CHAPTER 5

Central Limit Theorem

Not all populations are normally distributed and the reader should not assume them to be. However, for simplicity, the concepts in chapters 5 and 6 are presented using mostly normal or near normal populations.

*Non*normal populations are introduced in section 5.4, then demonstrated in chapters 7 and 8. ▼

5.1 Central Limit Theorem

One of the most remarkable theorems in statistics is called the **central limit theorem,** which is best explained through practical example.

Suppose a machine in a dress factory is set to cut pieces of silk material exactly to the length of 1000 mm. The pieces are then to be assembled into an outfit. Now, what are the chances the machine will cut every piece of silk material to precisely 1000 mm? Quite slim. Most of the cuts will be in the vicinity of 1000 mm, however many pieces will be shorter and many longer. Experience tells us that if a properly operating machine cuts millions and millions of pieces, the *histogram* representing the lengths of all these pieces of material may very well build into a shape closely resembling that of a normal distribution clustered around the average length of $\mu = 1000$ mm, and might look as follows:*

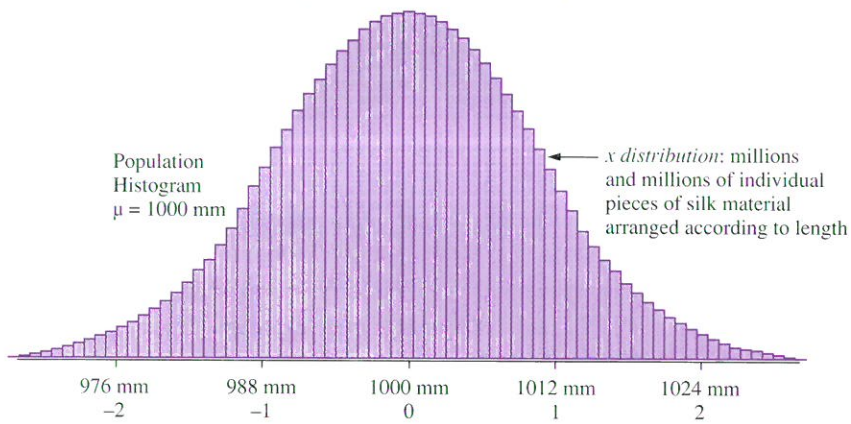

The standard deviation (σ) of a population such as this will vary from machine to machine. The speed of the cuts, even the length setting may affect the standard deviation, but let's say for our example we will ''invent'' a standard deviation of $\sigma = 12$ mm.

Now let us suppose, from these millions and millions of pieces we *randomly select* a sample of 36 pieces. Because our sample was *randomly* selected, we know from section 2.4 that:

$\bar{x} \approx \mu$ The sample average *is approximately equal to* the population average.

$s \approx \sigma$ The sample standard deviation *is approximately equal to* the population standard deviation.

*A single machine operating properly and uninterrupted will often produce goods whose measurement on a single characteristic, when recorded into a histogram, take on a shape strongly resembling that of a normal distribution.

So it should not be surprising that after we calculate $\bar{x}$, the average length of these 36 pieces in our sample, that this sample average ($\bar{x}$) might equal, say for instance, 998 mm, shown below.

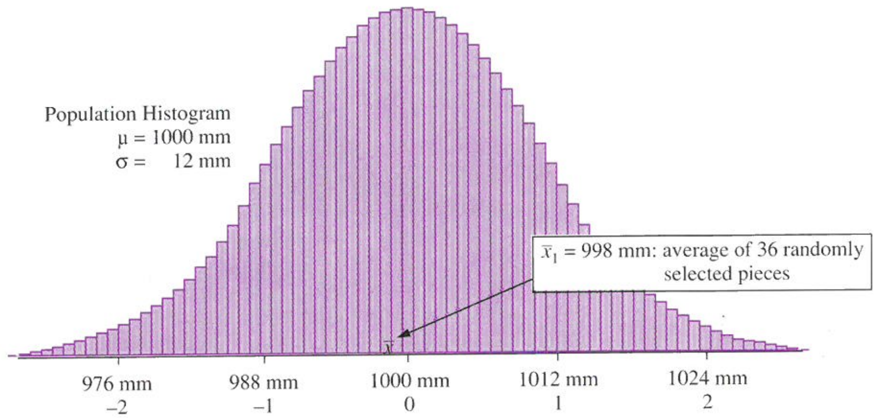

Since we know $\bar{x}$ should be *approximately equal to* μ, the question arises: would 998 mm be considered approximately equal to 1000 mm? In other words: how close must $\bar{x}$ be to be considered approximately equal to μ? For the answer to this question, we must take additional samples of 36 and actually calculate the values we get for $\bar{x}$. So, we randomly select a second sample of 36, then a third sample of 36, and even a fourth sample of 36. The new $\bar{x}$'s ($\bar{x}_2$, $\bar{x}_3$, and $\bar{x}_4$) are calculated and plotted along with $\bar{x}_1$ in the following diagram.

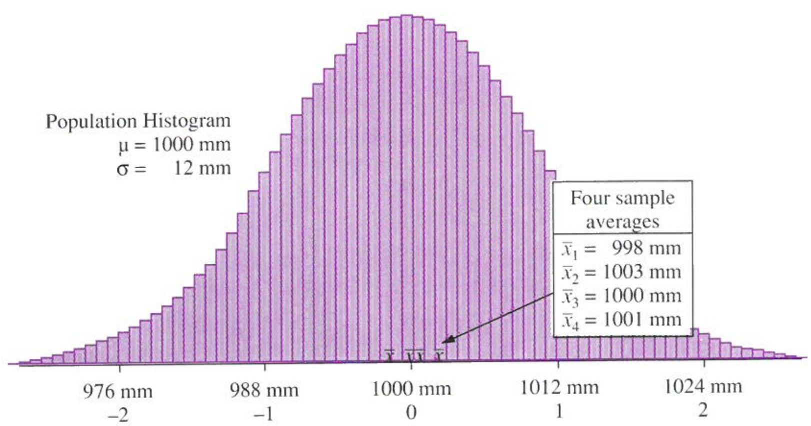

In the next sketch, we have added seven more sample averages ($\bar{x}$'s) to our plot. Note how the sample averages ($\bar{x}$'s) begin to ''pile up'' on the same readings, all in the vicinity of 1000 mm.

