

Workshop on: Noncentral Confidence Intervals and Power Analysis

Michael Smithson
Psychology, Australian National University

Overview

Definition:

Denote a population statistic by θ , whose value is unknown. We may define confidence intervals for values of θ given a confidence level of $100(1 - \alpha)\%$, where α lies between 0 and 1, and a sample size of N . Confidence intervals may have an upper limit or a lower limit, or both.

- An upper confidence limit (U) is a value that, under repeated random samples of size N , may be expected to exceed θ 's true value $100(1 - \alpha)\%$ of the time.
- A lower confidence limit (L) is a value that, under repeated random samples of size N , may be expected to fall below θ 's true value $100(1 - \alpha)\%$ of the time. The traditional confidence interval uses lower and upper limits that contain θ 's true value $100(1 - \alpha)\%$ of the time.

The interval often is written as $[L, U]$, and sometimes writers will express the interval and its confidence level by writing

$$\Pr(L \leq \theta \leq U) = 1 - \alpha.$$

Example 1.

Consider a pollster who collects a random sample of 1000 people from the voting population. Of the 1000 people, 410 said they would vote for candidate Aibici. So, the sample proportion voting for Aibici is

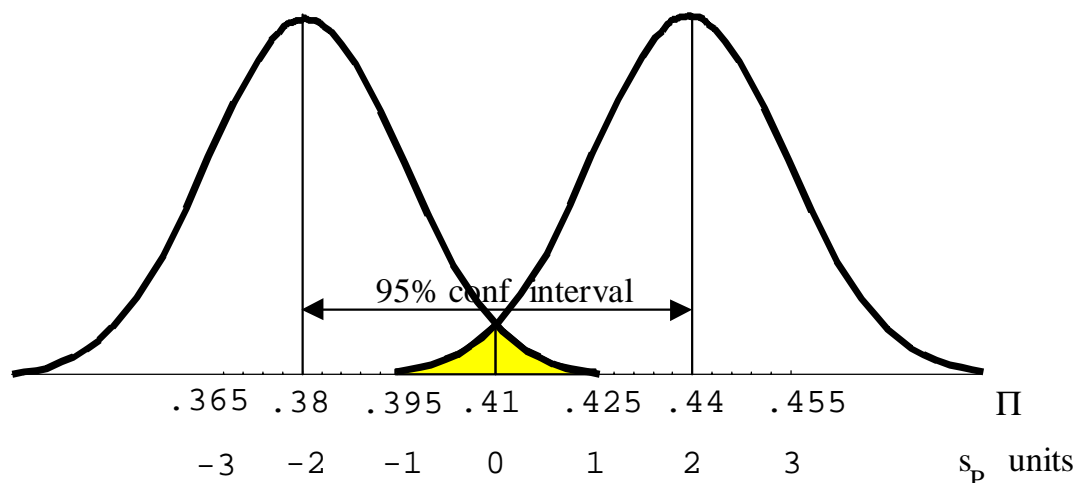
$$P = 410/1000 = 0.41.$$

The pollster would like to construct a 95% confidence interval for Π , the population value of the proportion. She uses the normal distribution as a large-sample approximation to the sampling distribution for the proportion, which has a mean of Π and a standard deviation s_P , with

$$s_P = \sqrt{\frac{P(1-P)}{N}}$$

To construct a $100(1 - \alpha)\%$ confidence interval given a sample P , we need to find the number of standard errors above and below P required to slice $\alpha/2$ from the tails of the normal distribution. Denote this number by $z_{\alpha/2}$.

- The lower limit of the confidence interval is the mean of a normal distribution for which $z_{\alpha/2}$ slices $\alpha/2$ from the upper tail. That mean is $P - w$, where $w = (z_{\alpha/2})(s_p)$.
- The upper limit of the confidence interval is the mean of a normal distribution for which $-z_{\alpha/2}$ slices $\alpha/2$ from the lower tail. That mean is $P + w$, where $w = (z_{\alpha/2})(s_p)$.
- The resulting confidence interval for Π is $[P - w, P + w]$.



