Paper 316-2012

A SAS® Macro for Computing Point Estimates and Confidence Intervals of Effect Sizes Associated with Mediation Analysis

Merlande Petit-Bois, University of South Florida, Tampa, FL
Thanh V. Pham, University of South Florida, Tampa, FL
Eun Kyeng Baek, University of South Florida, Tampa, FL
Jeffrey D. Kromrey, University of South Florida, Tampa, FL

ABSTRACT

Measures of effect size are recommended to communicate information on the strength of relationships. Such information supplements the reject/fail-to-reject decision obtained in statistical hypothesis testing. Because sample effect sizes are subject to sampling error, as is any sample statistic, computing confidence intervals for these statistics is a useful strategy to represent the magnitude of uncertainty about the corresponding population effect sizes. This paper provides a SAS macro that uses bootstrapping to compute confidence intervals for an effect size associated with mediation analysis models. Using SAS/IML, the macro produces point and interval estimates of an R-squared effect size that represents the proportion of variance accounted for by the mediated effect. This paper provides the macro programming language, as well as an example of the macro call and output. Finally, the results from a simulation study investigating the accuracy and precision of the estimates are presented.

Keywords: EFFECT SIZES, MEDIATION ANALYSIS, CONFIDENCE INTERVALS, SAS/IML®

INTRODUCTION

Gone are the days when social and behavioral science researchers should simply report obtained test statistics (e.g., t, F, χ^2) and their corresponding *p*-values. Over the years, interpreting the importance of scientific research based on the dichotomous reject or fail-to-reject decision resulting from hypothesis tests has become less popular among some disciplines such as psychology and education. Instead, researchers are encouraged to supplement hypothesis test results with measures of effect magnitude. In fact, according to Thompson (2007), 24 peer-reviewed journals had explicit editorial policies that required authors to include effect sizes or other measures of effect magnitude. A large part of this change began over a decade ago when the APA Task Force on Statistical Inference issued their statement that researchers should regularly report effect sizes, calculate confidence intervals, and use graphics to better communicate the nature of their findings for all primary outcomes (Wilkinson & APA Task Force on Statistical Inference, 1999). Unlike *p*-values that are used to determine if an observed effect or relationship may be accounted for by sampling error alone, effect sizes are used to estimate how large the effect or relationship is. Thus, when used together, not only can researchers make statements about the statistical significance of their findings, they also can report on the practical significance of their findings.

More recently, the scholarly discussion around effect sizes has evolved to also include recommendations and formulas for calculating and reporting confidence intervals around effect sizes (Cumming & Finch, 2001; Fidler & Thompson, 2001; Finch & French, 2010; Smithson, 2001; Thompson, 2002). The use of confidence intervals around effect sizes is particularly fruitful for meta-analytic research. For example, just as a confidence interval calculated around a sample mean can generate plausible values for the population mean, a confidence interval calculated around a sample effect size such as Cohen's *d*, *r*, or η^2 can describe plausible values for the population effect size.

Effect sizes for mediation analysis have been under investigation for many years. A promising approach was suggested by Fairchild, MacKinnon, Taborga, and Taylor (2009) that constructs an R^2 type of statistic to represent the proportion of variance accounted for by the mediated effect.

MEDIATION MODELS

The simplest mediation models require information on three variables: an independent variable, a dependent variable, and a mediating variable (see Figure 1). The role of the mediating variable is to transmit some or all of the effect of the independent variable on the dependent variable. Mediation models have gained importance in recent years as researchers investigate hypothesized causal mechanisms for observed effects.

The parameter estimates needed to obtain the estimated mediated effect in this model are obtained from two standardized regression equations (MacKinnon & Dwyer, 1993):

$$\hat{M} = aX$$
$$\hat{Y} = bM + cX$$

The coefficient *c* represents the direct effect of *X* on *Y* while controlling for *M*, and *b* represents the effect of *M* on *Y* while controlling for *X*. Finally, *a* is the effect of *X* on *M*. The total effect of *X* on *Y* is the sum of the direct effect (*c*) and the indirect, or mediated, effect (*ab*).

