

NOTES FOR SUMMER STATISTICS INSTITUTE COURSE

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

Day 3: Type I and II Errors, Power, and Multiple Inference

MAY 20 – 23, 2013

Instructor: Martha K. Smith

CONTENTS OF DAY 3

I. Type II error	3
II. Considering both types of error together	5
III. Deciding what significance level to use	9
IV. Power of a statistical procedure	13
Factors affecting power	19
Significance level	19
Sample size	20
Variance	23
Experimental design	25
Calculating sample size	25 and appendix
Detrimental effects of underpowered or overpowered studies	26
V. Common mistakes involving power	29
VI. The File Drawer Problem (“Publication Bias”)	33
VII. Multiple inference	39

(If time permits, we will start on some of the material listed for Day 4)

I. TYPE II ERROR

(Recall: Falsely rejecting a true null hypothesis is called a **Type I error**.)

Not rejecting the null hypothesis when in fact the alternate hypothesis is true is called a **Type II error**.

- Example 2 below provides a situation where the concept of Type II error is important.

New complication: "The alternate hypothesis" in the definition of Type II error may refer to the alternate hypothesis in a hypothesis test (a "general" alternate hypothesis), or it may refer to a "specific" alternate hypothesis.

Example: In a t-test for a sample mean μ , with null hypothesis " $\mu = 0$ " and alternate hypothesis " $\mu > 0$ ":

- We might talk about the Type II error relative to the general alternate hypothesis " $\mu > 0$ ".
- Or we might talk about the Type II error relative to the specific alternate hypothesis " $\mu = 1$ " (or " $\mu = 0.5$ ", or ...).
- Note that *the specific alternate hypothesis is a special case of the general alternate hypothesis*.

In practice, people often work with Type II error relative to a *specific* alternate hypothesis.

- In this situation, the probability of Type II error relative to the specific alternate hypothesis is often called β .
- In other words, β is the probability of making the *wrong* decision when the *specific alternate* hypothesis is *true*.
- The specific alternative is considered since it is more feasible to calculate β than the probability of Type II error relative to the general alternative.
- See the discussion of power below for related detail.

II: CONSIDERING BOTH TYPES OF ERROR TOGETHER

The following table summarizes Type I and Type II errors:

		Truth (for population studied)	
		Null Hypothesis True	Null Hypothesis False
Decision (based on sample)	Reject Null Hypothesis	<i>Type I Error</i>	<i>Correct Decision</i>
	Don't reject Null Hypothesis	<i>Correct Decision</i>	<i>Type II Error</i>

Analogy: Results of a criminal trial.

- The “null hypothesis” is "defendant is not guilty."
- The “alternate hypothesis” is "defendant is guilty."
- A Type I error would correspond to convicting an innocent person.
- Type II error would correspond to setting a guilty person free.
- The analogous table would be:


		Truth	
		Not Guilty	Guilty
Verdict	Guilty	<i>Type I Error -- Innocent person goes to jail (and maybe guilty person goes free)</i>	<i>Correct Decision</i>
	Not Guilty	<i>Correct Decision</i>	<i>Type II Error -- Guilty person goes free</i>

Note:

- This could be more than just an analogy if the verdict hinges on statistical evidence (e.g., a DNA test), and where rejecting the null hypothesis would result in a verdict of guilty, and not rejecting the null hypothesis would result in a verdict of not guilty.
- This analogy/example shows that *sometimes* a Type I error can be more serious than a Type II error. (However, this is *not always* the case).

The following diagram illustrates the Type I error and the Type II error

- against the specific alternate hypothesis " $\mu = 1$ "
- in a hypothesis test for a population mean μ ,
- with
 - null hypothesis " $\mu = 0$,"
 - alternate hypothesis " $\mu > 0$,"
 - and significance level $\alpha = 0.05$.



In the diagram,

- The blue (leftmost) curve is the *sampling* distribution of the test statistic *assuming the null hypothesis* " $\mu = 0$."
- The green (rightmost) curve is the *sampling* distribution of the test statistic *assuming the specific alternate hypothesis* " $\mu = 1$ ".
- The vertical red line shows the cut-off for rejection of the null hypothesis:
 - The null hypothesis is rejected for values of the test statistic to the *right* of the red line (and *not* rejected for values to the *left* of the red line).
- The area of the diagonally hatched region to the *right* of the red line and under the *blue* curve is the probability of type I error (α).
- The area of the horizontally hatched region to the *left* of the red line and under the *green* curve is the probability (β) of Type II error against the specific alternative.