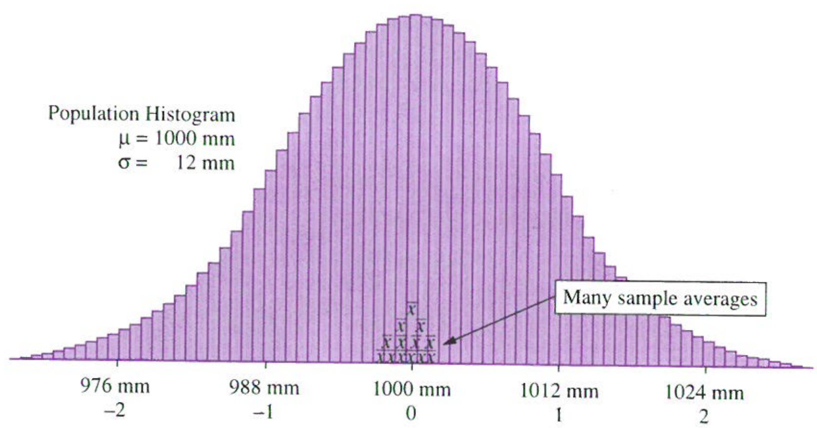

Now we run wild and randomly select thousands and thousands of samples, with each sample containing 36 pieces of material cut from the machine. For each sample of 36 pieces, we calculate the sample average such that, now, we have thousands and thousands of $\bar{x}$'s. Why on earth would anybody want to do this, you might ask? That's a difficult question to answer,[1] but somebody did and discovered something that, when put in combination with astute and sensible management, helped catapult numerous mid-sized businesses into gigantically successful worldwide empires. Two such empires are Proctor & Gamble and Intel. Management in these corporations use statistical techniques such as these in marketing research and technical analyses on a routine basis.

Okay, we now have thousands of $\bar{x}$'s. Now what? We group the results of all these thousands of *sample averages* and arrange them according to length into a **small histogram** (which we shall call the $\bar{x}$ distribution), which might look as follows:

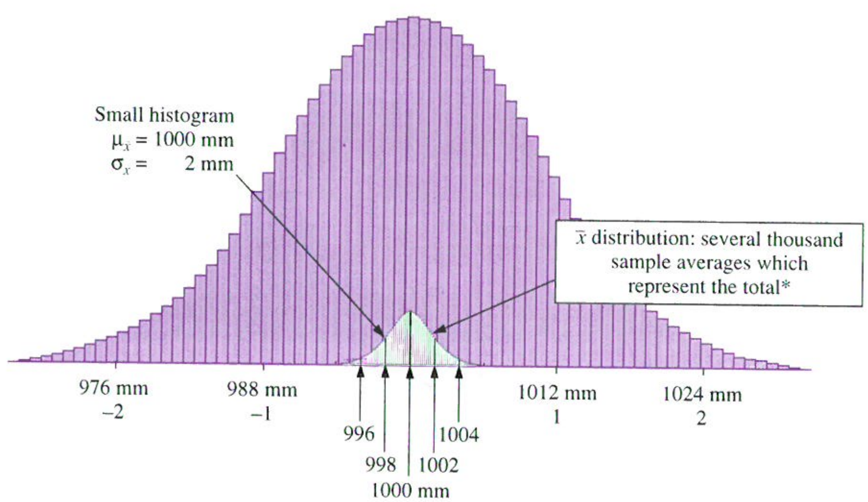

*Sampling distributions are based on the concept of sampling all possible different samples (of a fixed size) from a population. However, even small populations produce enormous numbers of different possible samples *(refer to endnote 2 for detailed discussion)*. However, usually after randomly selecting several hundred samples, the characteristics of a sampling distribution become quite clear. Sampling distributions in this section can be generated using Microsoft Excel (Tools, Data Analysis). For the given $\bar{x}$ distribution, fifteen thousand samples were randomly chosen, sample averages calculated and these values organized into a histogram represented above as the $\bar{x}$ distribution. The obtained values of $\mu_{\bar{x}}$ and $\sigma_{\bar{x}}$, the mean and standard deviation of one such sampling distribution, matched calculated values (formulas on next page) to approximately two decimal places.[2]

Perhaps now we can answer the question: how close does $\bar{x}$ have to be to be considered approximately equal to μ? We merely look at the results of the thousands and thousands of samples taken. Almost every sample average ($\bar{x}$) fell between 994 mm and 1006 mm. Are we saying it is impossible to get an $\bar{x}$ of, say, 1012 mm? No, not impossible, but highly improbable. In all the thousands of random samples actually taken, not one $\bar{x}$ even came near 1012 mm. In fact, most of the $\bar{x}$'s were clustered between 998 mm and 1002 mm with only a small few as far out as 994 mm or 1006 mm. Incredible!

And another thing. Did you notice the shape of the small histogram? Yes, a *normal distribution*. Which of course allows us to calculate areas that translate into probabilities. And this leads us to Theorem 5.1.

▼ Theorem 5.1: Properties of the $\bar{x}$ Distribution, Given Samples Drawn from a Normally Distributed Population

> If all possible different samples of a fixed sample size *n* are drawn from a normally distributed population with mean μ and standard deviation σ, then the distribution of $\bar{x}$'s will be normally distributed with mean μ and standard deviation $\sigma/\sqrt{n}$.

In other words, if the population is normally distributed, then

a. the $\bar{x}$ distribution will be normally distributed regardless of sample size, and
b. the mean and standard deviation of the $\bar{x}$ distribution are as follows:

$$\mu_{\bar{x}} = \mu$$
$$\sigma_{\bar{x}}{}^{*} = \sigma/\sqrt{n}$$

c. Furthermore, the z formula for the $\bar{x}$ distribution would be:

$$z = \frac{\bar{x} - \mu}{\sigma_{\bar{x}}}$$

Note that this theorem holds true for all sample sizes, for instance $n = 2$, $n = 3$ or any fixed sample size n provided, again, *the population is normally distributed.*

*$\sigma_{\bar{x}}$ is often referred to as the **standard error of the mean;** however, we will simply call it the standard deviation of the $\bar{x}$ distribution. Note: if the sample size n constitutes more than 5% of the population, then $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}\sqrt{\frac{N-n}{N-1}}$. The component $\sqrt{\frac{N-n}{N-1}}$ is referred to as the finite population correction factor.

In our "cutting machine" problem, since the population standard deviation (σ) is 12 mm and our sample size (n) is 36 pieces, $\sigma_{\bar{x}} = \sigma/\sqrt{n} = 12/\sqrt{36} = 2$ mm. Let's see how it works in the following example.

Example

A machine in a dress factory cuts pieces of silk material to an average length of $\mu = 1000$ mm with standard deviation $\sigma = 12$ mm. If one day we take a sample of size $n = 36$ pieces and calculate the average length of these thirty-six pieces and discover the sample average $\bar{x} = 995$ mm, (a) find the z score associated with $\bar{x} = 995$ mm, and (b) based on this z score, would you suspect the machine to be malfunctioning?

