

esize — Effect size based on mean comparison

| | | | |
|-----------------------------|--------------------------------------|--------------------------------|--------------------------------------|
| Description | Quick start | Menu | Syntax |
| Options | Remarks and examples | Stored results | Methods and formulas |
| References | Also see | | |

Description

`esize` calculates effect sizes for comparing the difference between the means of a continuous variable for two groups. In the first form, `esize` calculates effect sizes for the difference between the mean of *varname* for two groups defined by *groupvar*. In the second form, `esize` calculates effect sizes for the difference between *varname*₁ and *varname*₂, assuming unpaired data.

`esizei` is the immediate form of `esize`; see [\[U\] 19 Immediate commands](#). In the first form, `esizei` calculates the effect size for comparing the difference between the means of two groups. In the second form, `esizei` calculates the effect size for an *F* test after an ANOVA.

Quick start

Cohen's *d* and Hedges's *g* comparing the difference in means of *v* for two independent groups in *catvar*

```
esize twosample v, by(catvar)
```

As above, but with group data stored in *v1* and *v2*

```
esize unpaired v1==v2
```

As above, but use 90% confidence level

```
esize unpaired v1==v2, level(90)
```

Cohen's *d* and Hedges's *g* for means of *v* for groups in *catvar1* calculated over each level of *catvar2*

```
by catvar2: esize twosample v, by(catvar1)
```

Menu

`esize`

Statistics > Summaries, tables, and tests > Classical tests of hypotheses > Effect size based on mean comparison

`esizei`

Statistics > Summaries, tables, and tests > Classical tests of hypotheses > Effect-size calculator

Syntax

Effect sizes for two independent samples using groups

```
esize twosample varname [if] [in], by(groupvar) [options]
```

Effect sizes for two independent samples using variables

```
esize unpaired varname1 == varname2 [if] [in], [options]
```

Immediate form of effect sizes for two independent samples

```
esizei #_obs1 #_mean1 #_sd1 #_obs2 #_mean2 #_sd2 [, options]
```

Immediate form of effect sizes for *F* tests after an ANOVA

```
esizei #_df1 #_df2 #_F [, level(#)]
```

options

Description

Main

| | |
|-------------------------|--|
| <code>cohensd</code> | report Cohen's <i>d</i> (1988) |
| <code>hedgesg</code> | report Hedges's <i>g</i> (1981) |
| <code>glassdelta</code> | report Glass's Δ (Smith and Glass 1977) using each group's standard deviation |
| <code>pbcorr</code> | report the point-biserial correlation coefficient (Pearson 1909) |
| <code>all</code> | report all estimates of effect size |
| <code>unequal</code> | use unequal variances |
| <code>welch</code> | use Welch's (1947) approximation |
| <code>level(#)</code> | set confidence level; default is <code>level(95)</code> |

`by` is allowed with `esize`; see [D] [by](#).

Options

Main

`by(groupvar)` specifies the *groupvar* that defines the two groups that `esize` will use to estimate the effect sizes. Do not confuse the `by()` option with the `by` prefix; you can specify both.

`cohensd` specifies that Cohen's *d* (1988) be reported.

`hedgesg` specifies that Hedges's *g* (1981) be reported.

`glassdelta` specifies that Glass's Δ (Smith and Glass 1977) be reported.

`pbcorr` specifies that the point-biserial correlation coefficient (Pearson 1909) be reported.

`all` specifies that all estimates of effect size be reported. The default is Cohen's *d* and Hedges's *g*.

`unequal` specifies that the data not be assumed to have equal variances.

`welch` specifies that the approximate degrees of freedom for the test be obtained from Welch's formula (1947) rather than from Satterthwaite's approximation formula (1946), which is the default when `unequal` is specified. Specifying `welch` implies `unequal`.

`level(#)` specifies the confidence level, as a percentage, for confidence intervals. The default is `level(95)` or as set by `set level`; see [U] [20.8 Specifying the width of confidence intervals](#).