→ What happens to the Type II error probability (β) if we:

- a. increase α ?
- b. decrease α ?

III. DECIDING WHAT SIGNIFICANCE LEVEL TO USE

This should be done *before analyzing* the data -- *preferably before gathering the data*. This is important for (at least) two reasons:

- 1) The significance level desired is one criterion in deciding on an appropriate sample size.
 - See discussion of Power below.
- 2) If more than one hypothesis test is planned, additional considerations need to be taken into account.
 - See discussion of Multiple Inference below.

The choice of significance level should be based on the consequences of Type I and Type II errors:

1. If the *consequences of a Type I error are serious or expensive*, a very *small* significance level is appropriate.

Example 1: Two drugs are being compared for effectiveness in treating the same condition.

- Drug 1 is very affordable, but Drug 2 is extremely expensive.
- The null hypothesis is “both drugs are equally effective.”
- The alternate is “Drug 2 is more effective than Drug 1.”
- In this situation, a Type I error would be deciding that Drug 2 is more effective, when in fact it is no better than Drug 1, but would cost the patient much more money.
- That would be undesirable from the patient’s perspective, so a *small* significance level is warranted.

2. If the consequences of a Type I error are not very serious (and *especially if a Type II error has serious consequences*), then a *larger* significance level is appropriate.

Example 2: Two drugs are known to be equally effective for a certain condition.

- They’re also each equally affordable.
- However, there is some suspicion that Drug 2 causes a serious side effect in some patients, whereas Drug 1 has been used for decades with no reports of serious side effects.
- The null hypothesis is “the incidence of serious side effects in both drugs is the same”.
- The alternate is “the incidence of serious side effects in Drug 2 is greater than that in Drug 1.”
- Falsely rejecting the null hypothesis when it is in fact true (Type I error) would have no great consequences for the consumer.
- But a Type II error (i.e., failing to reject the null hypothesis when in fact the alternate is true, which would result in deciding that Drug 2 is no more harmful than Drug 1 when it is in fact more harmful) could have serious consequences from a consumer and public health standpoint.
- So setting a large significance level is appropriate.

Comments:

- Neglecting to think adequately about possible consequences of Type I and Type II errors (and deciding acceptable levels of Type I and II errors based on these consequences) *before* conducting a study and analyzing data is a **common mistake** in using statistics.
- Sometimes there are serious consequences of each alternative, so compromises or weighing priorities may be necessary.
 - The trial analogy illustrates this well: Which is better or worse, imprisoning an innocent person or letting a guilty person go free?
 - *This is a value judgment; value judgments are often involved in deciding on significance levels.*
 - *Trying to avoid the issue by always choosing the same significance level is itself a value judgment.*
- Different people may decide on different standards of evidence.
 - This is another reason why *it's important to report p-values even if you set a significance level.*
 - It's *not* enough just to say, "significant at the .05 level," "significant at the .01 level," etc. Such reporting is a **common mistake**.

- Sometimes different stakeholders have different interests that compete (e.g., in the second example above, the developers of Drug 2 might prefer to have a smaller significance level.)
 - This is another reason why it's important to report p-values in publications.
- See Wuensch (1994) for more discussion of considerations involved in deciding on reasonable levels for Type I and Type II errors.
- See also the discussion of Power below.
- Similar considerations hold for setting confidence levels for confidence intervals.

IV: POWER OF A STATISTICAL PROCEDURE

Overview

The *power* of a statistical procedure can be thought of as *the probability that the procedure will detect a true difference of a specified type*.

- As in talking about p-values and confidence levels, the reference category for "probability" is the sample.
- Thus, power is the probability that a randomly chosen sample
 - satisfying the model assumptions
 - will give evidence of a difference of the specified type when the procedure is applied,
 - if the specified difference does indeed occur in the population being studied.
- Note also that power is a conditional probability: the probability of detecting a difference, *if* indeed the difference does exist.

In many real-life situations, there are reasonable conditions that we would be interested in being able to detect, and others that would not make a practical difference.

Examples:

- If you can only measure the response to within 0.1 units, it doesn't really make sense to worry about falsely rejecting a null hypothesis for a mean when the actual value of the mean is within less than 0.1 units of the value specified in the null hypothesis.
- Some differences are of no practical importance -- for example, a medical treatment that extends life by 10 minutes is probably not worth it.