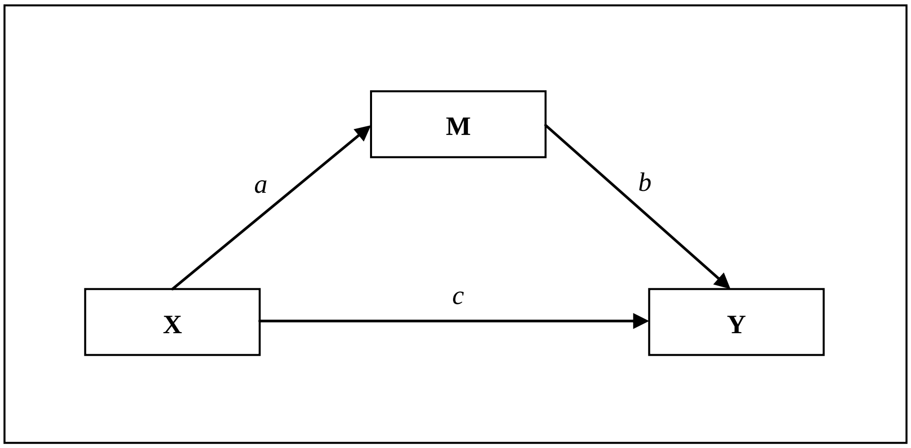

Figure 1. Simple mediation model with variables X = independent, Y = dependent, and M = mediating

EFFECT SIZES SUGGESTED FOR MEDIATION MODELS

Attempts to develop interpretable indices of effect size for mediation models have a history spanning several decades. Sobel (1982) suggested the use of ratios constructed from the indirect effect, direct effect, and total effect. For example, the ratio of the indirect effect to the total effect (ab/[ab+c]) is interpreted as the proportion of the total effect that is mediated. Hayes (2009) criticized this index because it is not bounded and because it can be negative (unlike proportions). Further, as the total effect (ab + c) approaches zero, this index rapidly approaches infinity. Taking the absolute values of both ab and c has been suggested as a variation of this index (Alwin & Hauser, 1975), but little work has been pursued in this direction (MacKinnon, Fairchild, & Fritz, 2007).A similar index is the ratio of the indirect effect to the direct effect (ab/c) but this index has the same limitations as the proportion mediated. MacKinnon, Warsi, and Dwyer (1995) investigated the sampling properties of these effect sizes for mediation using Monte Carlo methods. Their results suggested that sample sizes of at least 500 were needed for both point estimates and estimates of the sampling variability of the proportion mediated, and sample sizes of at least 2,000 were needed in some cases for the accurate estimation of the ratio of the indirect to the direct effect.

In addition to these ratio indices, effect size indices represented by partial r2 values and standardized regression slopes have been suggested. These indices, however, are inherently unsatisfactory because they represent only components of the mediation model rather than the model as an entirety (Fairchild, MacKinnon, Taborga, & Taylor, 2009).

Fairchild et al. (2009) suggested an effect size based on proportions of variance obtained from the multiple regression of *Y* on *X* and *M*, and related bivariate correlations:

$$R^2_{med} = r^2_{MY} - \left(R^2_{Y \cdot MX} - r^2_{XY}\right)$$

Using simulation methods, these authors demonstrated that this effect size estimate is an unbiased and consistent estimator of the corresponding population parameter. However, methods for the estimation of confidence intervals for this effect size were not provided. The SAS macro, CI_MEDIATE, provides a simple method for researchers to obtain these intervals.