Solution

We do this problem as we would do any normal curve problem, only now we are working in the $\bar{x}$ distribution. First calculate $\sigma_{\bar{x}}$, the standard deviation of the $\bar{x}$ distribution, then construct a clear visual indicating values for at least ±2 standard deviations in that distribution.

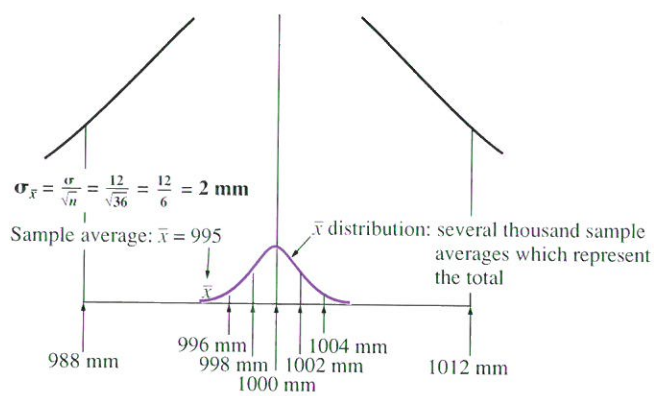

Next, we calculate the z score for a sample average $\bar{x} = 995$ mm.

$$z = \frac{\bar{x} - \mu}{\sigma_{\bar{x}}} = \frac{995 - 1000}{2} = -2.50$$

a. The z score for sample average $\bar{x} = 995$ mm is –2.50.
b. Certainly, we would be suspicious, since any sample average $\bar{x}$ more than two standard deviations from μ would be considered an unlikely occurrence if the machine was cutting properly. The question arises: did this unlikely occurrence occur or has μ, the average cut on the machine, changed? ■

Theorem 5.1 is a specific case of a more general theorem as follows.

▼ Central Limit Theorem

If all possible different samples of a fixed sample size n are drawn from a population of *any* shape with mean μ and standard deviation σ, then the $\bar{x}$ distribution will be normally distributed with mean μ and standard deviation $\sigma/\sqrt{n}$ provided the sample size n is greater than or equal to 30. That is, provided

$$n \geq 30$$

Thus, for any shaped population, provided $n \geq 30$

a. the $\bar{x}$ distribution will be normally distributed, and
b. the mean and standard deviation of the $\bar{x}$ distribution are as follows:

$$\mu_{\bar{x}} = \mu$$
$$\sigma_{\bar{x}} = \sigma/\sqrt{n}$$

Note that this theorem holds true for *any* shaped population provided $n \geq 30$. It is generally agreed among statisticians that a sample size of 30 or more is sufficient to assume a normal $\bar{x}$ distribution. We will use this as our guide.

5.2 Applying the Central Limit Theorem

The central limit theorem is one of the most remarkable achievements in statistics. It brought statistics out of the Dark Ages. Until its discovery and subsequent widespread application (starting about 100 years ago) we had only been able to estimate population characteristics with very large bodies of data—data that often took months or years to gather and sort and was often outdated before it was analyzed. Now we have at our disposal a precise mathematical way of estimating what is happening now. And because of this we can make better decisions. Let's see how it works in the following two examples.

Rounding Technique for This Textbook

As a general rule, work in three decimal places throughout the entire problem. Only round final answers to two decimal places. This rounding technique is important to ensure that everyone arrives reasonably close to the same answer. Use of the technique will grow increasingly more critical as we proceed through the material. Exception: z scores will *always* be presented in two decimal places, even in calculations.

Example

The National Institutes of Health agreed to supply active disease viruses, such as polio and AIDS, to research firms for the purpose of experimentation. A process is set up to automatically fill millions of small test tubes to an average of 9.00 milliliters of disease virus with standard deviation .35 milliliters. If we continually take random samples of 49 test tubes in each sample and calculate the average fill, $\bar{x}$, between what two values would you expect to find the middle 99% of all the sample averages ($\bar{x}$'s)?

Solution

Since we are concerned with the ''average'' fill of 49 test tubes and not with the contents of one test tube, we use the $\bar{x}$ distribution.

This is a typical ''working backward (given the area, find z)'' problem for the normal curve, only now we are dealing with the $\bar{x}$ distribution so we must first calculate $\sigma_{\bar{x}}$, the standard deviation of the $\bar{x}$ distribution, and list values for at least ±2 standard deviations.

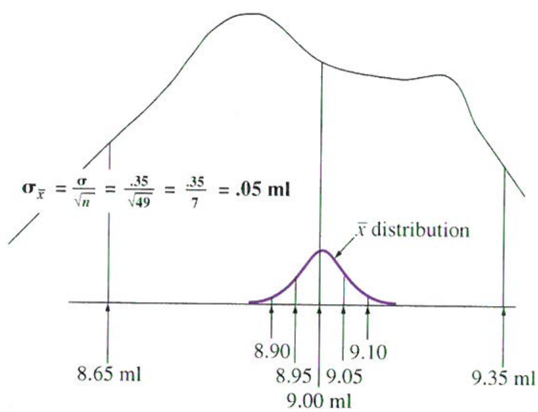

Since we know the area between the cutoffs is 99%, we merely look in the normal curve table for the corresponding z scores. *Remember:* the table reads ''half'' the normal curve, so we must look up an area of 49.5% ($\frac{1}{2}$ of 99%), which in decimal form is .4950.

According to the table, the corresponding z scores are -2.58 and $+2.58$ (note that .4950 fell precisely midway between two values in the table; in these cases, we round to the higher z score).

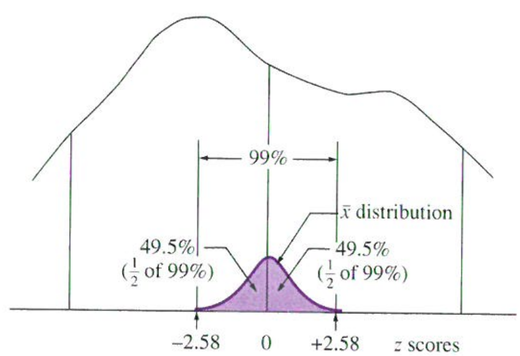

Substituting −2.58 and +2.58 into our formula, we solve for the cutoffs:

$$z = \frac{\bar{x} - \mu}{\sigma_{\bar{x}}} \qquad\qquad z = \frac{\bar{x} - \mu}{\sigma_{\bar{x}}}$$

$$-2.58 = \frac{\bar{x} - 9.00}{.05} \qquad\qquad +2.58 = \frac{\bar{x} - 9.00}{.05}$$

Solving for $\bar{x}$: Solving for $\bar{x}$:

$$\bar{x} = 8.87 \text{ ml} \qquad\qquad \bar{x} = 9.13 \text{ ml}$$

Graphically, the solution would appear as follows:

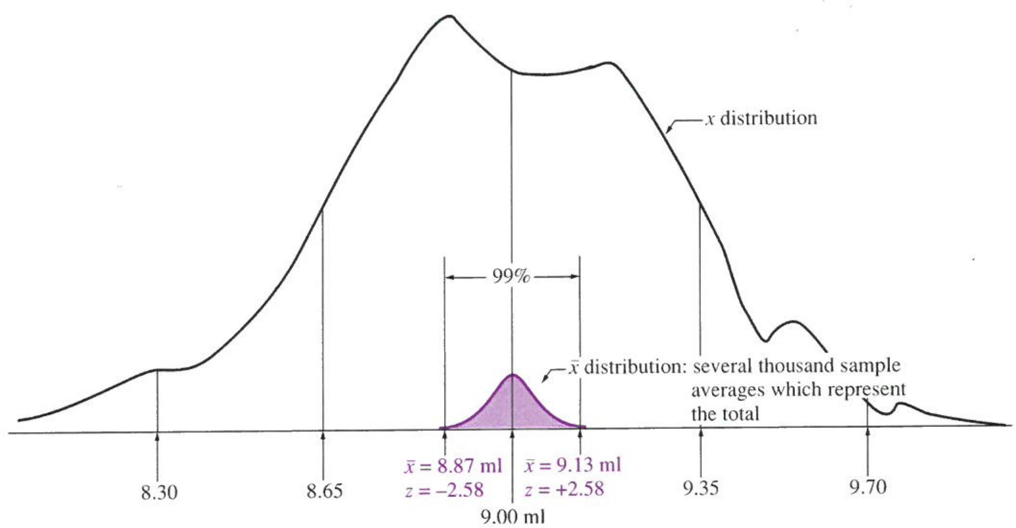

Answer

If we continually take random samples of 49 test tubes in each sample, and calculate the average fill ($\bar{x}$) in each of these samples, then 99% of all the sample averages ($\bar{x}$'s) would be expected to fall between $\bar{x}$ = 8.87 ml and $\bar{x}$ = 9.13 ml. ■

It is probably safe to say, if you randomly sample 49 test tubes from a properly functioning process, the average fill ($\bar{x}$) of this sample will fall between 8.87 ml and 9.13 ml. On any given day, if you obtain a sample average much *outside* this range, you might very well suspect the process is malfunctioning. Of course, we must keep in mind that 1% of the time (100% minus 99%), or approximately 1 out of every 100 times, a properly functioning process will produce a sample average ($\bar{x}$) outside this range, but since this is such a ''rare'' occurrence, it is probably wiser to check your filling operation for a malfunction.

Along with the value 95%, industry and research often use this value of 99% to establish a criteria for whether or not an operation may be malfunctioning. This is discussed at length in chapter 6.

Example

Brell Shampoo, an ''in-house'' brand, is marketed along with various other shampoos through a large national chain of convenience stores. In these stores, Brell's market share has remained relatively constant at $\mu = 24.00$ (meaning: on average 24.00% of the shampoo sold in these stores is Brell) with standard deviation 3.20. If we continually take random samples, each consisting of 64 stores, and calculate the average market share ($\bar{x}$) for each sample, what percentage of the sample averages ($\bar{x}$'s) would you expect to have a market share of *less than* 23.80?

Solution

Since we are concerned with the ''average'' market share in 64 stores and not the market share in an ''individual'' store, we use the $\bar{x}$ distribution. Remember: when using the $\bar{x}$ distribution we must first calculate $\sigma_{\bar{x}}$, the standard deviation of that distribution.

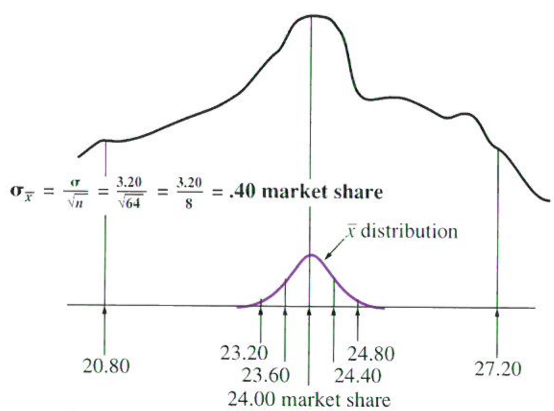

Next, we calculate the z score at the cutoff of $\bar{x} = 23.80$:

$$z = \frac{\bar{x} - \mu}{\sigma_{\bar{x}}} = \frac{23.80 - 24.00}{.40} = -.50$$

Since the area from $z = 0$ to $z = -.50$ is 19.15%, the area below $z = -.50$ must be 30.85% (50% minus 19.15%).

Graphically, the solution would appear as follows:

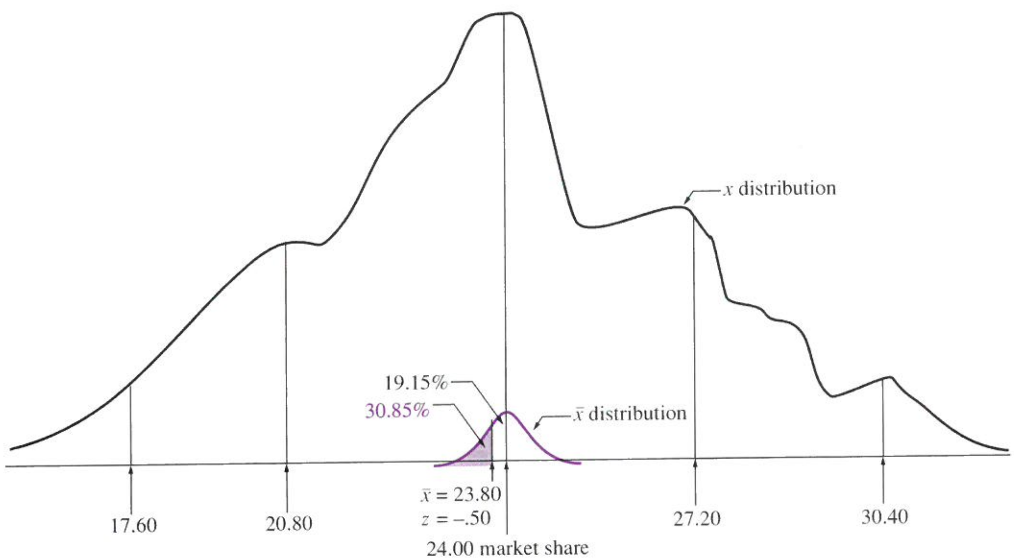

Answer

If we continually take random samples, each consisting of 64 stores, and calculate the average market share ($\bar{x}$) for each sample, then 30.85% of all the sample averages ($\bar{x}$'s) would be expected to have a market share of *less than* 23.80. ■

At this point a reminder may be needed: in the above problem concerning Brell Shampoo, we are not dealing with the market share in one outlet. We are dealing with the *average* market share in 64 randomly selected store outlets. The market share in one store can easily be, say, 19.00. However, the chances that the *average* market share in 64 randomly selected stores being 19.00 is nearly impossible.