In cases like these, neglecting power could result in one or more of the following:

- Doing more work or going to more expense than necessary
- Obtaining results that are meaningless
- Obtaining results that don't answer the question of interest.

Elaboration

For many *confidence interval procedures*, power can be defined as:

The probability (again, the reference category is “samples”) that the procedure will produce an interval with a half-width of at most a specified amount.

For a *hypothesis test*, power can be defined as:

The probability (again, the reference category is “samples”) of rejecting the null hypothesis under a specified condition.

Example: For a one-sample t-test for the mean of a population, with null hypothesis $H_0: \mu = 100$, you might be interested in the probability of rejecting H_0 when $\mu \geq 105$, or when $|\mu - 100| > 5$, etc.


As with Type II Error, we may think of power for a hypothesis test in terms of *power against a specific alternative* rather than against a general alternative.

Example: If we are performing a hypothesis test for the mean of a population, with null hypothesis $H_0: \mu = 0$ and alternate hypothesis $\mu > 0$, we might calculate the power of the test *against the specific alternative* $H_1: \mu = 1$, or *against the specific alternative* $H_3: \mu = 3$, etc.

The picture below shows three sampling distributions:

- The sampling distribution assuming H_0 (*blue; leftmost curve*)
- The sampling distribution assuming H_1 (*green; middle curve*)
- The sampling distribution assuming H_3 (*yellow; rightmost curve*)

The red line marks the cut-off corresponding to a significance level $\alpha = 0.05$.



→ Where would we reject/not reject the null hypothesis?

From the above, we conclude (how?) that:

- The area under the *blue* curve to the *right of the red line* is 0.05.
- The area under the *green* curve to the *right of the red line* is the probability of rejecting the null hypothesis ($\mu = 0$) if the specific alternative $H_1: \mu = 1$ is true.
 - In other words, this area is *the power of the test against the specific alternative $H_1: \mu = 1$* .
 - We can see in the picture that in this case, this power is greater than 0.05, but noticeably less than 0.50.
- Similarly, the area under the *yellow* curve to the *right of the red line* is *the power of the test against the specific alternative $H_3: \mu = 3$* .
 - Notice that the power in this case is much larger than 0.5.

This illustrates the general phenomenon that *the farther an alternative is from the null hypothesis, the higher the power of the test to detect it*.

→ See Claremont Graduate University WISE Project Statistical Power Demo for an interactive illustration.

Note:

- For most tests, it *is* possible to calculate the power against a specific alternative, at least to a reasonable approximation. (More below and in Appendix)
- It is *not* usually possible to calculate the power against a general alternative, since the general alternative is made up of infinitely many possible specific alternatives.

Power and Type II Error

Recall: The Type II Error rate β of a test against a specific alternate hypothesis test is represented in the diagram above as the area under the sampling distribution curve for that alternate hypothesis and to the *left* of the cut-off line for the test. Thus

$$\begin{aligned} \beta + (\text{Power of a test against a specific alternate hypothesis}) \\ &= \text{total area under sampling distribution curve} \\ &= 1, \end{aligned}$$

so

$$\text{Power} = 1 - \beta$$

Factors that Affect the Power of a Statistical Procedure

Power depends on several factors in addition to the difference to be detected.



1. Significance Level

This can be seen in the diagram illustrating power: *Increasing* the significance level α will move the red line to the *left*, and hence will *increase power*. Similarly, decreasing significance level decreases power.

2. Sample Size

Example: The pictures below each show the sampling distribution for the mean under the null hypothesis $\mu = 0$ (blue -- on the left in each picture) together with the sampling distribution under the alternate hypothesis $\mu = 1$ (green -- on the right in each picture), but *for different sample sizes*.

- The first picture is for sample size $n = 25$; the second picture is for sample size $n = 100$.
 - Why are the curves in the second graph skinnier?
- Note that both graphs are in the same scale. In both pictures, the blue curve is centered at 0 (corresponding to the null hypothesis) and the green curve is centered at 1 (corresponding to the alternate hypothesis).
- In each picture, the vertical red/orange line is the cut-off for rejection with $\alpha = 0.05$ (for a one-tailed test) -- that is, in each picture, the area under the *blue* curve to the right of the line is 0.05.
- In each picture, the area under the *green* curve to the right of the red line is the power of the test against the alternate depicted. Note that this area is *larger* in the second picture (the one with larger sample size) than in the first picture.



