2}{n}} = \sqrt{(\sigma^2)} \approx 0.927.$

Example 3. I install a new air-conditioning system in my house to cool the living room and bedroom. I set it to 20°, and let it run for some time. Suppose now I want to know the average temperature in a room.

Ideally, I need to measure the temperature at every point, and then compute the mean of this data; this has some problems -- there are infinite points in the room; and the temperature at any point changes with time. Suppose I measure the temperature at 5 different locations in each room. The readings are as follows:
Living Room: 18, 19, 20, 21, 22.
Bedroom: 19, 20, 20, 20, 19.

Suppose we now ask: what is the average temperature in the living room?
We cannot give an accurate answer, since we did not measure the temperature at each point. However, the few readings that we took can give us ***an estimate*** of the average temperature. Let's denote this estimate as:

$$m = \Sigma a_i / n = (18+19+20+21+22) / 5 = 20.$$

NOTE: m is not the true mean, μ, of the living-room temperature, it is an estimate.

The natural question to ask now is: how good an estimate is m?
The answer: it depends on a variety of factors, such as
- the distribution of temperature across the room (suppose some corner is really hot, but we did not measure anywhere close to it);
- how many data points we measured; etc.

In general, though, if the measuring points are randomly distributed across the room, then m is an unbiased estimator of μ. In other words, if we repeatedly pick 5 (or some number n) measuring points, and calculate m, it may actually be slightly more or less that μ; however, the average of a large number of such measures will coincide with the true mean, μ.

Likewise, since it is clear that the temperature at different points in the room is different, it is natural to ask what is the variance (or standard deviation) of the temperature. Once again, we do not have all the data to find the exact *stdev* σ, but we can certainly get an estimate, using the data we have.

Thus we write the stdev of the sample as: $s_n = \sqrt{\dfrac{\sum (a_i - m)^2}{n}}$.

Note that here we used the sample mean, m, instead of μ (which is unknown). Again, how good is s_n as an estimator of the true stdev σ?

Again, the same considerations apply as for the sample mean. Further, it turns out that s_n underestimates the true stdev σ on the average. In other words, s_n is a biased estimator of σ. Suppose s_1 is the stdev based on temperature of five (randomly selected) points in the room; s_2 is the same for another five random points, etc. Then, the average value of a large number of such calculations, $s_1, s_2, \ldots s_n$ will be lower than the true stdev σ for the

room. The mathematical proof of this is not difficult, but here we only provide the following result:

The **unbiased estimator of stdev** ***of a sample*** $= s = \sqrt{\frac{\sum (a_i - m)^2}{n-1}}$.

How do we set up a quality control system with sample testing?

Step 1. Identify those properties of the product that define quality.
Step 2. Find the best method/instrument to measure these properties.
Step 3. Validate the measurement system.
Step 4. Measure the product.

The simplest idea is to test every product (100% sampling). However, ***100% sample testing is often NOT practical***:
(i) It may be too expensive/time-consuming to measure each product. Consider, for example, a manufacturer of nuts and bolts (small, inexpensive items produced in very large numbers).
(ii) It may not be practical. Consider a company that manufactures soda cans (e.g. Coca Cola). Each sealed can should have a pressure of 50PSI (pounds per square inch) at 10°C. Obviously it is not practical to test the pressure in the can once it is sealed -- however, a small percentage of randomly selected cans may be picked out and tested. How many should we test? Which cans should we select?

In order to answer these questions, we should first have some idea about what type of PSI values we ***expect*** to see in our soda cans. Consider our soda can example. Suppose that we measure the exact pressure in many cans. We may expect to see values such as: 48.234, 49.213, 51.87, 50.023,… What is the probability that a can has pressure equal to 50 PSI? Clearly, this chance is zero (no matter how accurately we fill the can, it's exact pressure cannot exactly equal 50.00000 -- although it is possible that many cans have pressure in the range, say, 49.999 and 50.001. We may denote the pressure in a soda can as a variable that can take any real value -- we expect this value to be spread out in some range, say 15 PSI(≈ 1 atmosphere pressure, i.e. the can has a leak) to 95 PSI (above this level, the can will burst). Suppose we break this range into steps, e.g. [15, 17.5], [17.5, 20], … [87.5, 95], and plot the percentage of cans with pressure in each range -- the plot may look somewhat like Figure 2a below. Note that the sum of all the percentages = 1 (i.e. 100%). We could use smaller ranges, e.g. [15, 16], [16,17], … [94,95]. As we keep shrinking the range, in the limit, we will end up with a function called the ***probability density function (pdf)***, which may look somewhat like the Figure 2b below. As before, the area under the curve should equal 1 (i.e. 100%).

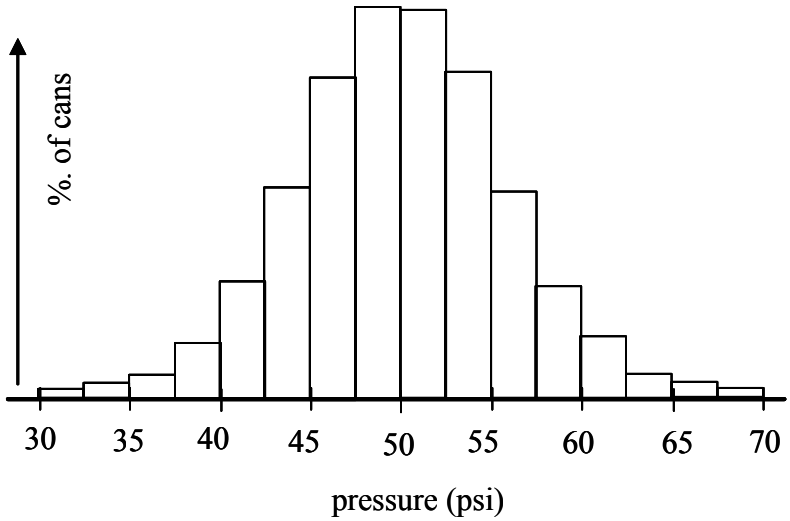

Figure 2 (a) Histogram showing % of soda cans with gas pressure in different ranges