Remarks and examples

stata.com

Remarks are presented under the following headings:

[Introduction](#)
[Estimating effect sizes](#)
[Immediate form](#)
[Video example](#)

Introduction

Whereas p -values are used to assess the statistical significance of a result, measures of effect size are used to assess the practical significance of a result. Effect sizes can be broadly categorized as “measures of group differences” (the d family) and “measures of association” (the r family); see [Ellis \(2010, table 1.1\)](#). The d family includes estimators such as Cohen’s d , Hedges’s g , and Glass’s Δ . The r family includes estimators such as the point-biserial correlation coefficient, η^2 , ε^2 , and ω^2 (also see `estat esize` in [R] [regress postestimation](#)). For an introduction to the concepts and calculation of effect sizes, see [Kline \(2013\)](#) and [Thompson \(2006\)](#). For a more detailed discussion, see [Kirk \(1996\)](#), [Ellis \(2010\)](#), [Cumming \(2012\)](#), [Grissom and Kim \(2012\)](#), and [Kelley and Preacher \(2012\)](#).

It should be noted that there is much variation in the definitions of measures of effect size ([Kline 2013](#)). As [Ellis \(2010, 27\)](#) cautions, “However, beware the inconsistent terminology. What is labeled here as g was labeled by Hedges and Olkin as d and vice versa. For these authors writing in the early 1980s, g was the mainstream effect-size index developed by Cohen and refined by Glass (hence g for Glass). However, since then g has become synonymous with Hedges’s equation (not Glass’s) and the reason it is called Hedges’s g and not Hedges’s h is because it was originally named after Glass—even though it was developed by Larry Hedges. Confused?”

To avoid confusion, `esize` and `esizei` closely follow the notation of [Hedges \(1981\)](#), [Smithson \(2001\)](#), [Kline \(2013\)](#), and [Ellis \(2010\)](#).

Estimating effect sizes

▷ Example 1: Effect size for two independent samples using by()

Suppose we are interested in question 1 from the fictitious `depression.dta`: “My statistical software makes me feel sad”. We might have conducted a t test to test the null hypothesis that there is no difference in response by sex. We could then compute various measures of effect size to describe the magnitude of the effect of sex.

```
. use http://www.stata-press.com/data/r15/depression
(Fictitious Depression Inventory data based on the Beck Depression Inventory)
. esize twosample qu1, by(sex) all
Effect size based on mean comparison
```

```
Obs per group:
Female =      712
Male =      288
```

| Effect Size | Estimate | [95% Conf. Interval] | |
|--------------------|-----------|----------------------|----------|
| Cohen's d | -.0512417 | -.1881184 | .0856607 |
| Hedges's g | -.0512032 | -.187977 | .0855963 |
| Glass's Delta 1 | -.0517793 | -.1886587 | .0851364 |
| Glass's Delta 2 | -.0499786 | -.1868673 | .086997 |
| Point-Biserial r | -.0232208 | -.0849629 | .0387995 |

Cohen's d , Hedges's g , and both estimates of Glass's Δ indicate that the score for females is 0.05 standard deviations lower than the score for males. The point-biserial correlation coefficient indicates that there is a small, negative correlation between the scores for females and males.

◀

□ Technical note

Glass's Δ has traditionally been estimated for experimental studies using the control group standard deviation rather than the pooled standard deviation. [Kline \(2013\)](#) notes that the choice of group becomes arbitrary for data arising from observational studies and recommends the reporting of Glass's Δ using each group standard deviation.

□

▷ Example 2: Effect size for two independent samples by a third variable

If we are interested in the same effect sizes from [example 1](#) stratified by race, we could use the `by` prefix with the `sort` option to accomplish this task.

```
. by race, sort: esize twosample qul, by(sex)
```

```
-> race = Hispanic
```

```
Effect size based on mean comparison
```

```
Obs per group:
Female =      88
Male =       45
```

| Effect Size | Estimate | [95% Conf. Interval] | |
|-------------------|-----------|----------------------|----------|
| Cohen's <i>d</i> | -.1042883 | -.463503 | .2553235 |
| Hedges's <i>g</i> | -.1036899 | -.4608434 | .2538584 |

```
-> race = Black
```

```
Effect size based on mean comparison
```

```
Obs per group:
Female =     259
Male =      95
```

| Effect Size | Estimate | [95% Conf. Interval] | |
|-------------------|-----------|----------------------|----------|
| Cohen's <i>d</i> | -.1720681 | -.4073814 | .063489 |
| Hedges's <i>g</i> | -.1717011 | -.4065127 | .0633536 |

```
-> race = White
```

```
Effect size based on mean comparison
```

```
Obs per group:
Female =     365
Male =     148
```

| Effect Size | Estimate | [95% Conf. Interval] | |
|-------------------|----------|----------------------|----------|
| Cohen's <i>d</i> | .0479511 | -.1430932 | .2389486 |
| Hedges's <i>g</i> | .0478807 | -.1428831 | .2385977 |