This illustrates the general situation:

Larger sample size gives larger power.

The reason is essentially the same as in the example: Larger sample size gives a narrower sampling distribution, which means there is less overlap in the two sampling distributions (for null and alternate hypotheses).

→ See Claremont University's Wise Project's Statistical Power Applet (http://wise.cgu.edu/powermod/power_applet.asp) for an interactive demonstration of the interplay between sample size and power for a one-sample z-test.



Note: Sample size needed typically increases at an increasing rate as power increases. (e.g., in the above example, increasing the sample size by a factor of 4 increases the power by a factor of about 2; the graphics aren't accurate enough to show this well.)

3. Variance

Power also depends on variance: *smaller variance yields higher power*.

Example: The pictures below each show the sampling distribution for the mean under the null hypothesis $\mu = 0$ (blue -- on the left in each picture) together with the sampling distribution under the alternate hypothesis $\mu = 1$ (green -- on the right in each picture), both with sample size 25, but for different standard deviations of the underlying distributions. (Different standard deviations might arise from using two different measuring instruments, or from considering two different populations.)

- In the first picture, the standard deviation is 10; in the second picture, it is 5.
- Note that both graphs are in the same scale. In both pictures, the blue curve is centered at 0 (corresponding to the null hypothesis) and the green curve is centered at 1 (corresponding to the alternate hypothesis).
- In each picture, the red/orange line is the cut-off for rejection with $\alpha = 0.05$ (for a one-tailed test) -- that is, in each picture, the area under the *blue* curve to the right of the line is 0.05.
- In each picture, the area under the *green* curve to the right of the line is the power of the test against the alternate depicted. Note that this area is *larger* in the second picture (the one with smaller standard deviation) than in the first picture.



→ See Claremont University's Wise Project's Statistical Power Applet at http://wise.cgu.edu/powermod/power_applet.asp or the Rice Virtual Lab in Statistics' Robustness Simulation at http://onlinestatbook.com/stat_sim/robustness/index.html for an interactive demonstration.

Note: Variance can sometimes be reduced by using a better measuring instrument, restricting to a subpopulation, or by choosing a better experimental design (see below).

4. Experimental Design

Power can sometimes be increased by adopting a different experimental design that has lower error variance. For example, stratified sampling or blocking can often reduce error variance and hence increase power. However,

- The power calculation will depend on the experimental design.
- The statistical analysis will depend on the experimental design. (To be discussed tomorrow.)
- For more on designs that may increase power, see Lipsey (1990) or McClelland (2000)

Calculating Sample Size to Give Desired Power: The dependence of power on sample size *in principle* lets us figure out beforehand what sample size is needed to detect a specified difference, with a specified power, at a given significance level, if that difference is really there.

- *In practice*, details on figuring out sample size will vary from procedure to procedure. See the Appendix for discussion of some of the considerations involved.

Detrimental Effects of Underpowered or Overpowered Studies

The most straightforward consequence of **underpowered** studies (i.e., those with low probability of detecting an effect of practical importance) is that *effects of practical importance are not detected*.


A second, more subtle consequence:

Underpowered studies result in *a larger variance of the estimates of the parameter being estimated*.

- For example, in estimating a population mean, the sample means of studies with low power have high variance.
- In other words, *the sampling distribution of sample means is wide*.

This is illustrated in the following picture:

- It shows the sampling distributions of the mean for a variable with zero mean when sample size $n = 25$ (red/orange) and when $n = 100$ (blue).
- The vertical lines toward the right of each sampling distribution show the cut-off for a one-sided hypothesis test with null hypothesis $\mu = 0$ and significance level $\alpha = .05$.
- Notice that:
 - The sampling distribution for the *smaller* sample size ($n = 25$) is *wider* than the sampling distribution for the larger sample size ($n = 100$).
 - Thus, *when the null hypothesis is rejected with the smaller sample size $n = 25$, the sample mean tends to be noticeably larger than when the null hypothesis is rejected with the larger sample size $n = 100$.*



This reflects the general phenomenon that *studies with low power have a larger chance of having a large sample mean than studies with high power.*

In particular, *when there is a Type I error (falsely rejecting the null hypothesis), the effect will appear to be stronger with a small sample size (lower power) than with a large sample size (higher power).*

- *This may suggest an exaggerated effect, or even one that is not there.*
- *Thus, when studies are underpowered, the literature is likely to be inconsistent and often misleading.*
- *This problem is increased because of the “File Drawer Problem” (discussed below).*

Overpowered studies waste resources.