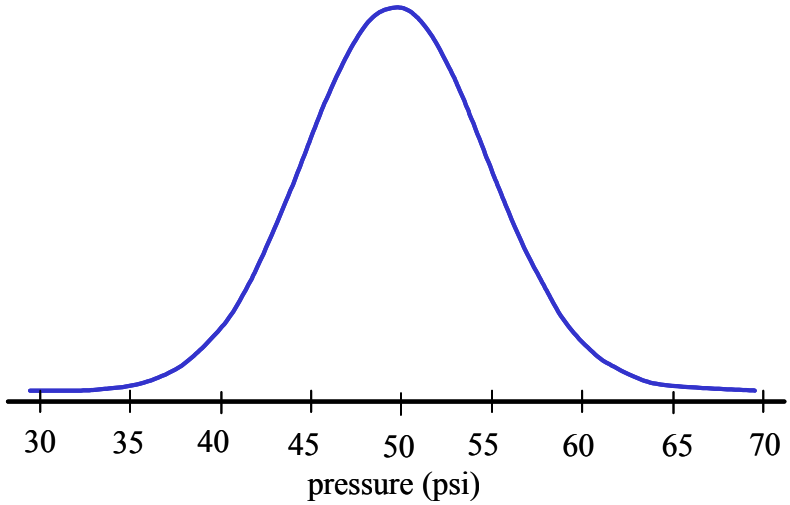

Figure 2 (b) The probability density function for gas pressure in soda cans

The pdf in Figure 2b above is a special shape -- it's called the **Normal** (or **Gaussian**) **Distribution Function**, and is often also called the bell curve because of its shape. The mathematical definition of this pdf, for a random variable, z, selected from a distribution with mean μ and stdev σ is: $\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(z-\mu)^2}{2\sigma^2}}$

Many types of things in nature have measurements that are distributed according to this shape -- e.g. the heights of all people in HK, the weight of soda in unopened Pepsi™ cans, etc. Why is this so? The answer is (most likely) related to a very important theorem in statistics, called the Central Limit Theorem. There are several forms of this theorem, two are important to us. Let's look at these.

The Central Limit Theorem:

Suppose that a random variable, X, is distributed according to some function with given mean, μ, and variance, σ^2 (the pdf may have any shape). If we randomly plot a histogram using many values of X, we expect to see the shape of the pdf. Let S_n = sum of n

randomly selected values of X; Suppose we draw a histogram of S_n. If S_n has a finite variance (which is true for anything we may measure in real life), then the Central Limit Theorem states that S_n will follow a normal distribution with mean = nS_n, and variance = $n\sigma^2$.

We will not prove this theorem here, but let's look at a simple example. Suppose, our variable can take value = -1, 0, 1, each with probability = 1/3. The pdf for this variable is shown in Figure 3 (three spikes, at -1, 0 and 1). Let S_2 = sum of two samples for X. Its distribution will look like Figure 4. Likewise, if we consider the sum of three values of X, S_3, and plot its pdf in Figure 5. you can see that as n increases, the distribution appears closer in shape to the familiar bell-shaped Gaussian curve!

$$X_1 = \begin{cases} -1, & \text{with probability } 1/3 \\ 0, & \text{with probability } 1/3 \\ 1, & \text{with probability } 1/3 \end{cases}$$

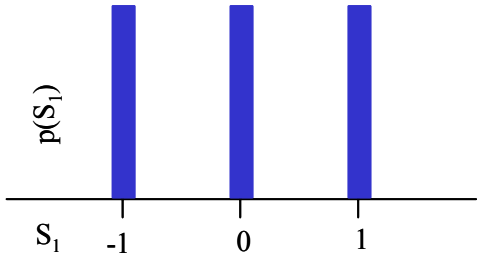

Figure 3. Distribution of a random variable, X_1 with three discrete values

X_1	X_2	$X_1 + X_2$
-1	-1	-2
-1	0	-1
-1	1	0
0	-1	-1
0	0	0
0	1	1
1	-1	0
1	0	1
1	1	2

$$X_1 + X_2 = \begin{cases} -2, & \text{with probability } 1/9 \\ -1, & \text{with probability } 2/9 \\ 0, & \text{with probability } 3/9 \\ 1, & \text{with probability } 2/9 \\ 2, & \text{with probability } 1/9 \end{cases}$$

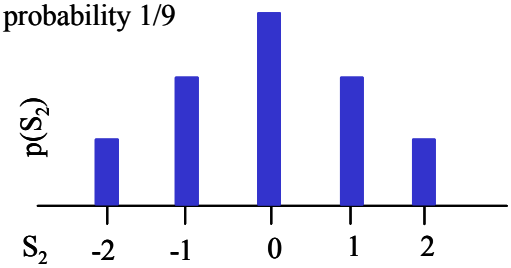

Figure 4. Distribution of random variable $S_2 = (X_1 + X_2)$

$$X_1 + X_2 + X_3 = \begin{cases} -3, & \text{with probability } 1/27 \\ -2, & \text{with probability } 3/27 \\ -1, & \text{with probability } 6/27 \\ 0, & \text{with probability } 7/27 \\ 1, & \text{with probability } 6/27 \\ 2, & \text{with probability } 3/27 \\ 3, & \text{with probability } 1/27 \end{cases}$$