▷ Example 3: Bootstrap confidence intervals for effect sizes

Simulation studies have shown that bootstrap confidence intervals may be preferable to confidence intervals based on the noncentral t distribution when the variable of interest does not have a normal distribution (Kelley 2005; Algina, Keselman, and Penfield 2006). Bootstrap confidence intervals can be easily estimated for effect sizes using the bootstrap prefix.

```
. use http://www.stata-press.com/data/r15/depression
(Fictitious Depression Inventory data based on the Beck Depression Inventory)
. set seed 12345
. bootstrap r(d) r(g), reps(1000) nodots nowarn: esize twosample qu1, by(sex)
Bootstrap results                                Number of obs   =       1,000
                                                Replications    =       1,000

command: esize twosample qu1, by(sex)
   _bs_1: r(d)
   _bs_2: r(g)
```

| | Observed Coef. | Bootstrap Std. Err. | z | P> z | Normal-based [95% Conf. Interval] | |
|-------|-------------------|------------------------|-------|-------|--------------------------------------|----------|
| _bs_1 | -.0512417 | .0742692 | -0.69 | 0.490 | -.1968066 | .0943233 |
| _bs_2 | -.0512032 | .0742134 | -0.69 | 0.490 | -.1966587 | .0942523 |

◀

▷ Example 4: Effect sizes for two independent samples using variables

Sometimes, the data of interest are stored in two separate variables. We can calculate effect sizes for the two groups by using the unpaired version of `esize`.

```
. use http://www.stata-press.com/data/r15/fuel
. esize unpaired mpg1==mpg2
Effect size based on mean comparison

                                Number of obs =          24

Effect Size | Estimate | [95% Conf. Interval]
-----|-----|-----
Cohen's d  | -.5829654 | -1.394934 .2416105
Hedges's g | -.5628243 | -1.34674 .2332631
```

◀

Immediate form

► Example 5: Immediate form for effect sizes for two means

Often we do not have access to raw data, but we are given summary statistics in a report or manuscript. To calculate the effect sizes from summary statistics, we can use the immediate command `esizei`. For example, [Kline \(2013\)](#) in table 4.2 shows summary statistics for a hypothetical sample where $\text{mean}_1 = 13$, $\text{sd}_1 = 2.74$, $\text{mean}_2 = 11$, and $\text{sd}_2 = 2.24$; there are 30 people in each group. We can estimate the effect sizes from these summary data using `esizei`:

```
. esizei 30 13 2.74 30 11 2.24
```

Effect size based on mean comparison

```
Obs per group:
Group 1 =      30
Group 2 =      30
```

| Effect Size | Estimate | [95% Conf. Interval] | |
|-------------------|----------|----------------------|----------|
| Cohen's <i>d</i> | .7991948 | .2695509 | 1.322465 |
| Hedges's <i>g</i> | .7888081 | .2660477 | 1.305277 |

◀

► Example 6: Immediate form for effect sizes for F tests after an ANOVA

`esizei` can also be used to compute η^2 , ε^2 , and ω^2 for F tests after an ANOVA. The following example from [Smithson \(2001, 623\)](#) illustrates the use of `esizei` for $\text{df}_{\text{num}} = 4$, $\text{df}_{\text{den}} = 50$, and $F = 4.2317$.

```
. esizei 4 50 4.2317, level(90)
```

Effect sizes for linear models

| Effect Size | Estimate | [90% Conf. Interval] | |
|-----------------|----------|----------------------|----------|
| Eta-Squared | .2529151 | .0521585 | .3603621 |
| Epsilon-Squared | .1931483 | | |
| Omega-Squared | .1903049 | | |

◀

Video example

[Tour of effect sizes](#)

Stored results

`esize` and `esizei` for comparing two means store the following in `r()`:

Scalars

| | |
|---------------------------|---|
| <code>r(d)</code> | Cohen's d |
| <code>r(lb_d)</code> | lower confidence bound for Cohen's d |
| <code>r(ub_d)</code> | upper confidence bound for Cohen's d |
| <code>r(g)</code> | Hedges's g |
| <code>r(lb_g)</code> | lower confidence bound for Hedges's g |
| <code>r(ub_g)</code> | upper confidence bound for Hedges's g |
| <code>r(delta1)</code> | Glass's Δ for group 1 |
| <code>r(lb_delta1)</code> | lower confidence bound for Glass's Δ for group 1 |
| <code>r(ub_delta1)</code> | upper confidence bound for Glass's Δ for group 1 |
| <code>r(delta2)</code> | Glass's Δ for group 2 |
| <code>r(lb_delta2)</code> | lower confidence bound for Glass's Δ for group 2 |
| <code>r(ub_delta2)</code> | upper confidence bound for Glass's Δ for group 2 |
| <code>r(r_pb)</code> | point-biserial correlation coefficient |
| <code>r(lb_r_pb)</code> | lower confidence bound for the point-biserial correlation coefficient |
| <code>r(ub_r_pb)</code> | upper confidence bound for the point-biserial correlation coefficient |
| <code>r(N_1)</code> | sample size n_1 |
| <code>r(N_2)</code> | sample size n_2 |
| <code>r(df_t)</code> | degrees of freedom |
| <code>r(level)</code> | confidence level |

`esizei` for F tests after ANOVA stores the following in `r()`:

Scalars

| | |
|--------------------------|-------------------------------------|
| <code>r(eta2)</code> | η^2 |
| <code>r(lb_eta2)</code> | lower confidence bound for η^2 |
| <code>r(ub_eta2)</code> | upper confidence bound for η^2 |
| <code>r(epsilon2)</code> | ε^2 |
| <code>r(omega2)</code> | ω^2 |
| <code>r(level)</code> | confidence level |

Methods and formulas

For the d family, the effect-size parameter of interest is the scaled difference between the means given by

$$\delta = \frac{(\mu_1 - \mu_2)}{\sigma}$$

One of the most popular estimators of effect size is Cohen's d , given by

$$\text{Cohen's } d = \frac{(\bar{x}_1 - \bar{x}_2)}{s^*}$$

where

$$s^* = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}}$$

Hedges (1981) showed that Cohen's d is biased and proposed the unbiased estimator

$$\text{Hedges's } g = \text{Cohen's } d \times c(m)$$

where $m = n_1 + n_2 - 2$ and

$$c(m) = \frac{\Gamma(\frac{m}{2})}{\sqrt{\frac{m}{2}}\Gamma(\frac{m-1}{2})}$$

Glass (Smith and Glass 1977) proposed an estimator for δ in the context of designed experiments,

$$\text{Glass's } \Delta = \frac{(\bar{x}_{\text{treated}} - \bar{x}_{\text{control}})}{s_{\text{control}}}$$

where s_{control} is the standard deviation for the control group.

As noted above, `esize` and `esizei` report two estimates of Glass's Δ : one using the standard deviation for group 1 and the other using the standard deviation for group 2:

$$\text{Glass's } \Delta_1 = \frac{(\bar{x}_1 - \bar{x}_2)}{s_1}$$

and

$$\text{Glass's } \Delta_2 = \frac{(\bar{x}_1 - \bar{x}_2)}{s_2}$$

For the r family, the effect-size parameter of interest is the ratio of the variance attributable to an effect and the total variance:

$$\eta^2 = \frac{\sigma_{\text{effect}}^2}{\sigma_{\text{total}}^2}$$

A popular estimator of η when there are two groups is the point-biserial correlation coefficient,

$$r_{\text{PB}} = \frac{t}{\sqrt{t^2 + \text{df}}}$$

where t is the t statistic for the difference between the means of the two groups, and df is the corresponding degrees of freedom. Satterthwaite's or Welch's adjustment (see [R] `ttest` for details) to the degrees of freedom can be used to calculate r_{PB} by specifying the `unequal` or `welch` option, respectively.

When more than two means are being compared, as in the case of an ANOVA with p groups, a popular estimator of effect size is the correlation ratio denoted η^2 (Fisher 1925; Kerlinger 1964). η^2 can be computed directly as the ratio of the $\text{SS}_{\text{effect}}$ and the SS_{total} or as a function of the F statistic with numerator degrees of freedom equal to df_{num} and denominator degrees of freedom equal to df_{den} .

$$\hat{\eta}^2 = \frac{F}{F + \text{df}_{\text{den}}/\text{df}_{\text{num}}}$$

Like its equivalent estimator R^2 , η^2 has an upward bias. Less biased estimators of effect size are ε^2 and ω^2 (Grissom and Kim 2012).

$$\varepsilon^2 = \frac{F - 1}{F + \text{df}_{\text{den}}/\text{df}_{\text{num}}} = \hat{\eta}^2 - \frac{\text{df}_{\text{num}}}{\text{df}_{\text{den}}}(1 - \hat{\eta}^2)$$

$$\widehat{\omega}^2 = \frac{F - 1}{F + (\text{df}_{\text{den}} + 1)/\text{df}_{\text{num}}}$$

To calculate $\widehat{\eta}^2$, $\widehat{\varepsilon}^2$, and $\widehat{\omega}^2$ directly after `anova` or `regress`, see `estat esize` in [R] `regress postestimation`.