- When human or animal subjects are involved, an overpowered study can be considered unethical.
 - For more on ethical considerations in animal studies, see Festing (2010) or Kilkenny et al (2010)
- More generally, an overpowered study may be considered unethical if it wastes resources.

A common practice is to compromise between over-power and under-power is to try for power around .80.

- *However, power needs to be considered case-by-case, balancing the risks of Type I and Type II errors.*

V: COMMON MISTAKES INVOLVING POWER

1. *Rejecting a null hypothesis without considering practical significance.*

- A study with large enough sample size will have high enough power to detect minuscule differences that aren't practically significant.
- Since power typically increases with increasing sample size, practical significance is important to consider.

2. *Accepting a null hypothesis when a result is not statistically significant, without taking power into account.*

- Power decreases with decreasing sample size.
- Thus *a small sample size may not be able to detect an important difference.*
- If there's strong evidence that the power of a procedure will indeed detect a difference of practical importance, then accepting the null hypothesis is appropriate.
 - However, it may be better to use a *test for equivalence*; see the Appendix for references.
- Otherwise "accepting the null hypothesis" is *not appropriate* -- all we can legitimately say then is that we fail to reject the null hypothesis.

3. *Being convinced by a research study with low power.*

As discussed above, *underpowered studies are likely to be inconsistent and are often misleading.*

4. *Neglecting to do a power analysis/sample size calculation before collecting data*

- If you use a sample size that's *too small* to detect a difference of practical significance, you may get a result that's not statistically significant even though there is a difference of practical significance.
 - Thus you've expended considerable effort to obtain a result that doesn't really answer the question of interest.
- If you use a sample size that is *larger than needed* to detect a relevant difference, you've also wasted resources.

5. *Neglecting to take multiple inference into account when calculating power.*

If more than one inference procedure is used for a data set, then power calculations need to take that into account. (*More on this below.*)

- Doing a power calculation for just one inference will result in an underpowered study. (*More on this tomorrow*)
- For more detail, see Maxwell and Kelley (2011) and Maxwell (2004)

6. Using standardized effect sizes rather than considering the particulars of the question being studied.

"Standardized effect sizes" (sometimes called "canned" effect sizes) are expressions involving more than one of the factors that needs to be taken into consideration in considering appropriate levels of Type I and Type II error in deciding on power and sample size. *Examples:*

- Cohen's effect size d is the ratio of the raw effect size (e.g., difference in means when comparing two groups) and a suitable standard deviation.
 - But each of these typically needs to be considered individually in designing a study and determining power; it's not necessarily the ratio that's important. (See Appendix)
- The correlation (or squared correlation) in regression.
 - The correlation in simple linear regression involves three quantities: the slope, the y standard deviation, and the x standard deviation.
 - Each of these three typically needs to be considered individually in designing the study and determining power and sample size.
 - In multiple regression, the situation may be even more complex.

For specific examples illustrating these points, see Lenth, (2000) and (2001)

7. Confusing retrospective power and prospective power.

- Power as defined above for a hypothesis test is also called *prospective* or *a priori* power.
 - It is a conditional probability, $P(\text{reject } H_0 \mid H_a)$, *calculated without using the data to be analyzed.*
 - In fact, it's best calculated before even gathering the data, and taken into account in the data-gathering plan.
- *Retrospective* power is calculated *after* the data have been collected, *using the data.*
- Depending on how retrospective power is calculated, it might be legitimate to use to estimate the power and sample size for a *future* study, but cannot legitimately be used as describing the power of the study from which it is calculated.
- However, some methods of calculating retrospective power calculate the power to detect the effect observed in the data -- which misses the whole point of considering practical significance. These methods typically yield simply a transformation of p -value. See Lenth (2000) for more detail.
- See Hoening and Heisley (2001) and Wuensch et al (2003) for more discussion and further references.

VI. THE FILE DRAWER PROBLEM (PUBLICATION BIAS)

Publication bias refers to the influence of the results of a study on whether or not the study is published.

Examples of how results might influence the publication decision:

- Whether or not the results are statistically significant.
- Whether or not the results are practically significant.
- Whether or not the results agree with the hopes or expectations of the researcher or sponsor.