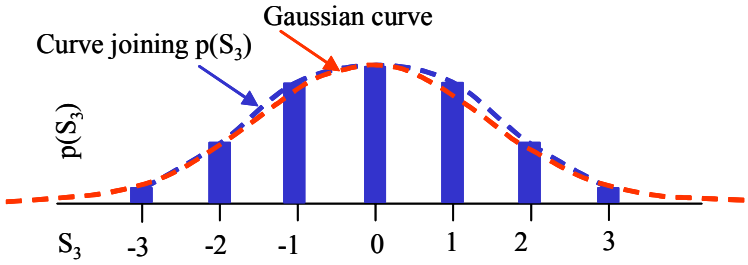

Figure 5. Distribution of $S_3 = X_1 + X_2 + X_3$ is approaching the shape of Gaussian curve

The above is true when, as in the example of Figures 3-5, the variables X_i are all from the *same distribution*. However, ***a weaker form of Central Limit Theorem exists:*** even if the values of X_i come from very different underlying distributions, their normalized sum, $S_n = (X_1 + X_2 + \ldots + X_n)/n$ approaches a normal distribution. In the figure below, $\overline{X_5} = (X_1$

$+ X_2 + \ldots + X_5)/5$, where X_i are selected randomly from five very different distributions, each with $\mu = 0$ and $\sigma = 1$. As you can see, the distribution of $\overline{X_5}$ follows the Gaussian curve.

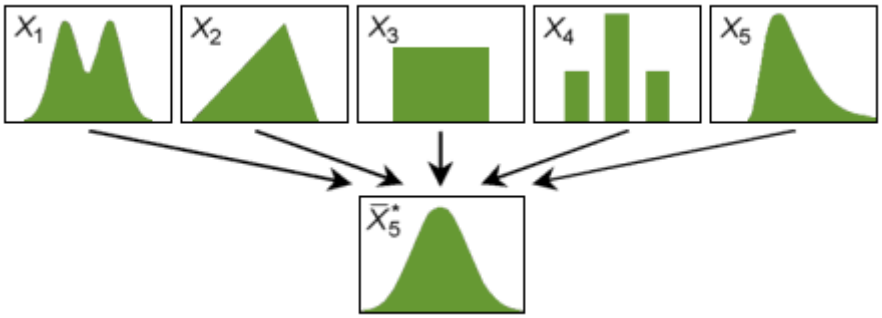

Figure 6. Central Limit Theorem: the normalized average of five independent random variables is still approximately normal despite the fact that the five random variables have very different distributions. [source: www.riskglossary.com]

Finally, there is a (somewhat weaker, but practically useful) form of the Central Limit Theorem that allows for even the mean and variance of the X_i to be different, and still the distribution of their normalized sum follows the Gaussian curve!

What is the use of all this?

Many, many practical measurements follow the normal (Gaussian) distribution. Consider the variations in the pressure in soda cans. It is likely that there are several (more or less independent) factors that cause the variations -- e.g. the amount of gas dissolved into the soda, the size of can, the amount of the liquid in the can, temperature during sealing, etc. Each of these factors may have their own pdf, yet the net effect on the pressure in the can is the *sum of these effect*s. By the central limit theorem, of there are enough number of such factors causing the variations, then it's quite likely that the net effect, namely the variation in pressure of the can, will be distributed normally!

Gaussian curve scaling

As we saw earlier, the expression for the Gaussian curve is quite complex, and in fact we cannot solve for this integral using analytical methods. So how to use the normal curve? To do so, we use a simple, two-step method:

Suppose we have a **random variable, *z***, that is **normally distributed with a mean μ and variance σ^2**. Then the **variable $(z - \mu)/\sigma$** will be **distributed normally, with a $\mu = 0$ and $\sigma = 1$**.

The shape of the standard Gaussian distribution with $\mu = 0$ and $\sigma = 1$ is fixed -- so its value at each point has been calculated (using approximation methods, but to very high accuracy), and is stored in **standard Normal tables**, denoted ***N(0, 1)***. Recall that the

Gaussian pdf $= \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(z-\mu)^2}{2\sigma^2}}$, therefore:

- The standard Gaussian pdf = $\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}$

- The standard Gaussian curve is symmetric about the z=0 axis

- The probability that a variable X ∈ N(0, 1) has value ≥ a is calculated by the area under the standard curve from z=a to z=∞ (see figure below).

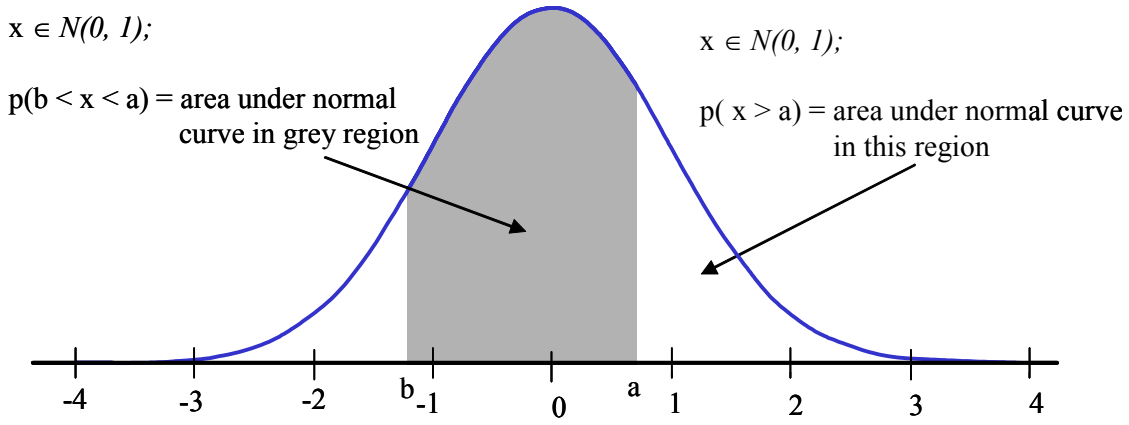

Figure 7. Interpretation and usage of the standard Gaussian

Therefore, most statistical and QC books will list a table, which gives the area under the curve, that is, the value of the $F(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{z^2}{2}} dz$ for different values of x. Your lecture notes webpage has a link to a pdf file with the numerical values of the area under this curve for different values of x. Below is an example of its use.