We already have $P = 410/1000 = 0.41$, so

$$s_p = \sqrt{\frac{.41(1-.41)}{1000}} = 0.01555.$$

A 95% confidence level entails $\alpha = .05$ so $\alpha/2 = .025$ must be sliced from the lower and upper tails of the normal distribution. The required number of standard errors above and below P to do this is $z_{\alpha/2} = 1.96$.

The half-width of the confidence interval is therefore

$$w = (z_{\alpha/2})(s_p) = (1.96)(0.01555) = 0.0305.$$

- The lower limit is $L = P - w = .41 - .0305 = .38$, and
- the upper limit is $U = P + w = .41 + .0305 = .44$.
- The 95% confidence interval for Π is therefore $[.38, .44]$.

Confidence Intervals as "Inverted" Significance Tests


Then 95% confidence interval in Example 1 contains all the "null hypothesis" values of Π that we cannot reject, using a two-tailed test and $\alpha = .05$.

- We could reject any hypothetical value of Π below .38 because $P = .41$ cuts off less than .025 of its corresponding distribution.
- Likewise, we could reject any hypothetical value of Π above .44.
- We could not reject any hypothetical value of Π from .38 to .44.

Example 2.

Suppose we are observing independent trials in an experiment that have only two possible outcomes, A and $\sim A$. We have observed $N = 15$ trials and 5 A's have occurred, so our observed proportion is $P = 5/15$.

Now suppose we would like to construct a 90% confidence interval for Π . This time we will use the binomial distribution as the sampling distribution for P.



Why .1417 and .5775?

- The heights of the bars in the sampling distribution for $\Pi = .1417$ from 5/15 on up are the probabilities of getting $P = 5/15$ or higher if Π actually is .1417. Those probabilities add up to 0.05.
- The heights of the bars in the sampling distribution for $\Pi = .5775$ from 5/15 on down are the probabilities of getting a sample statistic $P = 5/15$ or lower if population parameter Π actually is .8461. Those probabilities also add up to 0.05.

Central versus Noncentral Distributions

In both Examples 1 and 2 we are really sliding a sampling distribution up and down the axis by changing the value of Π until we find the appropriate lower and upper limits for the confidence interval.

However, because the binomial distribution has a unique shape for every value of Π , we need a computer program (or, in the past, very patient humans) to iteratively home in on these limits.

- With central or standardizable distributions, for any specified sample size and degrees of freedom the same standardized distribution may be used regardless of the hypothetical value of the population parameter.
- Noncentral or nonstandardizable distributions change shape depending on the value of the parameter, so a unique distribution must be computed for every parameter value under consideration.

Confidence Intervals, Significance Tests, Effect Size, and Power

Now suppose we want to test our sample mean against some hypothetical value, Π_h . The confidence interval above contains all the values of Π_h that we cannot reject at the α level. Any Π_h that lies outside this interval constitutes an hypothesis we can reject.

Three things will determine whether Π_h will lie inside or outside the confidence interval:

1. The size of the effect (how far Π_h is from the true population mean μ),
2. The confidence level ($1 - \alpha$), and
3. The sample size (N).

The lower the confidence level or the higher N, the narrower (more precise) the confidence interval.


All three factors influence statistical power.

Therefore, power depends on the effect-size we wish to detect and the precision of our confidence intervals (which in turn depend on α and N).

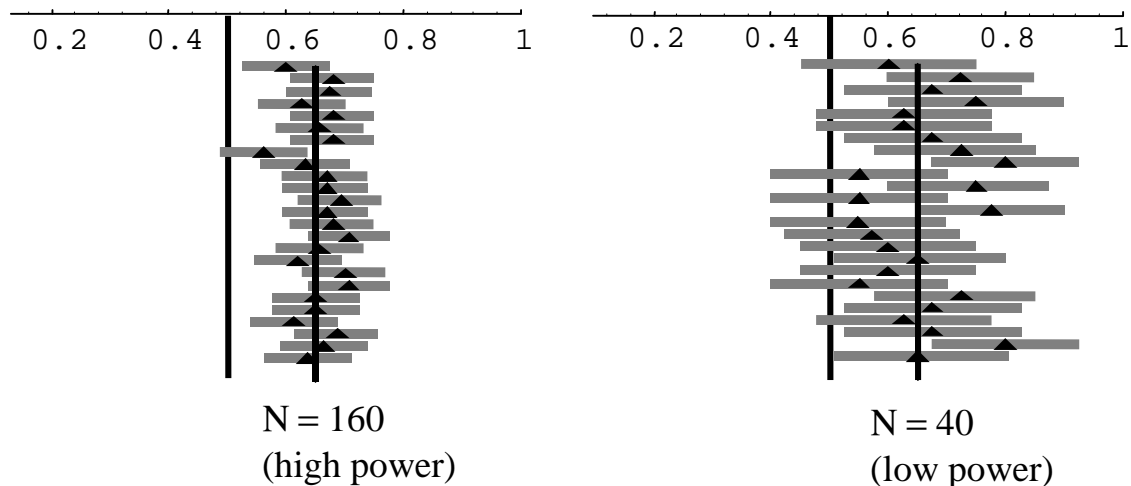
Example 3.

