

*NOTES FOR SUMMER STATISTICS INSTITUTE COURSE*

**COMMON MISTAKES IN STATISTICS –  
SPOTTING THEM AND AVOIDING THEM**

**Part II: Inference**

MAY 24 – 27, 2010

Instructor: Martha K. Smith

**CONTENTS OF PART 2**

Simple Random Samples, Part 2: More Precise Definition of Simple Random Sample	3
Why Random Sampling is Important	9
Overview of Frequentist Hypothesis Testing	10
Overview of Frequentist Confidence Intervals	16
Frequentist Hypothesis Tests and p-Values	24
p-values	27
Misinterpretations and Misuses of p-values	30
Type I and Type II Errors and Significance Levels	33
Connection between Type I Error and Significance Level	34
Pros and Cons of Setting a Significance Level	35
Type II Error	36
Considering Both Types of Error Together	37
Deciding what Significance Level to Use	41
Power of a Statistical Procedure	44
Elaboration	46
Power and Type II Error	48
Common Mistakes Involving Power	52

## SIMPLE RANDOM SAMPLES, PART 2: More Precise Definition of Simple Random Sample

In practice in applying statistical techniques, we are interested in *random variables* defined on the population under study. To illustrate, recall the examples in Simple Random Sample, Part 1:

1. In a medical study, the population might be all adults over age 50 who have high blood pressure.
2. In another study, the population might be all hospitals in the U.S. that perform heart bypass surgery.
3. If we're studying whether a certain die is fair or weighted, the population is all possible tosses of the die.

In these examples, we might be interested in the following random variables:

Example 1: The difference in blood pressure with and without taking a certain drug.

Example 2: The number of heart bypass surgeries performed in a particular year, or the number of such surgeries that are successful, or the number in which the patient has complications of surgery, etc.

Example 3: The number that comes up on the die.

If we take a *sample of units from the population*, we have a corresponding *sample of values of the random variable*.

For example, if the random variable (let's call it  $Y$ ) is “difference in blood pressure with and without taking the drug”, then the sample will consist of values we can call  $y_1, y_2, \dots, y_n$ , where

- $n$  = number of people in the sample from the population of patients
- The people in the sample are listed as person 1, person 2, etc.
- $y_1$  = the difference in blood pressures (that is, the value of  $Y$ ) for the first person in the sample,
- $y_2$  = the difference in blood pressures (that is, the value of  $Y$ ) for the second person in the sample
- etc.

*We can look at this another way*, in terms of  $n$  random variables  $Y_1, Y_2, \dots, Y_n$ , described as follows:

- The random process for  $Y_1$  is “pick the first person in the sample”; the value of  $Y_1$  is the value of  $Y$  for that person – i.e.,  $y_1$ .
- The random process for  $Y_2$  is “pick the second person in the sample”; the value of  $Y_2$  is the value of  $Y$  for that person – i.e.,  $y_2$ .
- etc.

The difference between using the small  $y$ 's and the large  $Y$ 's is that *when we use the small  $y$ 's we are thinking of a fixed sample of size  $n$  from the population, but when we use the large  $Y$ 's, we are thinking of letting the sample vary (but always with size  $n$ ).*

*Note:* The  $Y_i$ 's are sometimes called *identically distributed*, because they have the same probability distribution.

**Precise definition of simple random sample of a random variable:**

"The sample  $y_1, y_2, \dots, y_n$  is a simple random sample" means that the associated random variables  $Y_1, Y_2, \dots, Y_n$  are independent.

Intuitively speaking, "independent" means that *the values of any subset of the random variables  $Y_1, Y_2, \dots, Y_n$  do not influence the values of the other random variables in the list.*

*Recall:* We defined a random sample as one that is chosen by a random process.

- Where is the random process in the precise definition?

*Note:* To emphasize that the  $Y_i$ 's all have the same distribution, the precise definition is sometimes stated as, " $Y_1, Y_2, \dots, Y_n$  are independent, identically distributed," sometimes abbreviated as iid.

**Connection with the initial definition of simple random sample**

*Recall* the preliminary definition (from Moore and McCabe, *Introduction to the Practice of Statistics*) given in Simple Random Samples, Part 1:

*"A simple random sample (SRS) of size  $n$  consists of  $n$  individuals from the population chosen in such a way that every set of  $n$  individuals has an equal chance to be the sample actually selected."*

*Recall Example 3 above:* We are tossing a die; the number that comes up on the die is our random variable  $Y$ .