MACRO CI_MEDIATE

The macro CI_MEDIATE uses bootstrap sampling (that is, sampling with replacement from the original, observed set of observations) to estimate confidence intervals for the effect size of the mediated effect. Arguments to the macro include the name of the SAS data set that contains the sample of observations; the names of the variables in the SAS data set corresponding to the independent, dependent, and moderating variables; the number of bootstrap samples to generate; and the level of confidence required for the intervals. The macro output is a simple table that provides the sample estimate of the effect size and the endpoints of the requested confidence interval. The macro is written in SAS/IML to provide rapid execution. The macro CI_MEDIATE was developed based on an earlier macro (Preacher & Hayes, 2004) that provided bootstrap confidence intervals for the indirect effect in mediation models. CI_MEDIATE provides confidence intervals for the effect size rather than the indirect effect itself and incorporates a more efficient bootstrap algorithm based on the RANK function in PROC IML.

Users should first execute the macro command set that is provided below or reference the macro from an external file with the %include command. This will activate a command called *%CI_MEDIATE* with syntax:

```
%macro CI_MEDIATE(dataset=_last_,y=dv,x=iv,m=mv,nboot=5000, C_I = 95);
```

Where "*_last_*" is the name of an SAS dataset containing the data to be analyzed. If users don't specify the dataset name, this macro will automatically use the data set that was most recently created. "*dv"* is the name of the dependent variable in the dataset, "*iv"* is the name of the independent variable, and "*mv"* is the name of mediator. This macro is set up to run 5000 bootstrap samples as a default but users can specify the number of bootstrap samples by changing the number that follows the *"nboot="* option. If the specified number of bootstrap samples is less than 1000, the macro will increase it to 1000. The default confidence interval is set up at 95%. In order to change the desired confidence interval, users will change the number that follows the *"C_I="* option. This number can be any value between 1 and 99 to provide levels of confidence from 1% to 99%. If the input number is greater than 99 or smaller than 1, the macro will set up the level of confidence for the confidence interval to be 95%.

The macro will delete all cases from the analysis that are missing any of the three variables, where missing is defined as the default SAS missing value for numeric variables ("."). The macro also checks the correlation matrix of the independent, dependent, and mediating variables. If this matrix is singular, the macro will exit and provide users a message informing them about the singular matrix problem.

The SAS code for the macro is provided as the following.

```
%macro CI_MEDIATE(dataset=_last_,y=dv,x=iv,m=mv,nboot=5000, C_I = 95);
Title " ";
Proc iml;

/* Read SAS dataset into the matrix DATA */
   use &dataset where (&y ^= . & &x ^= . & &m ^= .);
   read all var {&y &x &m} into data;

/* Compute the sample size */
   n=nrow(data;

/* Initialize vector BOOT_RES to hold bootstrap values */
   btn = &nboot;
   if btn < 1000 then btn = 1000; /* minimum number of bootstrap samples */
   boot_res=j(btn,1,0);

/* Initialize matrix BOOT_DATA to hold bootstrap sample */
   boot_data = data;

/* Initialize indicator of singular correlation matrix */
   sample_rii_singular = 0;

/* Begin loop for bootstrap sampling */
do btni = 1 to btn;
      if (btni > 1) then do;
                  do ni = 1 to n;
                        v = int(uniform(0)*n)+1;
                        boot_data[ni,1:3]=data[v,1:3];
                  end;
      end;

      /* Reset the effect size for each loop*/
      effect_size = 0;

      /* Compute sample correlation matrix */
      xt=boot_data-j(n,1)*boot_data[:,];
      cv=(xt`*xt)/(n-1);
```

```
        sd=sqrt(diag(cv));
        detsd=det(sd);

        /* If all variances are greater than zero proceed with calculation */
        if detsd ^= 0 then do;
                    r=inv(sd)*cv*inv(sd);
                    rii = r[2:3,2:3];
                    ryi = r[2:3,1];
                    detrii = det(rii);

                    /* If R matrix is non-singular, compute effect size */
                    if detrii ^= 0 then do;
                                coefi = inv(rii)*ryi;
                                Rsquare = coefi`*ryi;
                                rmy = r[3,1];
                                rxy = r[2,1];
                                effect_size = rmy##2 - rsquare + rxy##2;

                                /* Store effect size in BOOT_RES */
                                boot_res[btni,1] = effect_size;
                    end;

                    /* If bootstrap R matrix is singular, draw another sample */
                    else if btni > 1 then btni = btni - 1;

                    /* If observed R matrix is singular, exit */
                    else do;
                          btni = btn;
                          sample_rii_singular = 1;
                    end;
              end;
        else if btni > 1 then btni = btni - 1;

  /* End loop for bootstrap sampling */
  end;


  /* If sample matrix was singular print message and exit */
  if sample_rii_singular = 1 then do;
     file print;
       put @35 "Correlation Error!!";
       put @30 ""//
            @30 "The variables in the sample data are linearly dependent."/
           @30 "The effect size and confidence interval cannot be estimated."/;
  end;

  /* If sample matrix was non-singular then proceed */
  else do;
              Seffect_size = boot_res[1,1];
              c_i = &c_i;

              /* Reset level of confidence if out of range */
              if c_i > 99 | c_i <1 then c_i = 95;

              /* Find interval endpoints */
              c_i_s = (100-c_i)/200;
              pctLO = round(btn#c_i_s);
              pctHI = round(btn#(1-c_i_s));
              r_vec = rank(boot_res);
              do btni = 1 to btn;
                    if r_vec[btni] = pctLO then low_limit = boot_res[btni,1];
                    if r_vec[btni] = pctHI then high_limit = boot_res[btni,1];
              end;

              /* Print table with results */
              file print;
```

```
            put @35 "Effect Size Measures for Mediation Analysis "///;
            put @30 "Independent Variable: " @75 "&x "/
                @30 "Dependent variable: "   @75 "&y "/
                @30 "Mediator variable: "    @75 "&m "//;

            put @30 "Sample Size: "                 @78 n 6. /
                @30 "Confidence Level: "            @75 c_i /
                @30 "Number of Bootstrap Samples: "       @78 btn 6. /
                @30 "Sample Effect Size: "          @78 Seffect_size 6.4 /
               @30 "Lower Limit: "                  @78 low_limit 6.4 /
                @30 "Upper Limit: "                 @78 high_limit 6.4 /;
end;
quit;

%MEND CI_MEDIATE;
```