Publication bias is also called the **file drawer problem**, especially when the nature of the bias is that studies which fail to reject the null hypothesis (i.e., that do not produce a statistically significant result) are less likely to be published than those that do produce a statistically significant result.

Several studies (see Sterling et al 1995, Song et al 2009, and Hopewell et al 2009) have found evidence of publication bias in the research literature.

Consequences of the File Drawer Problem:


1. Investigators may spend unnecessary effort conducting research on topics that have already been well researched but not reported because results were negative.

2. Effects that are not real may appear to be supported by research.

- Recall that if a significance level of 0.05 is set, then in repeated studies, about 5% of studies of a situation where the null hypothesis is true will falsely reject the null hypothesis
- Thus, *if just (or even predominantly) the statistically significant studies are published, the published record misrepresents the true situation.*

3. Furthermore, papers that are published because of Type I errors, if underpowered, may show an exaggerated effect size, increasing the misrepresentation.

- *Recall from the discussion of power:*




Some Methods Proposed to Detect Publication Bias:

1. Rosenthal (1979) proposed a method, based on probability calculations, for deciding whether or not a finding is "resistant to the file drawer threat."

- This method has become known as the *fail-safe file drawer (or FSFD) analysis*.
- It involves calculating a "fail-safe number" which is used to estimate whether or not the file-drawer problem is likely to be a problem for a particular review or meta-analysis.
- Scargle (2000) has criticized Rosenthal's method on the grounds that it fails to take into account the bias in the "file drawer" of unpublished studies, and thus can give misleading results.
- More recently, Simonsohn et al (2013) have pointed out that the prevalence of "p-hacking" (to be discussed tomorrow) invalidates Rosenthal's method.

2. Various types of plots have been used to try to detect publication bias. These plot some measure of precision against effect size, or vice-versa.

- *Example:*



- Some such plots are called "funnel plots" because they typically have a funnel shape.
- However, Lau et al (2006) point out some problems in using these plots.
- Recently, Simonsohn et al (2013) have proposed a method called "p-curving" to detect possible publication bias and/or p-hacking (to be discussed tomorrow).

3. Research registries have been instituted in some areas.

- For example, certain clinical trials are now required by law to be registered at the NIH database ClinicalTrials.gov.
- These are beginning to point to possible systemic problems, such as:

“We are finding that in some cases, investigators cannot explain their trial, cannot explain their data. Many of them rely on the biostatistician, but some biostatisticians can’t explain the trial design.

So there is a disturbing sense of some trials being done with no clear intellectual leader.”

*Deborah Zarin, Director, ClinicalTrials.gov,
quoted in interview in Marshall (2011)*

4. Additionally, full data may reveal a different story from what appears in published papers and conference proceedings.

- Although such data is increasingly becoming more available, obtaining it can often still be difficult or impossible.
- See Doshi et al (2012) for an example.
 - The editorial preface to this article says: “After publication of a Cochrane review into the effectiveness of oseltamivir [Tamiflu] in 2009, the reviewers got access to thousands of pages of previously unavailable data. [The authors] describe how it shook their faith in published reports and changed their approach to systematic reviews.”
 - The authors received over 3000 pages of study reports from one drug company, and over 25,000 pages from the European Medicines Agency.
 - The new review based on the additional data took the equivalent of two full-time researchers for 14 months.
 - They also point out how calculations based on electronic data bases may be questionable (e.g., because of lack of standardized definitions for complications).
- An accurate “history” of computational methods used is also an important source of data on research methods.
 - One method for facilitating this is sweave, <http://www.stat.uni-muenchen.de/~leisch/Sweave/>.

See the Appendix for suggestions for dealing with the File Drawer Problem.

VII. MULTIPLE INFERENCE

"Recognize that any frequentist statistical test has a random chance of indicating significance when it is not really present. Running multiple tests on the same data set at the same stage of an analysis increases the chance of obtaining at least one invalid result. Selecting the one "significant" result from a multiplicity of parallel tests poses a grave risk of an incorrect conclusion. Failure to disclose the full extent of tests and their results in such a case would be highly misleading."

Professionalism Guideline 8, *Ethical Guidelines for Statistical Practice*, American Statistical Association, 1997

Performing more than one statistical inference procedure on the same data set is called **multiple inference**, or **joint inference**, or **simultaneous inference**, or **multiple testing**, or **multiple comparisons**, or **the problem of multiplicity**.

Performing multiple inference *without adjusting the Type I error rate accordingly* is a **common error** in research using statistics.