- In terms of the preliminary definition, the population is all possible tosses of the die, and a simple random sample is  $n$  different tosses.
- The different tosses of the die are independent events (i.e., what happens in some tosses has no influence on the other tosses), which means that in the precise definition above, the random variables  $Y_1, Y_2, \dots, Y_n$  are indeed independent: The numbers that come up in some tosses in no way influence the numbers that come up in other tosses.

Compare this with example 2: The population is all hospitals in the U.S. that perform heart bypass surgery.

- Using the preliminary definition of simple random sample of size  $n$ , we end up with  $n$  *distinct* hospitals.
- This means that when we have chosen the first hospital in our simple random sample, *we cannot choose it again to be in our simple random sample.*
- Thus the events "Choose the first hospital in the sample; choose the second hospital in the sample; ... ," are *not* independent events: The choice of first hospital restricts the choice of the second and subsequent hospitals in the sample.
- If we now consider the random variable  $Y$  = the number of heart bypass surgeries performed in 2008, then it follows that the random variables  $Y_1, Y_2, \dots, Y_n$  are *not* independent.

**The Bottom Line:** *In many cases, the preliminary definition does not coincide with the more precise definition.*

More specifically, the preliminary definition allows sampling *without* replacement, whereas the more precise definition requires sampling *with* replacement.

**The Bad News:** *The precise definition is the one that is used in the mathematical theorems that justify many of the procedures of statistical inference. (More detail later.)*

**The Good News:**

- *If the population is large enough, the preliminary definition is close enough for all practical purposes.*
- *In many cases where the population is not "large enough," there are alternate theorems giving rise to alternate procedures using a "finite population correction factor" that will work*

Unfortunately, the question, "How large is large enough?" does not have a simple answer.

*The upshot:*

- 1) Using a "large population" procedure with a "small population" is a **common mistake**.
- 2) One more difficulty in selecting an appropriate sample, which leads to one more source of uncertainty.

### WHY RANDOM SAMPLING IS IMPORTANT

**Myth:** "A random sample will be representative of the population".

- This is false: a random sample might, by chance, turn out to be anything but representative.
- For example, it's possible (though unlikely) that if you toss a fair die ten times, all the tosses will come up six.
- If you find a book or web page that gives this reason, apply some healthy skepticism to other things it claims.

**A slightly better explanation that is partly true but partly urban legend:** "Random sampling prevents bias by giving all individuals an equal chance to be chosen."

- The element of truth: Random sampling will eliminate *systematic* bias.
- A practical rationale: This statement is often the best plausible explanation that is acceptable to someone with little mathematical background.
- However, this statement could easily be misinterpreted as the myth above.
- An additional, very important, reason why random sampling is important, at least in frequentist statistical procedures, which are those most often taught (especially in introductory classes) and used:

**The real reason:** *The mathematical theorems that justify most frequentist statistical procedures apply only to random samples.*

The next section elaborates.

### OVERVIEW OF FREQUENTIST HYPOTHESIS TESTING

Most commonly-used, frequentist hypothesis tests involve the following elements:

1. *Model assumptions*
2. *Null and alternative hypotheses*
3. *A test statistic.*

- This is something calculated by a rule from a sample.
- It needs to have the property that *extreme values of the test statistic cast doubt on the null hypothesis.*

4. *A mathematical theorem* saying, "If the model assumptions and the null hypothesis are both true, then the sampling distribution of the test statistic has this particular form."

*Note:*

- The *sampling distribution* is the probability distribution of the test statistic, when considering *all* possible suitably random samples of the same size. (More later.)
- The exact details of these four elements will depend on the particular hypothesis test.

**Example:** In the case of the large sample z-test for the mean with two-sided alternative, the elements are:

1. *Model assumptions:* We are working with simple random samples of a random variable  $Y$  that has a normal distribution.

2. *Null hypothesis:* "The mean of the random variable in question is a certain value  $\mu_0$ ."

*Alternative hypothesis:* "The mean of the random variable  $Y$  is not  $\mu_0$ ." (This is called the *two-sided alternative*.)

3. *Test statistic:*  $\bar{y}$  (the *sample mean* of a simple random sample of size  $n$  from the random variable  $Y$ ).