Suppose our $\Pi_h = 0.5$.

For sample sizes of 40, when $\Pi = .75$ rather than $.65$, power increases from $.5$ to $.9279$, meaning that we should correctly reject the null hypothesis of $\mu = .5$ about 93% of the time instead of just 50%.



Likewise for samples of size $N = 160$ our 95% confidence intervals would be about half as wide as those when $N = 40$. They would be far less likely to contain $.5$, as can be seen from a comparison between the two sets of 25 experiments in the figure below.



Sampling Distributions when the Null Hypothesis is False

Noncentral Distributions

- The most widely-used confidence intervals are based on "central" distributions such as the t, F, and χ^2 which would reflect reality if a null hypothesis were true.
- Many of our most useful statistics (e.g., Cohen's d and multiple R^2), however, are based on "noncentral" t, F, and χ^2 distributions that reflect effects of independent variables and/or deviations when the null hypothesis is false.

The noncentral t, F, and χ^2 distributions are generalizations of their usual (central) counterparts with an additional parameter known as a **noncentrality parameter**. The more false the null hypothesis, the larger the noncentrality parameter. The noncentrality parameter is therefore related to effect size and power analysis (Cohen's 1988 book uses them).

Noncentral χ^2 Distribution

Suppose we have a collection of K normally distributed random variables $\{X_i, i = 1, \dots, K\}$ whose means are all 0 and whose variances are all σ^2 . Then the sum of squares of these variables divided by σ^2 (i.e., $\sum X_i^2 / \sigma^2$) has a **central χ^2 distribution** with $K - 1$ degrees of freedom. This χ^2 variable will have a mean of $K - 1$.

But what if the means of the K random variables are not all 0?

Suppose instead that each X_i has its own mean μ_i .

Then it turns out that $\sum X_i^2 / \sigma^2$ has a **noncentral χ^2 distribution** with $r = K - 1$ degrees of freedom and a noncentrality parameter $\delta = \sum \mu_i^2 / \sigma^2$.

This χ^2 variable will have a mean of $K - 1 + \delta$.

One useful interpretation of δ is that it is the sum of squared standardized effects or deviations from 0. This interpretation may be directly applied to a fixed-effects ANOVA model and chi-square statistics for contingency table analysis.

Noncentral F Distribution

The **central F distribution** is the ratio of two central χ^2 variables.

A test for the equality of two variances, for instance, uses the F-ratio conventionally defined by $F = s_1^2/s_2^2$. What's missing from this definition is the fact that the ratio really uses the χ^2 variables s_1^2/σ_1^2 and s_2^2/σ_2^2 , with the null-hypothesis assumption that $\sigma_1^2 = \sigma_2^2$ so that they cancel each other out in the F-ratio.

If we reintroduce σ_1^2 and σ_2^2 then we may set up a noncentral F distribution in which $\sigma_1^2 \neq \sigma_2^2$. A simple way to do this is to set $\sigma_2^2 = 1$ so that σ_1^2 is considered as a multiple of σ_2^2 .

Thus, we put $F = (s_1^2/\sigma_1^2)/s_2^2$.

Then our **noncentral F distribution** is the ratio of a noncentral $\chi^2(df_1, \delta)$ variable and a central $\chi^2(df_2)$ variable.

It takes only one more step to apply this setup to a fixed-effects ANOVA. Consider a one-way experimental design with K conditions and n observations in each condition.

Let α_i be the treatment effect associated with the i^{th} condition, and σ^2 be the error variance.