Cohen's d , Hedges's g , and Glass's Δ have been shown to have a noncentral t distribution (Hedges 1981) with noncentrality parameter equal to

$$\lambda = \delta \sqrt{\frac{n_1 n_2}{n_1 + n_2}}$$

Confidence intervals are calculated by finding the noncentrality parameters λ_{lower} and λ_{upper} that correspond to

$$\Pr(\text{df}, \delta, \lambda_{\text{lower}}) = 1 - \frac{\alpha}{2}$$

and

$$\Pr(\text{df}, \delta, \lambda_{\text{upper}}) = \frac{\alpha}{2}$$

using the function `npnt(df, t, p)`. The noncentrality parameters are then transformed back to the effect-size scale:

$$\delta_{\text{lower}} = \lambda_{\text{lower}} \sqrt{\frac{n_1 + n_2}{n_1 n_2}}$$

and

$$\delta_{\text{upper}} = \lambda_{\text{upper}} \sqrt{\frac{n_1 + n_2}{n_1 n_2}}$$

(see Venables [1975]; Steiger and Fouladi [1997]; Cumming and Finch [2001]; Smithson [2001]).

Confidence intervals for the point-biserial correlation coefficient are calculated similarly and transformed back to the effect-size scale as

$$r_{\text{lower}} = \frac{\lambda_{\text{lower}}}{\sqrt{\lambda_{\text{lower}}^2 + \text{df}}}$$

and

$$r_{\text{upper}} = \frac{\lambda_{\text{upper}}}{\sqrt{\lambda_{\text{upper}}^2 + \text{df}}}$$

Following Smithson's (2001) notation, the F statistic is written as

$$F_{\text{df}_{\text{num}}, \text{df}_{\text{den}}} = f^2 (\text{df}_{\text{num}} / \text{df}_{\text{den}})$$

This equation has a noncentral F distribution with noncentrality parameter:

$$\lambda = f^2 (\text{df}_{\text{num}} + \text{df}_{\text{den}} + 1)$$

where $f^2 = \eta^2 / (1 - \eta^2)$.

Confidence intervals for $\hat{\eta}^2$ are calculated by finding the noncentrality parameters λ_{lower} and λ_{upper} for a noncentral F distribution that correspond to

$$\Pr(\text{df}_{\text{num}}, \text{df}_{\text{den}}, F, \lambda_{\text{lower}}) = 1 - \frac{\alpha}{2}$$

and

$$\Pr(\text{df}_{\text{num}}, \text{df}_{\text{den}}, F, \lambda_{\text{upper}}) = \frac{\alpha}{2}$$

using the function `npnF(df1, df2, f, p)`. The noncentrality parameters are transformed back to the $\hat{\eta}^2$ scale as

$$\hat{\eta}_{\text{lower}}^2 = \frac{\lambda_{\text{lower}}}{\lambda_{\text{lower}} + \text{df}_{\text{num}} + \text{df}_{\text{den}} + 1}$$

and

$$\hat{\eta}_{\text{upper}}^2 = \frac{\lambda_{\text{upper}}}{\lambda_{\text{upper}} + \text{df}_{\text{num}} + \text{df}_{\text{den}} + 1}$$

While confidence intervals for $\hat{\varepsilon}^2$ can be constructed using the same transformation that links it with $\hat{\eta}^2$, there are several arguments for not using them in practice. See [Smithson \(2003, 54\)](#) for further details.

Fred Nichols Kerlinger (1910–1991) was born in New York City. He studied music at New York University and graduated magna cum laude with a degree in education and philosophy. After graduation, he joined the U.S. Army and served as a counterintelligence officer in Japan in 1946. Kerlinger earned an MA and a PhD in educational psychology from the University of Michigan and held faculty appointments at several universities, including New York University. He was president of the American Educational Research Association and is best known for his popular and influential book *Foundations of Behavioral Research* (1964), which introduced Fisher's (1925) η^2 statistic to behavioral researchers.