Before discussing item 4 (the mathematical theorem), we first need to do two things:

#### I. Note on terminology:

- The **mean of the random variable**  $Y$  is also called the *expected value* or the *expectation* of  $Y$ .
  - It's denoted  $E(Y)$ .
  - It's also called the *population mean*, often denoted as  $\mu$ .
  - It is what we do *not* know in this example.
- A **sample mean** is typically denoted  $\bar{y}$  (read "y-bar").
  - It is calculated from a sample  $y_1, y_2, \dots, y_n$  of values of  $Y$  by the familiar formula  $\bar{y} = (y_1 + y_2 + \dots + y_n)/n$ .
- The sample mean  $\bar{y}$  is an *estimate* of the population mean  $\mu$ , but they are usually *not* the same.
  - Confusing them is a **common mistake**.
- Note that I have written, "the population mean" but "a sample mean".
  - A *sample* mean depends on the sample chosen.
  - Since there are many possible samples, there are many possible sample means.
  - However, there is only one *population* mean associated with the random variable  $Y$ .

II. We now need to step back and consider *all* possible simple random samples of  $Y$  of size  $n$ .

- For each simple random sample of  $Y$  of size  $n$ , we get a value of  $\bar{y}$ .
- We thus have a *new* random variable  $\bar{Y}_n$ :
  - The associated random process is "pick a simple random sample of size  $n$ "
  - The value of  $\bar{Y}_n$  is the sample mean  $\bar{y}$  for this sample
- Note that
  - $\bar{Y}_n$  stands for the new random variable
  - $\bar{y}$  stands for the value of  $\bar{Y}_n$ , for a particular sample of size  $n$ .
- The distribution of  $\bar{Y}_n$  is called the *sampling distribution of  $\bar{Y}_n$*  (or the *sampling distribution of the mean*).

Now we can state the theorem:


4. The *theorem* states: *If* the model assumptions are all true (i.e., if  $Y$  is normal and all samples considered are simple random samples), and *if in addition* the mean of  $Y$  is indeed  $\mu_0$  (i.e., if the null hypothesis is true), then
- The sampling distribution of  $\bar{Y}_n$  is normal
  - The sampling distribution of  $\bar{Y}_n$  has mean  $\mu_0$
  - The sampling distribution of  $\bar{Y}_n$  has standard deviation  $\sigma/\sqrt{n}$ , where  $\sigma$  is the standard deviation of the original random variable  $Y$ .

**More Terminology:**  $\sigma$  is called the *population standard deviation* of X; it is *not* the same as the *sample standard deviation*  $s$ , although  $s$  is an estimate of  $\sigma$ .

The following chart and picture summarize the theorem and related information:

	<b>Random variable Y (population distribution)</b>	Related quantity calculated from a sample $y_1, y_2, \dots, y_n$	<b>Random variable <math>\bar{Y}_n</math> (sampling distribution)</b>
<b>Mean</b>	<i>Population mean</i> $\mu$ ( $\mu = \mu_0$ <u>if</u> null hypothesis true)	<i>Sample mean</i> $\bar{y} = (y_1 + y_2 + \dots + y_n)/n$ $\bar{y}$ is an <u>estimate</u> of the population mean $\mu$	<i>Sampling distribution mean</i> $\mu$ ( $\mu = \mu_0$ <u>if</u> null hypothesis true)
<b>Standard deviation</b>	<i>Population standard deviation</i> $\sigma$	<i>Sample standard deviation</i> $s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\bar{y} - y_i)^2}$ $s$ is an <u>estimate</u> of the population standard deviation $\sigma$	<i>Sampling distribution standard deviation</i> $\sigma/\sqrt{n}$

*Comment:* The model assumption that the sample is a simple random sample (in particular, that the  $Y_i$ 's as defined earlier are independent) is used to prove that the sampling distribution is normal and (even more importantly) that the standard deviation of the sampling distribution is  $\sigma/\sqrt{n}$



*Consequences* (More detail later):

- I. If the conclusion of the theorem is true, the sampling distribution of  $\bar{Y}_n$  is narrower than the original distribution of Y
  - In fact, the conclusion of the theorem gives us an idea of just how narrow it is, depending on  $n$ .
  - This will allow us to construct a hypothesis test.
- II. The only way we know the conclusion is true is if we know the hypotheses of the theorem (the model assumptions) are true.
- III. Thus: *If the model assumptions are not true, then we do not know that the theorem is true, so we do not know that the hypothesis test is valid.*

In the example (large sample z-test for a mean), this translates to:

***If the sample is not a simple random sample, or if the random variable is not normal, then the reasoning establishing the validity of the test breaks down.***

**Comments:**

1. Different hypothesis tests have different model assumptions.
  - Some tests apply to random samples that are not simple.
  - For many tests, the model assumptions consist of *several* assumptions.
  - If *any one* of these model assumptions is not true, we do not know that the test is valid.
2. Many techniques are *robust to some departures from at least some model assumptions*.
  - This means that if the particular assumption is not too far from true, then the technique is still approximately valid.
  - More later.
3. *Using a hypothesis test without paying attention to whether or not the model assumptions are true and whether or not the technique is robust to possible departures from model assumptions is a common mistake in using statistics.* (More later.)