EXAMPLE OF MACRO CI_MEDIATE

An example of a call to the macro in a SAS job stream and the resulting output is provided in this section. The following data step creates a SAS data set called 'example' containing 12 observations with a single predictor variable (predictor), a mediating variable (mediator), and an outcome variable (outcome).

```
data example;
       input predictor mediator outcome;
cards;
1 2 3
2 1 2
3 3 1
4 5 4
5 4 8
6 9 9
7 8 7
8 7 6
9 6 5
10 2 3
11 1 2
12 3 1
;
```

To estimate the effect size associated with the mediation analysis and to compute a confidence interval around the effect size. The following statement will execute the macro:

```
%CI_MEDIATE (dataset = example, y = outcome, x = predictor, nboot = 10000,m =
mediator, c_i = 99);
```

The macro arguments provide the name of the SAS data set containing the data (example) and the SAS names of the predictor, mediator, and outcome variables. In this example, it was not necessary to specify the name of the SAS data set because the default value (_LAST_) would have referenced the data set example. The number of bootstrap samples, called `nboot`, is set to ten thousand in this example. A 99% confidence interval, called `c_i`, is calculated around the effect size of the example.

Figure 2 shows the output from CI_MEDIATE for the example data. The names of the variables analyzed (labeled Independent Variable, Dependent Variable, and Mediator Variable) are provided on the output page (predictor, outcome, and mediator, respectively). The number of observations in the sample is reported, followed by the requested level of confidence and the number of bootstrap samples used in the estimation of the confidence interval. The sample effect size associated with the mediation analysis of the example data, -.0105, indicates that only approximately 1% of the variance in the outcome is attributable to the indirect effect of the predictor through the mediator. The lower and upper limits of the 99% confidence interval (-.1479 and 0.6458, respectively) indicate that the researcher may be 99% confident that the population effect size for the mediation effect is between these two values. The relatively broad confidence interval results from having only 12 observations in the sample.

```
                 Effect Size Measures for Mediation Analysis


     Independent Variable:                         predictor
     Dependent variable:                           outcome
     Mediator variable:                            mediator


     Sample Size:                                         12
     Confidence Level:                                    99
     Number of Bootstrap Samples:                      10000
     Sample Effect Size:                               -.0105
     Lower Limit:                                      -.1479
     Upper Limit:                                      0.6358
```