Example 2. A manufacturer of long life milk estimates that the life of a carton of milk (i.e. before it goes bad) is normally distributed with a mean = 150 days, with a stdev = 14 days. What fraction of milk cartons would be expected to still be ok after 180 days?

Solution: Let Z = 180 days. $(Z - \mu)/\sigma$ = (180 - 150)/14 ≈ 2.14; from the tables, we see that corresponding to Z = 2.14 we have an area = 0.9838. So we expect that the fraction of milk cartons that will be ok beyond $Z = 180$ days, or $Z = \mu + 2.14\sigma$, is 1 - 0.9838 = 0.0162.

Definition: ***distribution of the mean of a sample***
Suppose some random variable X is normally distributed with a mean μ and stdev σ. Suppose we take a random sample of n values of X, and it has mean m. If we repeatedly take random samples of size n, each time we shall obtain a different mean, m_i. Thus we can think of m itself as a random variable. Then this random variable m_i is normally distributed, with mean = μ, and stdev = $\sigma/\sqrt{n}$.

In other words, the reliability of m as an estimate of μ is related to $\sigma/\sqrt{n}$, which is therefore also called the *standard error of the mean.*

Example: Quality control of raw materials
A logistics company buys Shell-C brand diesel for its trucks. They know that each full tank of fuel allows the truck to travel 510 Km, with a standard deviation of 31 Km. Suppose a new seller promises to provide a cheaper fuel, Caltex-B, with the claim that it will give similar mileage as the Swell-C.
(i) What is the probability that the *mean distance* traveled over 40 full-tank journeys of Shell-C is between 500 Km and 520 Km?
(ii) Suppose that the *mean distance* covered by a sample of 40 full-tank journeys using Caltex-B turns out to be 495 Km. What is the probability that Caltex-B is equivalent to Shell-C?

Solution:
(i) From the results above, the mean distance covered by a sample of the paint will be normally distributed, with $\mu=510$, and for a sample size n=40, the $\sigma_{40}= \sigma/\sqrt{40} = 31/\sqrt{40} \approx$ 4.9 Km.
Thus we want to find the area under the standard normal curve between:
$z= (500 -510)/4.9 \approx -2.04$ and $z = (520 - 510)/4.9 \approx 2.04$.

The area = 1 - ((1 - 0.9793) + (1 - 0.9793)) = 0.9586.
In other words, the probability that the mean distance covered over 40 full tank journeys lies between 510 ± 10 Km is almost 95.86%.

(ii) Let us compute the probability that the mean distance over 40 journeys is ≤ 495. A sample mean of 495 corresponds to $z = (495 - 510)/4.9 \approx -3.06$. From the tables, the probability that the mean distance of 40 full tank journeys using fuel similar to Shell-C will be larger than 495 $m^2 = 0.9989$.
Thus the likelihood that the new fuel is equivalent to the older one is (1 - 0.9989) = 0.0011, which is less than 1%.

———

The above example shoes how sampling can be used to compare quality of samples of known materials over time. In this case, it was useful that we knew the values of μ and σ.

Various results in statistical theory allow us to make similar analysis regarding statistics that we collect to measure quality of some product/process under different conditions (For example: suppose we know the value of m for the actual distribution, but don't know the σ; how can we use the sample *s* instead of the unknown σ to make logical quality assessments? etc.). If you take a course in applied statistics, or in statistical quality control, you will study such issues in detail.

There are many types of quality control issues we can investigate, depending on the company and what it is doing. Here, we only look at one or two typical problems that are encountered.

Hypothesis Testing
(a) Suppose that a machine is set to fill bottles with soda. The machine operator sets the soda pumping pressure and time, and claims that the average volume of soda in each

bottle will be 250 ml. By only measuring data from some (not all bottles), how can we determine whether the settings are correct or not?

(b) Suppose further that the operator claims that his settings are accurate enough, such that no more than 0.01% of the bottles will have less than 245 ml of soda. How do we test the accuracy of this claim?

Let us restate the claim made in (a) as follows: the mean volume of soda in the bottles ≥ 250ml.
How do we evaluate whether this claim is true or false?

Suppose we randomly pick up 5 bottles, measure the volume in each, and compute the mean, which turns out to be 246 ml.

We do know that the actual volume varies a little bit (and we shall assume that it is normally distributed). Of course, 246 < 250, so we are tempted to reject the claim of the operator -- namely, we are tempted to say that the mean volume of soda in the bottles is in fact less than 250. *But of course, it is* ***possible*** *that the mean of this sample being a little low is entirely due to chance, and in fact the machine settings are good enough, and the claim is indeed true even though the mean of this sample turned out to be low*!

In setting the machine to produce acceptable quality filling, we need to first determine a ***criterion that will allow us to make decisions*** whether the machine settings need to be changed or not. Perhaps we may determine that if $\mu \geq 250$, we accept the current settings; if on the other hand, $\mu \leq 248$, we are sure that the consumer council will catch us and possibly take legal action against us -- so we must reject the claim, and adjust the machine settings. Finally, if μ is in a small range just under the claimed value, e.g. if it is in the range, say 248-250ml, we may be indifferent -- namely, we may decide not to take any action at this time, and it is unlikely that there will be any customer dissatisfaction or regulatory action against the company.

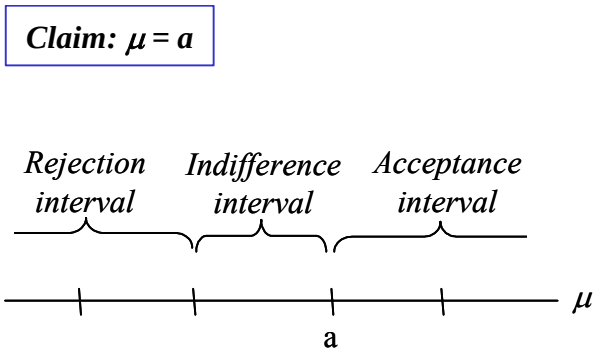

Figure 8. Decision intervals

Now, if the true mean for the soda volume is known, it is easy to act based on the accepted ***decision criterion***. But because we don't know μ, we shall have to *estimate it using samples*. Because we use estimates instead of actual values, there is always some chance that we make the decision. There are four possible cases to consider.