Random Selection

Another reminder: the above mathematical procedure, or any mathematical procedure we discuss involving sampling, is based on the critically important process of random selection. Samples are chosen, in a way, similar to how lottery winners are selected. Each person (or store, in our case) has an equal chance of being selected on every pick. For instance, if you drive through Delaware and select 64 store outlets, this is *not* a random sample and these mathematical procedures cannot be used. For a random sample, you must have access to every store outlet in the country on each and every pick, say for instance, through a master list of all the stores. Then, each selection must be made as if ''blindfolded,'' such that each store has an equal chance of being picked (see chapter 1 for further discussion).

5.3 How n and σ Affect $\sigma_{\bar{x}}$

The standard deviation of the $\bar{x}$ distribution, $\sigma_{\bar{x}}$, is influenced by two factors. The first is σ, the population standard deviation, and the second is n, the sample size, according to the formula $\sigma_{\bar{x}} = \sigma/\sqrt{n}$.

How σ Affects $\sigma_{\bar{x}}$

The relationship between σ and $\sigma_{\bar{x}}$ is direct. If σ increases, $\sigma_{\bar{x}}$ increases. If σ decreases, $\sigma_{\bar{x}}$ decreases. Furthermore, σ and $\sigma_{\bar{x}}$ increase or decrease in the same multiple. In other words, if σ doubles, $\sigma_{\bar{x}}$ doubles. If σ triples, $\sigma_{\bar{x}}$ triples. If σ and $\sigma_{\bar{x}}$ are respectively 12 mm and 2 mm, and σ quadruples to 48 mm, then $\sigma_{\bar{x}}$ quadruples to 8 mm.

In practical terms, however, σ is a fixed item. One normally cannot manipulate the population standard deviation, σ, to influence $\sigma_{\bar{x}}$. But its relationship to $\sigma_{\bar{x}}$ is still important for an understanding of advanced work.

How n Affects $\sigma_{\bar{x}}$

The relationship between your sample size, n, and $\sigma_{\bar{x}}$ is more complex and of more concern since we can often control n. Note that $\sigma_{\bar{x}}$ varies inversely as the square root of n according to the formula

$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}.$$

First, this means, as n increases, $\sigma_{\bar{x}}$ decreases.

Second, this means, as n increases to 4 times its original value, $\sigma_{\bar{x}}$ decreases to $\frac{1}{\sqrt{4}}$ or $\frac{1}{2}$ its original value.

If n increases to 9 times its original value, then $\sigma_{\bar{x}}$ decreases to $\frac{1}{\sqrt{9}}$ or $\frac{1}{3}$ its original value.

If n increases to 16 times its original value, then $\sigma_{\bar{x}}$ decreases to $\frac{1}{\sqrt{16}}$ or $\frac{1}{4}$ its original value.

And so on.

To understand the impact of increasing sample size, let us again use the cutting machine problem in the following example.

Example

Using the cutting machine problem with $\mu = 1000$ mm and $\sigma = 12$ mm, compare the $\bar{x}$ distribution when you change your sample size from 36 pieces in each sample to four times this value (144 pieces in each sample) and then again to sixteen times this value (576 pieces in each sample).

Solution

If we continually take random samples of size $n = 36$, then the resulting $\bar{x}$ distribution would have a standard deviation $\sigma_{\bar{x}} = \frac{12}{\sqrt{36}} = 2$ mm.

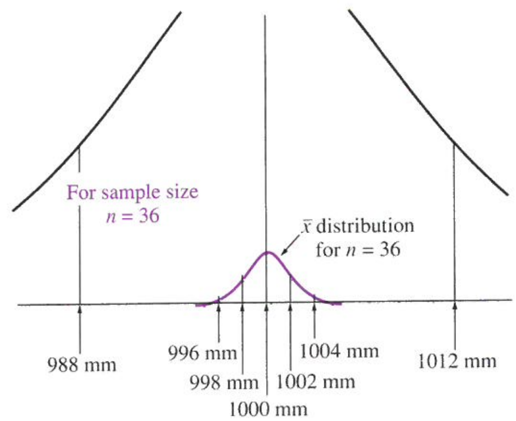

If we continually take random samples of size $n = 144$, then the resulting $\bar{x}$ distribution would have a standard deviation $\sigma_{\bar{x}} = \frac{12}{\sqrt{144}} = 1$ mm.

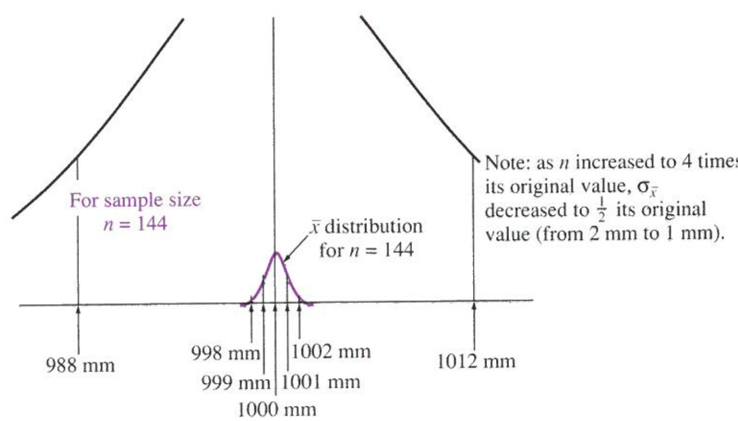

If we continually take random samples of size $n = 576$, then the resulting $\bar{x}$ distribution would have a standard deviation $\sigma_{\bar{x}} = \dfrac{12}{\sqrt{576}} = \dfrac{1}{2}$ mm.