William Lee Hays (1926–1995) was born in Clarksville, Texas. He studied mathematics and psychology at Paris Junior College in Paris, Texas, and at East Texas State College. He earned BS and MS degrees from North Texas State University. Upon completion of his PhD in psychology at the University of Michigan, he joined the faculty, where he eventually became associate vice president for academic affairs. In 1977, Hays accepted an appointment as vice president for academic affairs at the University of Texas at Austin, where he remained until his death in 1995. Hays is best known for his book *Statistics for Psychologists* (1963), which introduced the ω^2 statistic (and is actually denoted here by ε^2).

References

- Algina, J., H. J. Keselman, and R. D. Penfield. 2006. Confidence interval coverage for Cohen's effect size statistic. *Educational and Psychological Measurement* 66: 945–960.
- Cohen, J. 1988. *Statistical Power Analysis for the Behavioral Sciences*. 2nd ed. Hillsdale, NJ: Erlbaum.
- Cumming, G. 2012. *Understanding the New Statistics: Effect Sizes, Confidence Intervals, and Meta-Analysis*. New York: Routledge.
- Cumming, G., and S. Finch. 2001. A primer on the understanding, use, and calculation of confidence intervals that are based on central and noncentral distributions. *Educational and Psychological Measurement* 61: 532–574.
- Ellis, P. D. 2010. *The Essential Guide to Effect Sizes: Statistical Power, Meta-Analysis, and the Interpretation of Research Results*. Cambridge: Cambridge University Press.

- Fisher, R. A. 1925. *Statistical Methods for Research Workers*. Edinburgh: Oliver & Boyd.
- Grissom, R. J., and J. J. Kim. 2012. *Effect Sizes for Research: Univariate and Multivariate Applications*. 2nd ed. New York: Routledge.
- Hays, W. L. 1963. *Statistics for Psychologists*. New York: Holt, Rinehart & Winston.
- Hedges, L. V. 1981. Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics* 6: 107–128.
- Huber, C. 2013. Measures of effect size in Stata 13. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2013/09/05/measures-of-effect-size-in-stata-13/>.
- Kelley, K. 2005. The effects of nonnormal distributions on confidence intervals around the standardized mean difference: Bootstrap and parametric confidence intervals. *Educational and Psychological Measurement* 65: 51–69.
- Kelley, K., and K. J. Preacher. 2012. On effect size. *Psychological Methods* 17: 137–152.
- Kerlinger, F. N. 1964. *Foundations of Behavioral Research*. New York: Holt, Rinehart & Winston.
- Kirk, R. E. 1996. Practical significance: A concept whose time has come. *Educational and Psychological Measurement* 56: 746–759.
- Kline, R. B. 2013. *Beyond Significance Testing: Statistics Reform in the Behavioral Sciences*. 2nd ed. Washington, DC: American Psychological Association.
- Pearson, K. 1909. On a new method of determining correlation between a measured character A, and a character B, of which only the percentage of cases wherein B exceeds (or falls short of) a given intensity is recorded for each grade of A. *Biometrika* 7: 96–105.
- Satterthwaite, F. E. 1946. An approximate distribution of estimates of variance components. *Biometrics Bulletin* 2: 110–114.
- Smith, M. L., and G. V. Glass. 1977. Meta-analysis of psychotherapy outcome studies. *American Psychologist* 32: 752–760.
- Smithson, M. 2001. Correct confidence intervals for various regression effect sizes and parameters: The importance of noncentral distributions in computing intervals. *Educational and Psychological Measurement* 61: 605–632.
- . 2003. *Confidence Intervals*. Thousand Oaks, CA: Sage.
- Steiger, J. H., and R. T. Fouladi. 1997. Noncentrality interval estimation and the evaluation of statistical models. In *What If There Were No Significance Tests?*, ed. L. L. Harlow, S. A. Mulaik, and J. H. Steiger, 221–257. Mahwah, NJ: Erlbaum.
- Thompson, B. 2006. *Foundations of Behavioral Statistics: An Insight-Based Approach*. New York: Guilford Press.
- Venables, W. 1975. Calculation of confidence intervals for noncentrality parameters. *Journal of the Royal Statistical Society, Series B* 37: 406–412.
- Welch, B. L. 1947. The generalization of 'student's' problem when several different population variances are involved. *Biometrika* 34: 28–35.

Also see

- [R] **bittest** — Binomial probability test
- [R] **ci** — Confidence intervals for means, proportions, and variances
- [R] **mean** — Estimate means
- [R] **oneway** — One-way analysis of variance
- [R] **prtest** — Tests of proportions
- [R] **sdtest** — Variance-comparison tests
- [R] **ttest** — *t* tests (mean-comparison tests)