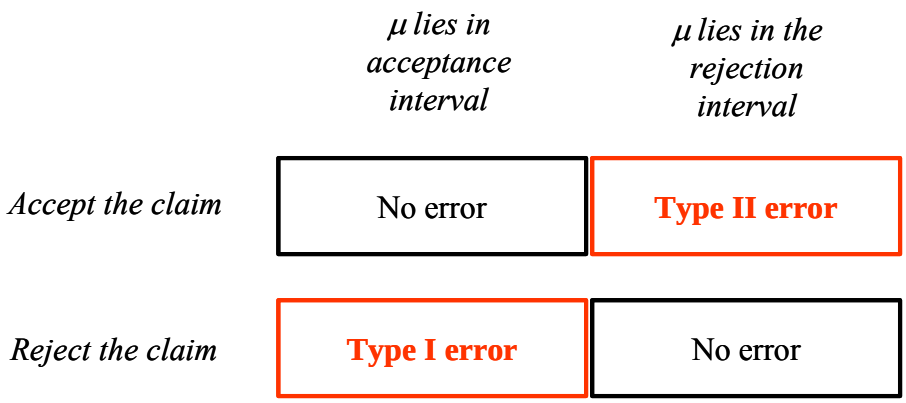

Figure 9. Possible cases and main error types in testing claims (hypothesis)

Example: Hypothesis testing and quality control *[source: Montgomery]*
Soda bottles are manufactured to be able withstand an internal pressure of at least 175 psi before they fail (i.e. leak or burst). The bottling company wants to know if the mean pressure strength of a batch (called a *lot*) exceeds 175 psi. From previous experience, it is known that the stdev of the pressure strength is 10 psi.

The bottling company wants to test the ***Hypothesis: $\mu > 175$***.

Suppose we decide to test a sample of 25 randomly selected bottles on a Hydrostatic pressure machine. The machine slowly increases the pressure in the bottle till the bottle fails. For the sample, the mean failing pressure, m = 182 psi.

From our earlier results, we know that $z = (m - \mu)/(\sigma/\sqrt{n})$ follows the standard normal distribution.

Suppose we want to be 95% certain against making a Type I error -- that is, the actual (but unknown) $\mu > 175$, but, based on the data from our sample, we reject the lot.

The only issue here is that we don't know μ. However, we want to guard against the chance of making a Type I error (i.e. a false rejection), namely, the probability that $m \geq 182$ when the $\mu \leq 175$. Notice that the probability of making such an error is highest when $\mu = 175$. Therefore, it is sufficient to consider that the acceptance interval is contracted to a single point, $\mu = 175$; now, what is the probability that we get a sample mean of 182, for this value of m?
Looking at the std normal tables for $z = (182 - 175)/(10/\sqrt{25}) = 3.50$, we get a probability of less than 0.001. That is, our confidence that the sample came from a lot with mean larger than 175 with likelihood > 99.9%, which is better than the 95% guarantee required.

3.1. Control Charts and Statistical Process Control

If you visit any production factory, you will often notice some graphs stuck onto the bulletin boards in various departments. These graphs look very similar -- even in companies that are making very different types of products. There are several types of

control charts, but it is sufficient to illustrate how they work using a single example: the $\bar{x}$-control chart (traditional statistics use $\bar{x}$ to indicate the mean.

Example *[source: Montgomery]*
Suppose a company is manufacturing piston rings. A critical dimension of the manufacturing process is the inside diameter. The machines are programmed to produce piston rings with a mean diameter = 74mm, and it is known that the standard deviation is 0.01 mm.

Figure 10. Piston rings (these are important components in automobile engines)

If all goes well, then all rings will be of acceptable quality; however, over time, some process(es) may go out of control, causing variations. Therefore we wish to keep track of the product quality, but we don't want to measure each ring.

The set up: Every hour, a random sample of 5 rings is measured, and the mean value of the inside diameter, $\bar{x}$, is recorded. For convenient visualization, the $\bar{x}$ data over time is plotted on a chart called a control chart.

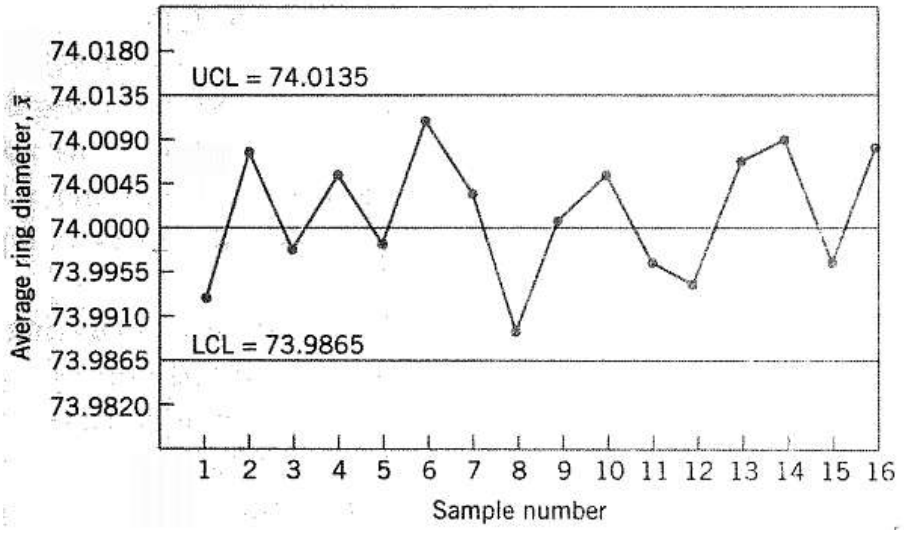

Figure 11. $\bar{x}$ control chart for piston-ring inside diameter [Source: Montgomery]

Notice that the chart has two extra lines, called the Upper Control Limit (UCL), and the Lower Control Limit (LCL) lines. These are critical to the use of the chart, so let's see how these are set.