**OVERVIEW OF****FREQUENTIST CONFIDENCE INTERVALS**

*Before continuing the discussion of hypothesis tests, it will be helpful to first discuss the related concept of confidence intervals.*

**The General Situation:**

- We are considering a *random variable*  $Y$ .
- We are interested in a certain *parameter* (e.g., a proportion, or mean, or regression coefficient, or variance) associated with the random variable  $Y$ .
- We do not know the value of the parameter.
- *Goal 1:* We would like to estimate the unknown parameter, using data from a sample.
- *Goal 2:* We would also like to get some sense of how good our estimate is.

***The first goal is usually easier than the second.***

*Example:* If the parameter we are interested in estimating is the *mean of the random variable*, we can estimate it using a *sample mean*.



**The idea for the second goal** (getting some sense of how good our estimate is):

- Although we typically have just one sample at hand when we do statistics, *the reasoning used in frequentist (classical) inference depends on thinking about all possible suitable samples of the same size  $n$ .*
- Which samples are considered "suitable" will depend on the particular statistical procedure to be used.
- Each statistical procedure has *model assumptions* that are needed to ensure that the reasoning behind the procedure is sound.
- The model assumptions determine which samples are "suitable."

*Example:* The parameter we are interested in estimating is the population mean  $\mu = E(Y)$  of the random variable  $Y$ . The confidence interval procedure is the large-sample z-procedure.

- The model assumptions for this procedure are: The random variable is normal, and samples are simple random samples.
- Thus in this case, "suitable sample" means "simple random sample".
- We will also assume  $Y$  is normal, so that the procedure will be legitimate.
- We'll use  $\sigma$  to denote the standard deviation of  $Y$ .
- So we collect a simple random sample, say of size  $n$ , consisting of observations  $y_1, y_2, \dots, y_n$ .
  - For example, if  $Y$  is "height of an adult American male", we take a sample random sample of  $n$  adult American males;  $y_1, y_2, \dots, y_n$  are their heights.
- We use the sample mean  $\bar{y} = (y_1 + y_2 + \dots + y_n)/n$  as our estimate of  $\mu$ .
  - This is an example of a **point estimate** -- a numerical estimate with no indication of how good the estimate is.

- To get an idea of how good our estimate is, we look at *all possible simple random samples of size  $n$  from  $Y$* .
  - In the specific example, we consider all possible simple random samples of adult American males, and for each sample of men, the list of their heights.
- One way we can get a sense of how good our estimate is in this situation is to consider the *sample means  $\bar{y}$  for all possible simple random samples of size  $n$  from  $Y$* .
  - This amounts to defining a new random variable, which we will call  $\bar{Y}_n$  (read  $Y$ -bar sub  $n$ ).
  - We can describe the random variable as  $\bar{Y}_n$  as "sample mean of a simple random sample of size  $n$  from  $Y$ ", or perhaps more clearly as: "pick a simple random sample of size  $n$  from  $Y$  and calculate its sample mean".
  - Note that *each value of  $\bar{Y}_n$  is an estimate of the population mean  $\mu$* .
- This new random variable  $\bar{Y}_n$  has a distribution. This is called a *sampling distribution*, since it arises from considering varying samples.
  - The values of  $\bar{Y}_n$  are all the possible values of sample means  $\bar{y}$  of simple random samples of  $Y$  -- i.e, the values of our estimates of  $\mu$ .
  - *The distribution of  $\bar{Y}_n$  gives us information about the variability (as samples vary) of our estimates of the population mean  $\mu$* .
  - There is a mathematical theorem that proves that *if the model assumptions are true, the mean of the sampling distribution is also  $\mu$* .
  - This theorem also says that the sampling distribution is normal, and its standard deviation is  $\frac{\sigma}{\sqrt{n}}$
  - The chart and picture below summarize some of this information.