- For example, A. M. Strasak et al (2007) examined all papers from 2004 issues of the *New England Journal of Medicine* and *Nature Medicine* and found that 32.3% of those from *NEJM* and 27.3% from *Nature Medicine* were "Missing discussion of the problem of multiple significance testing if occurred."
- These two journals are considered the top journals (according to impact figure) in clinical science and in research and experimental medicine, respectively.

The Problem

Recall: If you perform a hypothesis test using a certain significance level (we'll use 0.05 for illustration), and if you obtain a p-value less than 0.05, then there are three possibilities:

1. The model assumptions for the hypothesis test are not satisfied in the context of your data.
2. The null hypothesis is false.
3. Your sample happens to be one of the 5% of samples satisfying the appropriate model conditions for which the hypothesis test gives you a Type I error – i.e., you falsely reject the null hypothesis.

Now suppose you are performing *two* hypothesis tests, *using the same data* for both.

- Suppose that in fact all model assumptions are satisfied and *both* null hypotheses are true.
- *There is in general no reason to believe that the samples giving a Type I error for one test will also give a Type I error for the other test.*
- See Jerry Dallal's Simulation
- So we need to consider the *joint Type I error rate*

Joint Type I error rate: This is the probability that a randomly chosen sample (of the given size, satisfying the appropriate model assumptions) will give a Type I error for *at least one* of the hypothesis tests performed.

The joint Type I error rate is also known as the **overall Type I error rate**, or **joint significance level**, or the **simultaneous Type I error rate**, or the **family-wise error rate (FWER)**, or the **experiment-wise error rate**, etc.

- The acronym FWER is becoming more and more common, so will be used in the sequel, often along with another name for the concept as well.

An *especially serious* form of neglect of the problem of multiple inference is the one alluded to in the quote from the ASA ethics page:

- Trying several tests and reporting just one significant test, without disclosing how many tests were performed or correcting the significance level to take into account the multiple inference.
- *Don't do it!*
- To help you remember: Remember Jelly Beans, <http://xkcd.com/882/>

Multiple inference with confidence intervals

The problem of multiple inference also occurs for confidence intervals.

- In this case, we need to focus on the *confidence level*.
- *Recall:* A 95% confidence interval is an interval obtained by using a procedure that, for 95% of all suitably random samples, of the given size from the random variable and population of interest, produces an interval containing the parameter we are estimating (assuming the model assumptions are satisfied).
- In other words, the procedure does what we want (i.e. gives an interval containing the true value of the parameter) for 95% of suitable samples.
- *If we're using confidence intervals to estimate two parameters, there's no reason to believe that the 95% of samples for which the procedure "works" for one parameter (i.e. gives an interval containing the true value of the parameter) will be the same as the 95% of samples for which the procedure "works" for the other parameter.*
- If we're calculating confidence intervals for more than one parameter, we can talk about the **joint (or overall or simultaneous or family-wise or experiment-wise) confidence level**.
- For example, a group of confidence intervals (for different parameters) has an **overall 95% confidence level** (or **95% family-wise confidence level**, etc.) if the intervals are calculated using a procedure which, for 95% of all suitably random samples, of the given size from the population of interest, produces for *each* parameter in the group an interval containing that parameter (assuming the model assumptions are satisfied).

What to do about multiple inference

Unfortunately, there's no simple formula to cover all cases:

- Depending on the context, the samples giving Type I errors for two tests might be the same, they might have no overlap, or they could be somewhere in between.
- Various techniques for bounding the FWER (joint Type I error rate) have been devised for various special circumstances.
- Only two fairly general methods (Bonferroni and False Discovery Rate) will be discussed here.
- For more information on other more specialized methods, see, e.g., Hochberg and Tamhane (1987) and Miller (1981)
- See Efron (2010) for both an account of the history of the subject and discussion of recent developments in dealing with multiple inference, especially in large data sets.

Bonferroni method:

Fairly basic probability calculations show that *if the sum of the individual Type I error rates for different tests is less than α , then the overall ("family-wise") Type I error rate (FWER) for the combined tests will be at most α .*

- So, for example, if you are performing five hypothesis tests and would like an FWER (overall significance level) of at most 0.05, then using significance level 0.01 for *each* test will give an FWER (overall significance level) of at most 0.05.
- Similarly, if you are finding confidence intervals for five parameters and want an overall confidence level of 95%, using the 99% confidence level for each confidence interval will give you overall confidence level at least 95%. (Think of confidence level as $1 - \alpha$.)