σ = 0.01, and n = 5; therefore $\bar{x}$ is normally distributed with $\sigma_{\bar{x}}$ = 0.01/√5 = 0.0045. If the process is in control with mean diameter=74mm, then we expect 100(1 - α)% of the sample means to fall between the range of [74 - $Z_{\alpha/2}$(0.0045), 74 + $Z_{\alpha/2}$(0.0045)]. By convention, $Z_{\alpha/2}$ is set to a value of 3 (hence the name ***3-sigma control limits***). In our

example, we get the control limits = 74 ± 3(0.0045), so UCL = 74.0135, and LCL = 73.9865.

Interesting points:

1. If the sample size is larger, then the control limit lines move close together. [Why?]

2. Statistically, if a point on the $\bar{x}$ chart stays within the limits, it is equivalent to our confidence that we shall not make a Type I or Type II error regarding the hypothesis that our samples come from a population with $\mu = 74$. You can check the standard tables to verify that the probability of making a type I error = 0.0027, when we use $Z_{\alpha/2} = 3$.
3. Most control charts will also specify the sample size. If the sample size is larger, then the control chart will be able to identify smaller shifts in the process [Why?].

4. Often, in addition to the UCL and LCL, two more lines are drawn corresponding to $\pm 2\,\sigma_{\bar{x}}$ level. The use of these lines is discussed below.

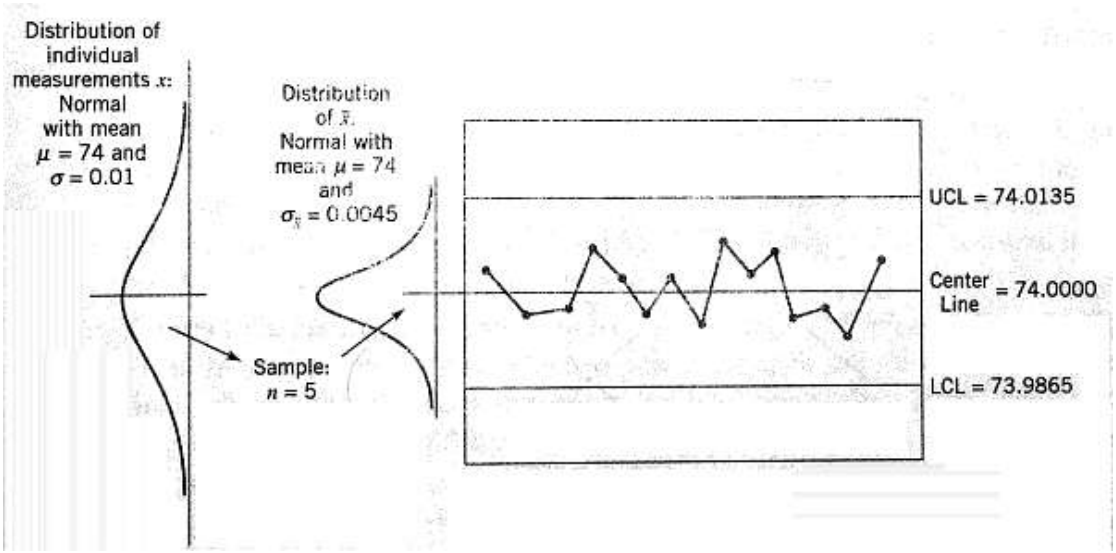

Fig. 12. Relation between sample variance, $\sigma_{\bar{x}}$ and the control limits [Source: Montgomery]

How to use Control Charts?

Usually, control charts are plotted regularly and posted on the factory floor. They provide a visible indication whenever a process goes, or will go 'out of control'. How does it work? Below are commonly used guidelines related to control charts usage. The process is out of control if any one or more of the criteria is met.

1. One or more points outside of the control limits. This pattern may indicate:
    - A special cause of variance from a material, equipment, method, or measurement system change.
    - Error in measurement of a part or parts.
    - Miscalculated or mis-plotted data points.
    - Miscalculated or mis-plotted control limits.
2. A run of eight points on one side of the center line. This pattern indicates a shift in the process output from changes in the equipment, methods, or materials or a shift in the measurement system.

3. Two of three consecutive points outside the 2-sigma warning limits but still inside the control limits. This may be the result of a large shift in the process in the equipment, methods, materials, or operator or a shift in the measurement system.
4. Four of five consecutive points beyond the 1-sigma limits.
5. An unusual or nonrandom pattern in the data.
    1. A trend of seven points in a row upward or downward. This may show
        - Gradual deterioration or wear in equipment.
        - Improvement or deterioration in technique.
    2. Cycling of data can indicate
        - Temperature or other recurring changes in the environment.
        - Differences between operators or operator techniques.
        - Regular rotation of machines.
        - Differences in measuring or testing devices that are being used in order.
6. Several points near a warning or control limit.

To conclude this section, I will mention that the first systematic use of this simple, but very effective technique was made by Dr. Walter A. Shewhart, and therefore sometimes control charts are also called Shewhart charts.

4.1. Robust design and the Design of Experiments

SPC and Sampling techniques basically center on the improvement of the product (by rejecting poor quality raw materials, for example), or process (by identifying and correcting root cause of variations.
A different approach was proposed formally by Genichi Taguchi, a quality control expert with the Japanese telecommunications company, NTT. A highly simplified version of his main idea is that quality can be controlled by improving the product design itself. We shall begin with an interesting example, and then give some basic insights about how it leads to a statistical approach.