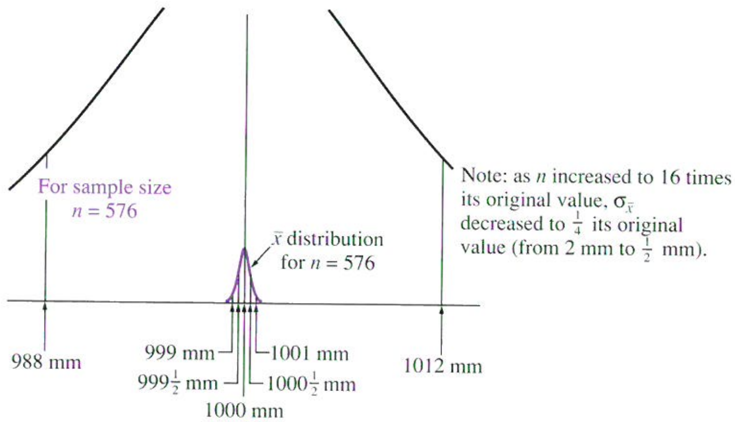

Note that a substantial increase in the sample size (n) is necessary to produce only a modest decrease in $\sigma_{\bar{x}}$. We had to increase our sample size to 16 times its original value (from 36 pieces to 576 pieces) to decrease $\sigma_{\bar{x}}$ to $\frac{1}{4}$ of its original value (from 2 mm to $\frac{1}{2}$ mm).

To fully understand the impact of sample size changes, let's see how it affects the location of the middle 95% of the $\bar{x}$'s in the next problem. ■

Example Using the cutting machine problem with $\mu = 1000$ mm and $\sigma = 12$ mm, calculate where the middle 95% of the $\bar{x}$'s would be expected to fall for the three situations in the last example, that is, for $n = 36$, for $n = 144$, and for $n = 576$.

Solution The completed solutions are as follows:

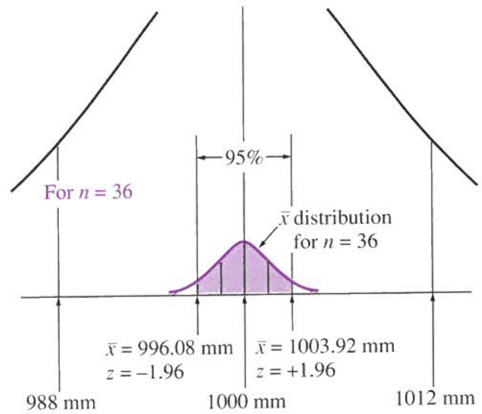
95%
For n = 36
x̄ distribution
for n = 36
x̄ = 996.08 mm
z = −1.96
x̄ = 1003.92 mm
z = +1.96
988 mm
1000 mm
1012 mm

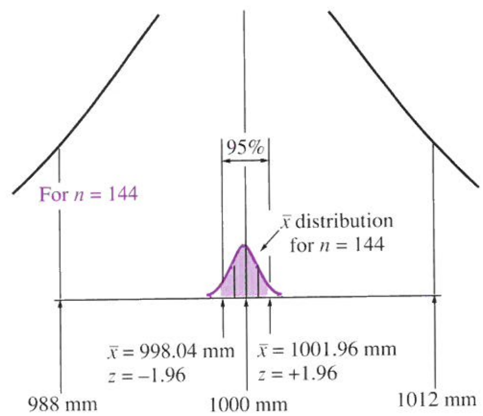
95%
For n = 144
x̄ distribution
for n = 144
x̄ = 998.04 mm
z = −1.96
x̄ = 1001.96 mm
z = +1.96
988 mm
1000 mm
1012 mm

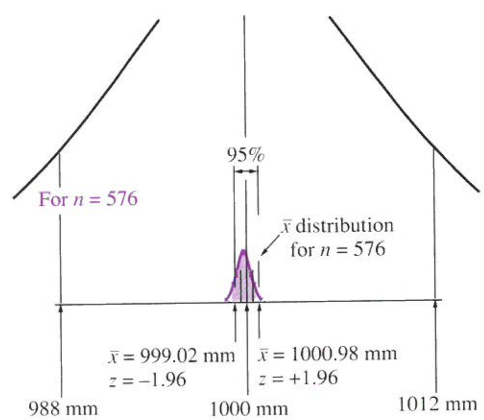
95%
For n = 576
x̄ distribution
for n = 576
x̄ = 999.02 mm
z = −1.96
x̄ = 1000.98 mm
z = +1.96
988 mm
1000 mm
1012 mm

Notice the extreme compression of the $\bar{x}$'s when you increase your sample size to 16 times its original value. At $n = 36$, note that 95% of the $\bar{x}$'s clustered between 996.08 mm and 1003.92 mm. However, at $n = 576$, the sample averages ($\bar{x}$'s) drew in closer to 1000 mm, with 95% of the $\bar{x}$'s now clustered between 999.02 mm and 1000.98 mm.

The importance of this compression of $\bar{x}$'s when the sample size is increased will become apparent when we discuss controlling statistical errors in chapter 6.

5.4 Central Limit Theorem Applied to Nonnormal Populations

The amazing thing about the central limit theorem is that it applies to any shaped population (normal or nonnormal), provided your sample size is 30 or more ($n \geq 30$). That is,

> The $\bar{x}$'s will distribute normally around μ for any shaped population, provided
>
> $$n \geq 30$$

Let's look at some examples.

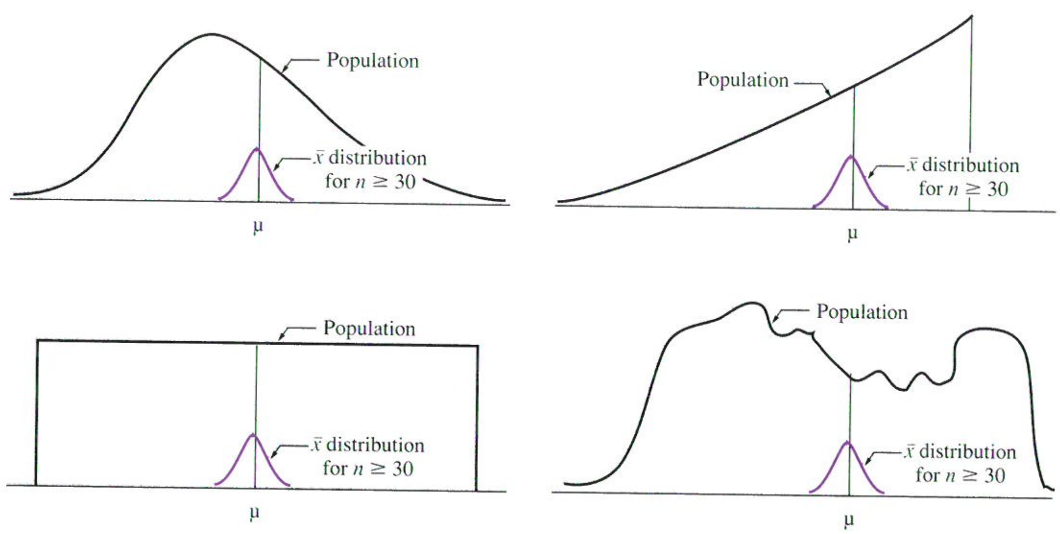

Note that in all cases, when the sample size was 30 or more ($n \geq 30$), the $\bar{x}$'s distributed normally around μ. And we can use the methods and techniques described in this chapter to predict where the $\bar{x}$'s will fall. Does this imply the $\bar{x}$'s will *not* distribute normally around μ if n is less than 30 ($n < 30$)? Yes and no. If your population is normal, the $\bar{x}$'s distribute normally around μ no matter what the size of n (even for $n = 5$ or $n = 2$). However, for many *non*normal populations, the $\bar{x}$'s do *not* assume the normal shape unless the sample size is at least 20 or 25 or in some extreme cases much more. A good rule of thumb is, assume the $\bar{x}$'s normally distributed for sample sizes of $n \geq 30$, unless the population has a highly unusual shape (say for instance, an extraordinarily skewed distribution), in which case, n may have to be more than 30 to ensure a normal $\bar{x}$ distribution. Further discussion on this is in sections 7.3 and 8.4.