Figure 2. Output from CI_MEDIATE

SIMULATION WORK

To investigate the accuracy and precision of the effect size index and the bootstrap confidence intervals, Monte Carlo methods were used in a study conducted Baek, Petit-Bois, Pham, and Kromrey (2012). The use of simulation methods allowed for the control and manipulation of research design factors and the incorporation of sampling error into the analyses. Population correlation matrices providing population effect sizes ranging from -.72 to .77 were investigated and sample sizes ranging from 50 to 500 were simulated. For each sample generated, the sample estimate of effect size was calculated and a percentile bootstrap was used to estimate the endpoints of a 95% confidence interval (Efron & Gong, 1983; Efron & Tibshirani, 1986; Stine, 1990). The outcomes examined included the statistical bias and RMSE of the point estimates of effect size, and the estimated coverage probabilities and interval widths of the bootstrap confidence intervals. For each condition examined in the Monte Carlo study, 5,000 samples were simulated. Study results indicated that the point estimates of effect size present very little statistical bias across the conditions examined and the interval estimates maintain near nominal coverage probabilities, although both indices improve with larger samples (see Figures 3 and 4). Interval widths and RMSE, however, suggest that substantial sampling error is evident with the smallest samples (see Figures 5 and 6).

STATISTICAL BIAS

As expected, the sample size was the factor most strongly associated with variation in bias (η^2 = .23). Study results indicate that the point estimates of effect size present very little statistical bias across the conditions examined. The distributions of statistical bias by sample size are presented in Figure 3. Evident in this figure is the reduction in statistical bias as sample size increased. In addition, the variability in the bias estimates was notably smaller in the larger samples. For example, when the sample size is 500, the bias is approximately 0, with very little variability across conditions examined at this sample size.

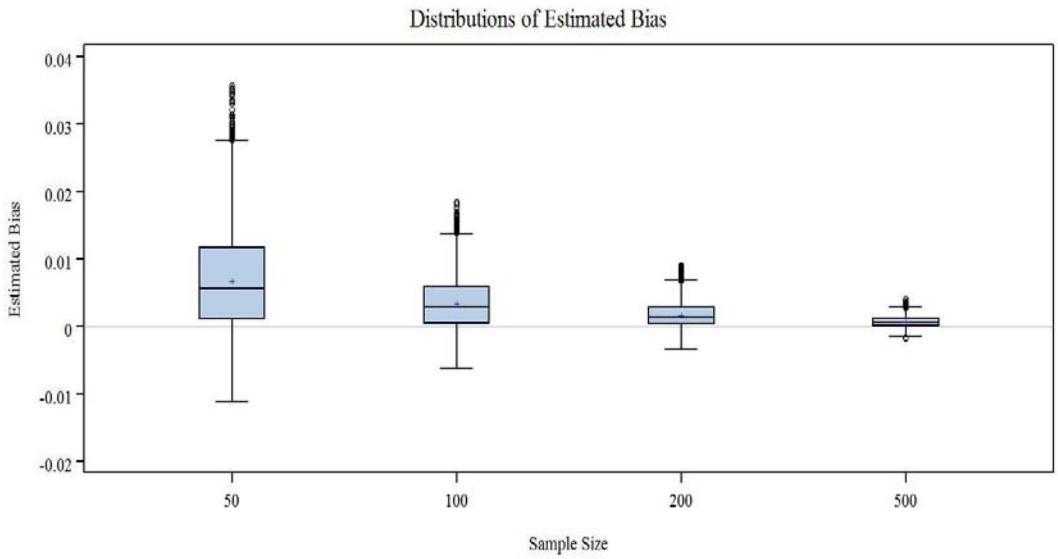

Figure 3. Distributions of Estimated Bias in Point Estimates of Effect Size

The distributions of statistical bias by population effect size are presented in Figure 4. As is illustrated by the figure, the more negative effect sizes had more positive statistical bias. The inverse was also true; the moderately positive effect sizes had negative bias, but the negative bias is less extreme than the positive bias for the negative effect sizes. The variability was greatest for the moderate negative effect sizes, while the positive effect sizes tended to have less variability. The statistical bias approached 0 for effect size of 0.2.