Example: The INA Tile Company
The Ina tile company has a kiln (oven) which is used to bake large number of refractory tiles as shown in the figure below (tiles have many uses, for example, the entire building of HKUST is covered with ceramic tiles). Ideally, we would like all tiles to have the same dimensions after baking, but there is much variation in the sizes of the tiles. It is known that the variation is a result of the baking process.

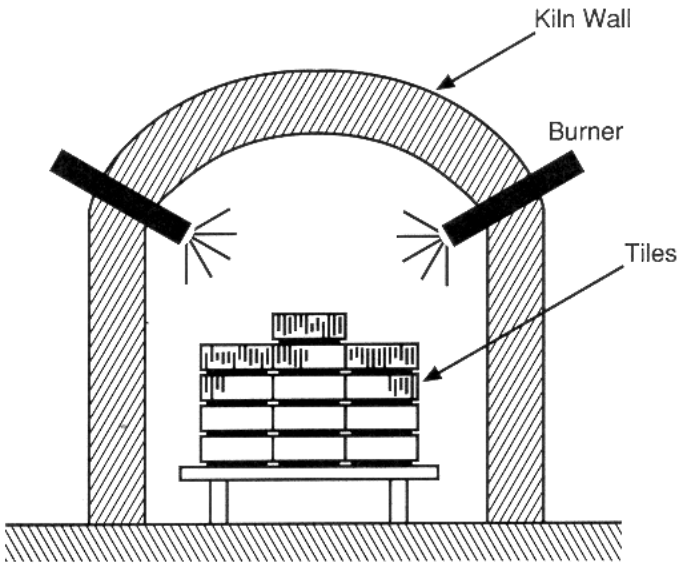

Figure 13. Tiles being baked in a kiln

We could examine each produced tile, and reject those that do not meet the dimensional specs -- but this is expensive and wasteful (due to large rejection rates).

The SPC approach is to identify the cause of the problem. In this case, a study of the problem indicated that the variations were due to differences in the temperature profiles experienced by the tiles on the outside with the tiles on the inside (Figure 14).

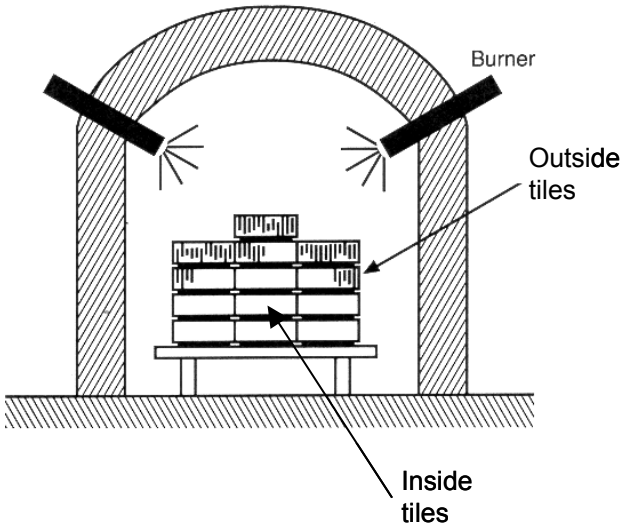

Figure 14. Different temperature profiles in different regions of kiln

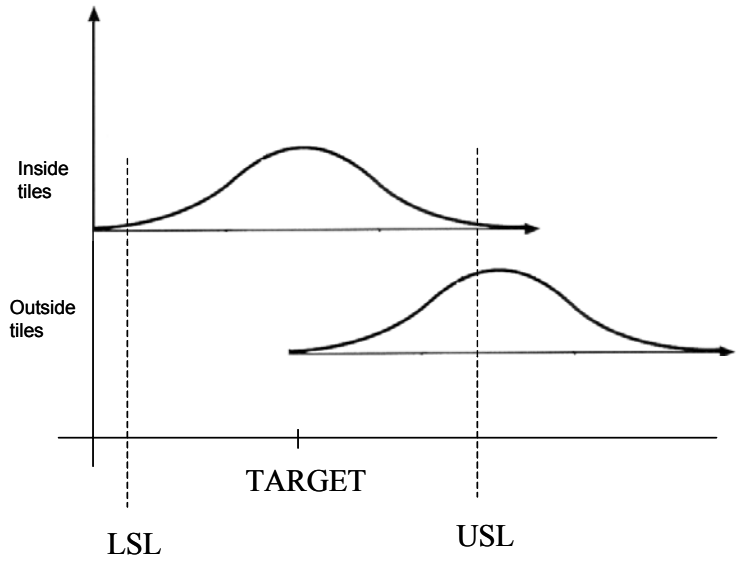

Figure 15. The variations in sizes of tiles from different regions in kiln

If there is some technique to reduce such differential profiles, for example, by the addition of special air circulation fans inside the kiln, perhaps one can achieve tighter dimensional control on the tiles, as shown in the figure below.

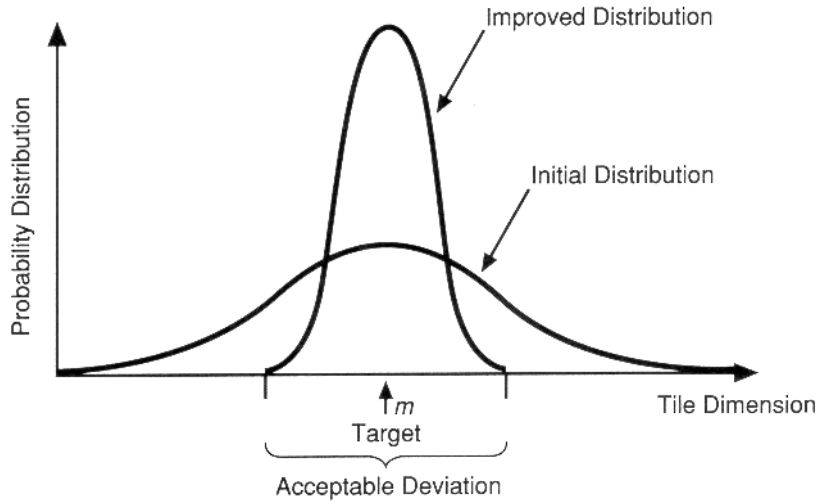

Figure 16.Reduce dimensional variations by eliminating cause ➔ reduce temp variations

However, such a re-design approach indicates expensive modifications to the factory infrastructure. Taguchi methods introduced a systematic approach that could help to eliminate the problem of product variations in the most economical fashion.