Summary

Central limit theorem: For any shaped population, if you were to select all possible different samples of a fixed sample size greater than 30, the sample averages ($\bar{x}$'s) would build into the shape of a normal distribution, called an $\bar{x}$ distribution, such that:

1. The mean of the $\bar{x}$ distribution is μ, the mean of the population, and
2. The standard deviation of the $\bar{x}$ distribution, called sigma x-bar ($\sigma_{\bar{x}}$) is equal to the population standard deviation (σ) divided by the square root of your sample size (n). In other words,

$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$$

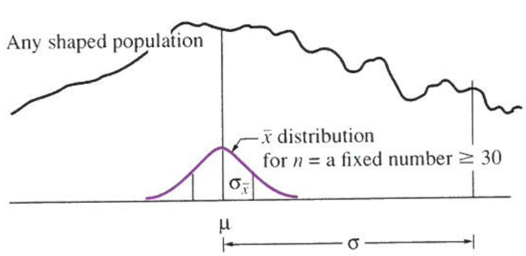

Theorem 5.1: Furthermore, if the population is normal, the $\bar{x}$ distribution will be normal for **any** sample size, even $n = 2$, $n = 3$ and so on.

Factors that affect $\sigma_{\bar{x}}$, the spread of the $\bar{x}$ distribution, are as follows.

σ affects the spread directly, that is, if σ *increases* to three times its value, $\sigma_{\bar{x}}$ will *increase* to three times its value.

n affects the spread inversely as its square root. If n *increases* to 16 times its value, then $\sigma_{\bar{x}}$ will *decrease* to $1/\sqrt{16}$ or 1/4 its value.

The amazing thing about the central limit theorem is that it applies to almost any shaped population provided the sample size is equal to or greater than 30. In other words, the $\bar{x}$'s will distribute normally around μ for almost any shaped population provided.

$$n \geq 30$$

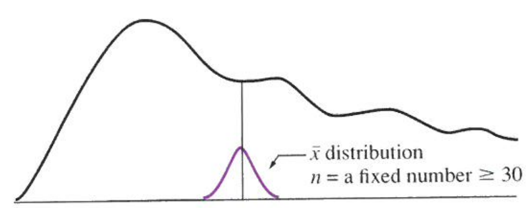

Only in populations with quite unusual shapes (such as one with a highly extended skew) may we have to sample sometimes more than 30 to be assured of a normally distributed $\bar{x}$ distribution.

5.1 - 5.8 (2008)

Exercises

Note that full answers for exercises 1–5 and abbreviated answers for odd-numbered exercises thereafter are provided in the Answer Key.

5.1 A machine cuts pieces of silk material to an average length of 1000 mm with standard deviation 12 mm. Between what two lengths would we expect to find the middle 95% of all the sample averages ($\bar{x}$'s)?

Assume sample size:

a. $n = 36$
b. $n = 144$
c. $n = 576$

5.2 A machine cuts pieces of silk material to an average length of 1000 mm with standard deviation 12 mm. Between what two lengths would we expect to find the middle 99% of all the sample averages ($\bar{x}$'s)?

Assume sample size:

a. $n = 36$
b. $n = 144$
c. $n = 576$

5.3 A national institute supplies active disease viruses for medical research. A process is set to fill small test tubes at an average of 9.00 ml with standard deviation of .35 ml. If we continually take random samples of 49 test tubes in each sample, what percentage of the $\bar{x}$'s would you expect to fall *below* 8.92 ml or *above* 9.08 ml?

5.4 Brell Shampoo, an "in-house" brand, is marketed through a large national chain of convenience stores. This chain also carries other national brands of shampoo. Brell's "in-house" market share is $\mu = 24.00$ (meaning: on average 24.00% of the shampoo sold in these stores is Brell) with standard deviation 3.20. If we continually take random samples, each sample consisting of 64 stores, and calculate the average market share ($\bar{x}$) in each sample, below what market share would we find the *lowest 1%* of all the sample averages ($\bar{x}$'s)?

5.5 The credit manager of a large sports shop made a statement at an important board meeting that the average age of their customers is 32 years old with standard deviation 4 years.

a. If the credit manager is correct and we were to continually take random samples of 100 customers each, what percentage of the $\bar{x}$'s would you expect to be between $\bar{x} = 31.5$ years old and $\bar{x} = 32.5$ years old?

b. If the credit manager is correct and we were to take *one* random sample of 100 customers, what is the probability the average of this one sample ($\bar{x}$) would be between 31.5 years old and 32.5 years old?

c. Between what two ages would you expect to find the middle 95% of all the sample averages ($\bar{x}$'s)?

5.6 Gaunt Health Farms, based on a survey of records of all visitors for the last five years, claim an average weight loss of 12.0 lb with standard deviation of 2.4 lb.

a. If you took a random sample of 36 from these records, what is the probability the sample average ($\bar{x}$) will be less than 11.0 lb?

b. If you took a random sample of 36 and your sample average, $\bar{x}$, was indeed less than 11.0 lb, would you be suspicious of their claim that the average visitor weight loss is 12.0 lb?

5.7 Bad-debt accounts are a serious source of profit drain for all businesses, but especially for the fashion industry, which deals with the risky whims of the public. Ralph Weetz Co., a distributor of women's blouses to small boutiques, was one such company. A computer tally of all bad-debt accounts of the past few decades reveals an accumulation of thousands of bad-debt accounts, with the average amount owed of $\mu = \$550.00$ and standard deviation $\sigma = \$75.90$.

a. Assuming a normal distribution, if we randomly selected one bad-debt account, what is the probability this one bad-debt account is between \$538.00 and \$562.00?

b. If we took a random sample of 40 bad-debt accounts and calculated the average amount owed ($\bar{x}$) in this sample, what is the probability this sample average ($\bar{x}$) will be between \$538.00 and \$562.00?

c. Assuming $n = 34$, with what probability can we assert a sample average ($\bar{x}$) will fall within \$20.00 of $\mu = \$550.00$?