Then the F statistic with $K - 1$ and $K(n - 1)$ degrees of freedom has a noncentrality parameter

$$\delta = n \sum \alpha_i^2 / \sigma^2,$$

so that δ/n is the sum of squared standardized effects.

A confidence interval for δ also gives us a confidence interval for the sum of those effects.

Noncentral t Distribution

The **central t distribution** is a simple function of the ratio of two random variables (say, W and V).

W has a normal distribution with a mean of 0 and standard deviation of 1.

V has a central $\chi^2(r)$ distribution. Then if we put

$$T = W/(V/r)^{1/2},$$

T has a central t distribution with r degrees of freedom.

In its most common application T is the sample mean divided by the sample standard error. If we rewrite the one-sample t statistic as

$$T = \frac{(\bar{X} - \mu)/(\sigma/\sqrt{N})}{s/\sigma}$$

then we can see that W is the numerator of T and

$$V = SS/\sigma^2$$

where SS is the sums of squares for X.

The assumption that gives this t statistic a central t distribution is that W has a mean of 0 (i.e., the null hypothesis is true).

If W has a mean of $E/(\sigma/\sqrt{N}) \neq 0$, then T has a **noncentral t distribution** with a noncentrality parameter

$$\delta = E/(\sigma/\sqrt{N}),$$

which is directly related to the standardized effect size.

Reminder: Cohen's d is the difference between the sample mean and the population mean divided by the sample standard deviation:

$$d = \frac{\bar{X} - \mu}{s}$$

Cohen's d measures the difference between the sample mean and the population mean in standard deviation units so it is scale-free. For the one-sample case,

$$t = d\sqrt{N}.$$

So,

$$d = \delta/\sqrt{N}.$$

This interpretation of δ may be applied not only to confidence intervals for Cohen's d but also for any standardized planned contrast.

Noncentral Confidence Intervals: The Basics

The Inversion Confidence Interval Principle

Briefly, the main idea is to use the observed value of a test statistic (e.g., an observed t-value) to initiate the search for the lower and upper limits to a $1 - \alpha$ confidence interval for the noncentrality parameter, which is then converted into a confidence interval for an effect size. That is, we move through three stages:

test statistic \rightarrow noncentrality parameter \rightarrow effect statistic

Example 4.

A fictitious social facilitation study from Smithson (2000: 194-196) involves measurements of twelve 800-metre runners' times under two conditions (no one present vs one person present). The mean difference is 0.983 seconds faster for the someone-present condition, and the sample standard deviation is 1.284.

Step 1.

Using the conventional $\alpha = .05$, we would ordinarily calculate

$$t(11) = 0.983 / (1.284 / \sqrt{12}) = 2.65,$$

ascertain that it is larger than $t_{\alpha/2} = 2.201$, and thereby conclude that we could reject the null hypothesis.


We could construct a 95% confidence interval (CI) around the mean difference. However, if we wanted to compare this social facilitation effect with that for some other event (a foot-race over another distance, or swimming or cycling) then we would also want a 'dimensionless' effect-size measure such as Cohen's d:

$$d = \frac{\bar{X} - \mu}{s} = 0.983 / 1.284 = 0.766.$$

To construct a 95% CI for Cohen's d requires us to use the noncentral t distribution.

Step 2.

We already know the observed $t(11) = 2.65$. This also gives us a sample estimate of the noncentrality parameter δ . At this point, we require a computer routine to home in on a lower limit δ_L , corresponding to the noncentral t distribution for which $t(11) = 2.65$ cuts off $\alpha/2$ of its upper tail, and an upper limit for δ_U corresponding to the noncentral t distribution for which 2.65 cuts off $\alpha/2$ of its lower tail.



- For a noncentral t-distribution with $\delta = 0.3604$ we would observe a sample $t(11) > 2.65$ just 2.5% of the time; and
- For a noncentral t-distribution with $\delta = 4.8494$ we would observe a sample $t(11) < 2.65$ just 2.5% of the time.

Step 3.

We already know that

$$d = \delta/\sqrt{N},$$

so we immediately have a 95% CI for d , namely

$$[0.3604/\sqrt{12}, 4.8494/\sqrt{12}] = [0.1040, 1.3999].$$

According to Cohen's (1988) effect-size guidelines,

$d \leq 0.2$ is "small,"

$0.2 < d < 0.5$ is "medium," and

$d \geq 0.8$ is "large."