The method requires a careful identification of the response parameter we are measuring, and all control parameters as well as sources of noise. In case of Ina Tile Co., the following were identified:

Response: Tile dimension
Control Parameters (tile design):
Amount of Limestone

Fineness of additive
Amount of Agalmatolite
Type of Agalmatolite
Raw material Charging Quantity
Amount of Waste Return
Amount of Feldspar

Noise parameter was the temperature gradient.

By a careful set of experiments, the company was able to determine that by changing one of the control parameters, namely, 'Amount of Limestone', the sensitivity of the tile dimension to temperature gradients in the kiln was minimized. By increasing the limestone content, the new size variations that were obtained for tiles in the different regions of the kiln were as shown in the figure below.

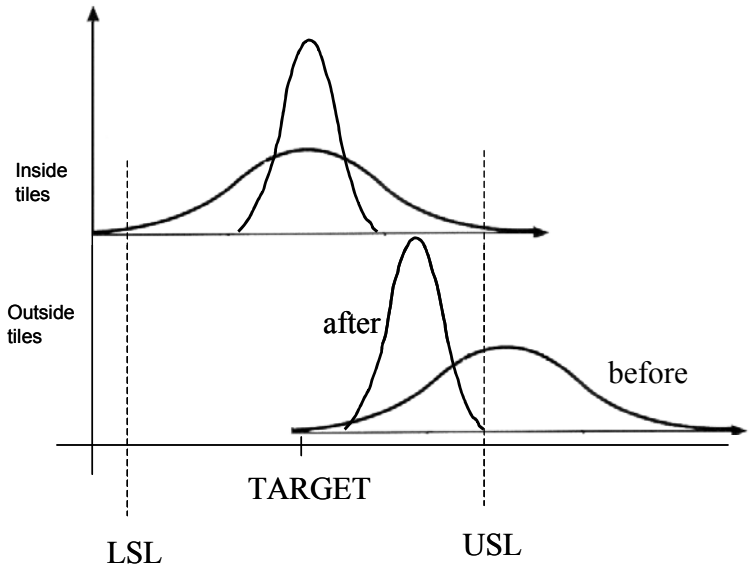

Figure 17. Tile size variations before and after changing limestone content.

In this case, this design change yielded tremendous benefit, not only by saving expensive kiln re-design costs, but also because limestone was the cheapest ingredient of the tiles !

Thus, the robust design method is defined as:
A method of designing a process or product aimed at reducing the variability (deviations from product target performance) by lowering sensitivity to noise.

4.2. The Design of Experiments

There are numerous examples of the application of Taguchi's method in modern industry. There is a well studied statistical method that can be used in each case. In the Ina Tile example, it is easy to see that the final design was great, but *how did the company know that the high-limestone design was the best one*? The answer is that that they did not know! What they did was to list all controllable parameters, and used different combinations of different input materials in each run. Then they tried out all the combinations before finding the best one. Here, 'all combinations' is clearly infinite; so we would like to try out only a subset of design combinations, but try to get more information out of the limited test we perform. In other words, we will experiment with a

subset of all possibilities. Therefore the approach we use is called the '***Design of Experiments***'. The following figure demonstrates the basic idea.

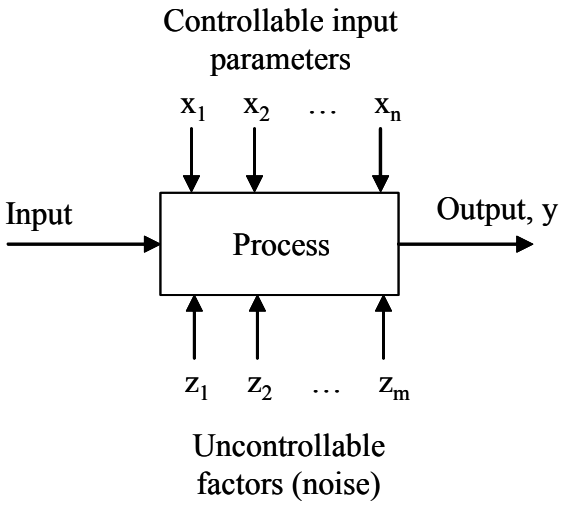

Figure 18. Parameters that affect the output from a Design of Experiments view

Typical **objectives** when we conduct such a **series of experiments** are:

(i) Determine which input variables have the most influence on the output;
(ii) Determine what value of x_i's will lead us closest to our desired value of y;
(iii) Determine where to set the most influential x_i's so as to reduce the variability of y;
(iv) Determine where to set the most influential x_i's such that the effects of the uncontrollable variables (z_i's) are minimized.

As you can see, the process of Design of experiments is related to the study of variations in inputs, and the resulting variability of the output y. The statistical technique used to do this is called **AN**alysis **O**f VAriance, or ANOVA. Unfortunately, we cannot go deeper into the study of ANOVA in this course. The main idea is that ANOVA allows us to analyse (namely, break into smaller components) the observed variance from the result of a series of experiments. The analysis allows us to draw inference about each component -- what is each component caused by. Thus, we can find out the relative effect/influence of variations in the inputs on the output. Clearly, this is useful in answering questions related to the stated objectives of the Design of Experiments.