5.8 A nationwide marketing study concluded the average age of horror film moviegoers is 17.4 years old with standard deviation 2.7 years.

a. Assuming a normal distribution, what percentage of horror film movie goers nationwide would you expect to be over 18.0 years old?

b. If we continually take random samples of 81 horror film movie goers nationwide and calculate the sample average ($\bar{x}$) for each sample, what percentage of the sample averages ($\bar{x}$'s) would you expect to be over 18.0 years?

c. If you took a random sample of 81 horror film movie goers, what is the probability the sample average ($\bar{x}$) would be over 18.0 years?

5.9 In a certain year, the nationwide SAT verbal score averaged $\mu = 430$ with standard deviation $\sigma = 96$. Answer the following assuming SAT scores are continuous over the scale 200 to 800.

a. Assuming a normal distribution, between what two values would you expect to find the middle 90% of SAT verbal scores?

b. If we continually take random samples of 144 students, and calculate the average SAT verbal score ($\bar{x}$) in each sample, between what two values would you expect to find the middle 90% of the sample averages ($\bar{x}$'s)?

c. If we continually take random samples of 42 students, and calculate the average SAT verbal score ($\bar{x}$) in each sample, between what two values would you expect to find the middle 90% of the sample averages ($\bar{x}$'s)?

5.10 In a certain year, the nationwide SAT mathematics score averaged $\mu = 470$ with standard deviation $\sigma = 96$. Answer the following assuming SAT scores are continuous over the scale 200 to 800.

a. Assuming a normal distribution, between what two values would you expect to find the middle 98% of SAT mathematics scores?

b. If we continually take random samples of 256 students, and calculate the average SAT mathematics score ($\bar{x}$) in each sample, between what two values would you expect to find the middle 98% of the sample averages ($\bar{x}$'s)?

c. If we continually take random samples of 30 students, and calculate the average SAT mathematics score ($\bar{x}$) in each sample, between what two values would you expect to find the middle 98% of the sample averages ($\bar{x}$'s)?

5.11 Medical doctors in Kansas City are known to work on average $\mu = 54.7$ hours per week with standard deviation $\sigma = 6.8$ hours.

a. Assuming a normal distribution, what is the probability a doctor will work less than 53.0 hr/wk?

b. What is the probability a random sample of 55 doctors will yield an average, $\bar{x}$, of less than 53.0 hr/wk?

c. With what probability can we assert a random sample average ($\bar{x}$) will be between 53.0 and 56.0 hr/wk, based on $n = 55$?

5.12 A survey indicated the average yearly salary of entry-level women managers in St. Paul to be $\mu = \$56{,}700$ with standard deviation $\sigma = \$7{,}200$.

a. Assuming a normal distribution, what is the probability a woman manager's entry-level salary will exceed \$58,000?

b. What is the probability a random sample of 50 women managers will yield an average entry-level salary ($\bar{x}$) exceeding \$58,000?

c. Assuming $n = 50$, with what probability can we assert a sample average ($\bar{x}$) will fall between \$55,000 and \$58,000?

d. Assuming $n = 42$, with what probability can we assert a sample average ($\bar{x}$) will fall within \$1500 of $\mu = \$56{,}700$?

Endnotes

1. Actually, the earliest experiments were based on Carl Gauss's work on measurement error in the field of astronomy (1809, 1816, 1823; Encke, 1832, 1834). Laplace formulated the theoretical underpinnings (1810–1812) and, later, experiments in biological measurement, agriculture and machine output confirmed these early findings.

2. Sampling distributions are based on the concept of sampling all possible different samples (of a fixed sample size) from a given population. Theory assumes that each sample is drawn without replacement and the pieces of material returned to the population and the process repeated until all possible different samples are obtained. However, even small populations using small sample sizes yield enormous numbers of different possible samples. For instance, a population with only $N = 52$ values using sample sizes of $n = 5$ produce approximately 2.6 million different possible samples. Ask anyone who plays the card game 5-card poker. From a deck of 52 playing cards, a person is dealt a random selection of 5 cards. There are C(52,5) = 2,598,960 different possible hands. In statistics, four aces and the two of clubs is considered a different hand (or a different sample of $n = 5$) than four aces and the three of clubs.

In the "cutting machine" example presented, suppose the population had 1,000,000 values (that is, the lengths of one million pieces of material) and when plotted looked as shown in Example A below.

If we continually draw samples of size $n = 36$ pieces, then C(1000000,36) = 10^{174} different samples are possible. To give you a feel for the magnitude of the number 10^{174}, it is thought the entire universe contains less than 10^{100} atoms. In fact, if we plotted 10^{174} sample means on the same frequency scale as the population, it might look as shown in Example B below.

Note that even the tails of the population produce an enormous number of different possible samples. For instance, the 50 smallest population values indicated by the box in the lower left corner of the sketch produce C(50,36) = 10^{12} different possible samples. Thus we would have to plot 10^{12} sample averages ($\bar{x}$'s) at the very tail of the population. Keep in mind, the entire population histogram below contains only 10^6 values which represent the lengths of the pieces of material.

To present the entire 10^{174} sample means on the same frequency scale as the population would produce an $\bar{x}$ distribution rising to a height out of our known universe.

In the representation we presented in section 5.1, we used fifteen thousand sample averages generated by computer simulation to represent this total. As we would expect, most $\bar{x}$'s clustered rather close to μ.

Often in the literature, sampling and population distributions are presented as probability distributions (or probability density functions), that is, each distribution is said to contain an area (a probability density) of 1.00, no matter how many values are contained in the distribution. Since both the population and sampling distribution contain the same area (1.00), the use of overlapping sketches can be awkward and often the two distributions are separated. This allows for a number of advantages, however, it forgoes the opportunity of directly comparing the standard deviations of the two distributions on the same scale. For this and other instructional purposes, we represent all sampling distributions as modified frequency sketches: several thousand sample values which represent the total.

Example A

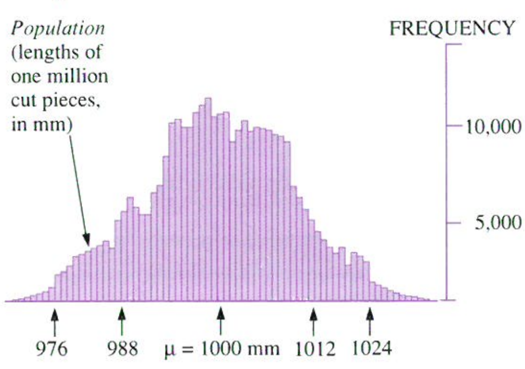

Example B