The Bonferroni method can be used as a fallback method when no other method is known to apply.

- However, if a method that applies to the specific situation is available, it will often be better (less conservative) than the Bonferroni method, so calculate by both methods and compare.
- Holm's procedure (which depends on the Bonferroni idea, but in a more sophisticated way) is a relatively easy (e.g., on a spreadsheet) method that gives higher power than the basic Bonferroni method. It's described various places on the web – e.g., http://en.wikipedia.org/wiki/Holm%E2%80%93Bonferroni_method

The basic Bonferroni method is also useful for apportioning the overall Type I error between different types of inference.

- For example, if three confidence intervals and two hypothesis tests are planned, and an overall Type I error rate of .05 is desired, then using 99% confidence intervals and individual significance rates .01 for the hypothesis tests will achieve this.
- This can also be used to apportion Type I error rate between *pre-planned inference* (the inference planned as part of the design of the study) and “data-snooping” inferences (inferences based on looking at the data and noticing other things of interest; more below).
- However, *this apportioning should be done before analyzing the data.*

Whichever method is used, *it's important to make the calculations based on the number of tests that have been done, not just the number that are reported.*

- Remember Jelly Beans!
- See Data Snooping tomorrow for more discussion.

False discovery rate:

An alternative to bounding Type I error was introduced by Benjamini and Hochberg (1995): bounding the *False Discovery Rate*.

The ***False Discovery Rate*** (FDR) of a group of tests is the *expected value of the ratio of falsely rejected hypotheses to all rejected hypotheses.*

("Expected value" refers to the mean of a distribution. Here, the distribution is the sampling distribution of the ratio of falsely rejected hypotheses to all rejected hypotheses tested.)

Note:

- The family-wise error rate (FWER) focuses on the possibility of making *any* Type I error among all the inferences performed.
- The false discovery rate (FDR) tells you what *proportion* of the *rejected* null hypotheses are, *on average*, really false.
- Bounding the FDR rather than the FWER may be a more reasonable choice when many inferences are performed, especially if there is little expectation of harm from falsely rejecting a null hypothesis.
- Thus it is increasingly being adopted in areas such as micro-array gene expression experiments or neuro-imaging.
- However, these may involve variations rather than the original definition given above.

As with the FWER, there are various methods of actually bounding the false discovery rate.

- For the original false discovery rate, see Benjamini and Hochberg (1995), Benjamini and Yekutieli (2001), and Benjamini and Yekutieli (2005)
- For variations of false discovery rate, see Efron (2010).

Subtleties and controversies

Bounding the overall Type I error rate (FWER) will reduce the power of the tests, compared to using individual Type I error rates.

- Some researchers use this as an argument against multiple inference procedures.
- The counterargument is the argument for multiple inference procedures to begin with: Neglecting them will produce excessive numbers of false findings, *so that the "power" as calculated from single tests is misleading.*
 - See Maxwell and Kelley (2011) and Maxwell (2004) for more details.
- Bounding the False Discovery Rate (FDR) will usually give higher power than bounding the overall Type I error rate (FWER).

Consequently, it's important to consider the particular circumstances, as in considering both Type I and Type II errors in deciding significance levels.

- *In particular, it's important to consider the consequences of each type of error in the context of the particular research.*

Examples:

1. A research lab is using hypothesis tests to screen genes for possible candidates that may contribute to certain diseases.
 - Each gene identified as a possible candidate will undergo further testing.
 - If the results of the initial screening are not to be published except in conjunction with the results of the secondary testing, *and* if the secondary screening is inexpensive enough that many second level tests can be run, then the researchers could reasonably decide to ignore overall Type I error in the initial screening tests, since there would be no harm or excessive expense in having a high Type I error rate.
 - However, if the secondary tests were expensive, the researchers would reasonably decide to bound either family-wise Type I error rate or False Discovery Rate.

2. Consider a variation of the situation in Example 1:
 - The researchers are using hypothesis tests to screen genes as in Example 1, but plan to publish the results of the screening *without* doing secondary testing of the candidates identified.
 - In this situation, ethical considerations warrant bounding either the FWER or the FDR -- *and* taking pains to emphasize in the published report that these results are just of a preliminary screening for possible candidates, and that these preliminary findings need to be confirmed by further testing.