Although $d = 0.766$ is a moderately large effect, our CI tells us that the true effect-size could plausibly range from quite small (0.1040) to very large (1.3999). That's because our sample size is small ($N = 12$).

Example 5: Let's Do It!

Suppose you are reading an article that claims to have found a way of increasing adult mathematical aptitude. They use a mathematics test that is normed for the population from which their random sample came and find that the mean for their sample is significantly higher than the population mean ($t(35) = 3.6$, $p = .005$).

On the basis of just this information, we can answer these questions:

1. How big is the effect in terms of Cohen's d ?
2. What is a 99% CI for Cohen's d ?
3. What is the power of this study to detect an effect size of $d = 0.2$, 0.5 , or 0.8 when $\alpha = .01$?

1. First, convert $t(35)$ to d by realising that since $df = 35$, $N = 36$ and so $d = t/\sqrt{N} = 3.6/6 = 0.6$.
about medium-sized according to Cohen's criteria.

2. Now, launch **NoncT.sav**, **NoncT2.sps**, and **T2D.sps** from the **SPSS** folder in the **Workshop** folder. Bring the NoncT.sav data editor forward as the active window.

- In the **tval** column enter 3.6 as the value of the observed t statistic.
- In the **df** column enter 35 as the df .
- In the **conf** column enter 0.99 as the confidence level.
- Now make the **NoncT2.sps** window the active window.
- Under the **Run** menu select **All** to run the syntax.

You should find that in the **lc2** column you have .7770 and in the **uc2** column you have 6.3844. These are the lower and upper limits of the 99% CI for δ , the noncentrality parameter.

You may convert these limits to a 99% CI for d by dividing each of them by $\sqrt{N} = 6$, to get [0.1295, 1.0641].

To get SPSS to do this for you, in the **T2D.sps** window select the syntax
`COMPUTE LOWD = LC2/SQRT(DF+1) .`
`COMPUTE HIGHD = UC2/SQRT(DF+1) .`
`EXECUTE.`

and run that. You should find two new variables (**lowd** and **highd**) containing the appropriate CI limits.

3. Note that in the **power** column it says .8029, meaning that this study has power of .8029 to detect an effect-size of 0.6 when $\alpha = .01$.

To find out the power to detect an effect-size of $d = 0.2, 0.5, \text{ or } 0.8$ under this condition, follow these steps:

- Create three new rows of data in the NoncT.sav window, all with **df** = 35 and **conf** = 0.99. Leave **tval** blank.
- Create a new variable called **d**. Type in the values 0.2, 0.5, and 0.8 for your three new rows.
- Convert these **d** values to ones for **tval** by selecting the following syntax in the **T2D.sps** window and running it:

```
COMPUTE TVAL = D*SQRT(DF+1) .
EXECUTE.
```

You should see tvals of 1.2, 3.0, and 4.8.
- Now run the **NoncT2.sps** syntax. You should find that the power column contains these values:
.0761, .6109, and .9768.

In short, the power of this study to detect small to medium effects is not very high.

And something extra!

You can also find out what the 99% CIs would be like if the observed effect sizes were 0.2, 0.5, or 0.8. To do this:

- In the **T2D.sps** window select the syntax
- ```
COMPUTE LOWD = LC2/SQRT(DF+1) .
COMPUTE HIGHD = UC2/SQRT(DF+1) .
EXECUTE.
```
- and run that.

You should find lowd and highd now display the appropriate CI limits for  $d$ . You can see that for  $d = 0.2$  you would not be able to reject the null hypothesis, and you can barely do so for  $d = 0.6$ .

Notice also that although power increases as we go from  $d = 0.2$  to 0.8, the width of the CIs does not decrease (Why not?? Hint: What increases power but doesn't make CIs narrower?).

## Using Power and Precision to Design Studies

### Relationship between Noncentral CIs and Power Analysis

Exploring the effect of sample size and effect size on the width and location of CIs may be used to good effect in designing studies. Since this is usually a role played by power analysis, how are noncentral CI widths and power related?

All else being equal, the narrower the CI and/or the greater the effect size, the greater the power. Narrower CIs are obtained with large N and/or low confidence.

Power analysis involves the same sampling distributions as noncentral CIs. Nonetheless, noncentral CIs provide different information from power analysis; high power does not entail precise estimation nor vice-versa.

### Example 6: Let's Do It!

Suppose we have two criteria for a study using a one-sample t-test:

1. The study should detect  $d \geq 0.3$  when  $\alpha = .01$  with power of at least 0.9, and
2. The 99% CI for an observed  $d$  near 0.3 should be no wider than 0.30.

First, create a table of candidate sample sizes:

- In the **NoncT2.sav** data-editor, create six rows of data.
- In the **conf** column enter 0.99 for all rows.
- In the **df** column enter 49, 99, 149, 199, 249, and 299 respectively.
- Enter values for **d** of 0.3 for all six rows.

Then...

- Convert the **d** values to **tval** by selecting the syntax in the **T2D.sps** window and running it:

```
COMPUTE TVAL = D*SQRT(DF+1) .
EXECUTE.
```

You should see **tvals** ranging from 2.1213 to 5.1962.

- Now run the **NoncT2.sps** syntax. You should find in the power column that for  $df = 149$  power is .8548. So a sample size of about 150 would fulfil the first criterion.
  - In the **T2D.sps** window select the syntax
- ```
COMPUTE LOWD = LC2/SQRT(DF+1) .
COMPUTE HIGHD = UC2/SQRT(DF+1) .
EXECUTE.
```

and run that. You should find that the width of the 99% CI for d does not narrow down to 0.30 until $df = 299$ (i.e., $N = 300$). The second criterion is clearly much more stringent than the first.

Noncentral CIs for ANOVA and Regression

Fixed-Effects Factorial ANOVA

Confidence intervals for effects in a fixed-effects ANOVA are based on the noncentral F distribution. When the null hypothesis is false, $F(df1,df2)$ has a noncentral F distribution with a noncentrality parameter δ that may be estimated by

$$\delta = (df1/df2)(F)(df1 + df2 + 1).$$

The effect-size, measured by partial η^2 , is related to δ and F by:
 $\eta^2 = \delta/(\delta + df1 + df2 + 1) = (df1)F/[df2 + (df1)F].$

Example 7: Let's Do It!

Suppose we have a 2x7 factorial ANOVA with 4 observations per cell.

- Row effect $F(1,42) = 6.00$ ($p = .0186$),
- Column effect $F(6,42) = 2.65$ ($p = .0285$), and
- Interaction effect $F(6,42) = 2.50$ ($p = .0369$).

The researchers could be forgiven for supposing that the row effect was the "most significant," followed by the column effect.

But partial η^2 values tell a different story. From the formula above we find

- Row effect $\eta_{\alpha}^2 = .1250$,
- Column effect $\eta_{\beta}^2 = .2746$, and
- Interaction effect $\eta_{\gamma}^2 = .2632$.

The row effect is, in fact, the smallest effect.

Now suppose we construct 90% CIs for these η^2 values.

First, launch **NoncF.sav**, **NoncF3.sps**, and **F2R2.sps** from the **SPSS** folder in the **Workshop** folder. Bring the NoncF.sav data editor forward as the active window.

- Create three rows of data.
- In the **fval** column enter the F-values 6.00, 2.65, and 2.50.
- In the **df1** column enter 1, 6, and 6.
- In the **df2** column enter 42 for all three rows.
- In the **conf** column enter 0.90 as the confidence level.
- Now make the **NoncF3.sps** window the active window.
- Under the **Run** menu select **All** to run the syntax.

You should see values in the lc2 and uc2 columns corresponding to the δ intervals and in the lr2 and ur2 columns corresponding to the η^2 intervals shown below:

Row effect	δ_α interval = [0.5208, 17.1211],	η_α^2 interval = [.0117, .2801];
Column effect	δ_β interval = [0.8699, 27.2893],	η_β^2 interval = [.0174, .3577];
Interaction effect	δ_γ interval = [0.4621, 25.8691],	η_γ^2 interval = [.0093, .3455].

The CI is narrowest for the row effect because the row effect has $df1 = 1$ and $df2 = 42$, whereas the column effect and interaction effects both have $df1 = 6$ and $df2 = 42$. This is what is driving the differences in the significance-levels of the F tests.

Power and One-Sided vs Two-Sided CIs

The usual F-test is a one-tailed test. Power for the F-test is therefore associated with just one tail, and also with a one-sided CI.

In **NoncF3.sps**, power is computed using $\alpha = (1 - CONF) / 2$.

For instance, setting $CONF = .90$ gives us power for $\alpha = .05$ and a 95% one-tailed CI as well as a 90% two-tailed CI.

Example 8: Let's Do It!

Let's compare the power of our 2x7 factorial design to detect $\eta^2 = .125$ for the row versus the column effect, when $\alpha = .05$.

- In the second row of **NoncF.sav**, change **r2** value to .1250 (the same as the value in the first row).
- In **F2r2.sps**, select the following syntax:

```
COMPUTE   FVAL = (R2/(1-R2))*(DF2/DF1) .
EXECUTE .
```

and run it. You should see **fval** for the second row change to 1.000.
- Run the syntax in the **NoncF3.sps** window. You should find power for the first row is .6878, but for the second row is only .4068.
- As a bonus, you'll find the lower limit of a one-sided 95% CI for η^2 in the lr2 column-- .0117 for the row effect but 0 for the column effect.

Regression

Constructing CIs or conducting power analysis for R^2 requires that we stipulate whether our regression setup involves fixed-score or random-score regressors. . Fixed scores are preplanned or given whereas random scores arise from random sampling.

- Fixed scores: Noncentral F is the sampling distribution of the F-statistic and we therefore get R^2 via the Inversion Confidence Interval Principle as in ANOVA.
- Random scores: The noncentral F distribution can no longer be used here. In this case we have Fisher's (1928) distribution for R^2 , whence we use a program like Steiger & Fouladi's R2 for constructing our CIs.

Power will be higher and the CIs narrower under fixed-score regression simply because there are fewer sources of sampling error.

Fixed-score Regression:

Suppose we have two sets of predictors, A and B, where

w = # of predictors in set A

u = # of predictors in set B.

Define $df1 = u$ and $df2 = N - u - w - 1$.

Then the F-statistic that tests the significance of the contribution of set B after set A has already been included in the model is related to its noncentrality parameter by

$$\delta = (df1/df2)(F)(df1 + df2 + 1).$$

The effect-size, measured by partial R^2 , is related to δ and F by:

$$R^2 = \delta / (\delta + df1 + df2 + 1) = (df1)F / [df2 + (df1)F].$$

If $w = 0$ then we get a CI for multiple R^2 for a model with u predictors in it.

Example 9: Let's Do It!

Tabachnick and Fidell's (1996: 174-193) regression example has number of visits to health professionals (timedrs) being predicted by three independent variables.

The sample size is $N = 465$, so $df1 = 3$ and $df2 = 461$.

Their SPSS output tells us that $F(3,461) = 92.901$ and multiple $R^2 = .3768$.

Let's construct the 90% CI around R^2 assuming a fixed-score regression model.

- Pick a row in the **NoncF.sav** data editor, and in the **fval** column enter the F-value 92.901.
- In the **df1** column enter 3.
- In the **df2** column enter 461.
- In the **conf** column enter 0.90 as the confidence level.
- Run the **NoncF3.sps** syntax.

You should get a 90% CI for the noncentrality parameter of [217.3965, 342.5157], which translates to a CI for R^2 of [.3186, .4242].

R2, on the other hand, is what we use for the random-score model. It gives us a wider 90% CI of [.31462, .43165].

- To run R2, open the R2 folder and double-click on the **R2** icon.
- Hit any key, and the ensuing dialog box will tell you how to use the arrow keys to get around in the menu.
- Once you are in the **Confidence Interval** item under the **Options** menu, you will need to enter number of observations (465), number of variables (4), R-squared (.3768), and Confidence Level (.9).
- When you select **Go** and hit Enter, you'll need to hit another key to display the results. The Lower Limit and Upper Limit are the limits for the two-sided CI; the Lower Bound of .32767 is the lower limit on a one-sided CI.
- To exit R2, arrow to the **File** menu, hit Enter, arrow down to the **Exit** option, and hit Enter.