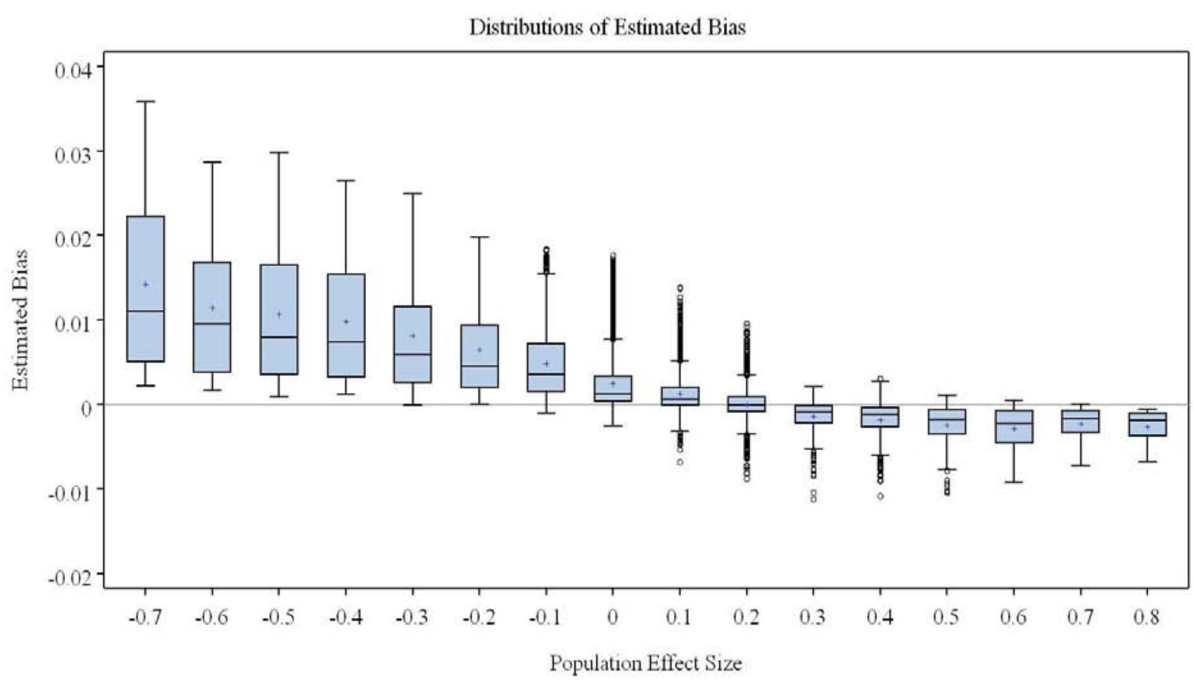

Figure 4. Distributions of Bias by Population Effect Size

CONFIDENCE INTERVAL COVERAGE

The sample size was the factor most strongly associated with variation in interval coverage as expected. The distributions of interval coverage by sample size are presented in Figure 5. Evident in this figure is the increase in mean probability of confidence interval coverage as sample size increased. In addition, the variability in the coverage of the intervals was notably smaller in the larger samples.

The distributions of confidence interval coverage by population effect size are presented in Figure 6. The highest coverage probabilities on average were obtained with a population effect size of zero. However, the average coverage and variability were maintained consistently near nominal coverage probability across various population effect sizes. A notably large dispersion in coverage probabilities was observed with a population of effect size of zero.

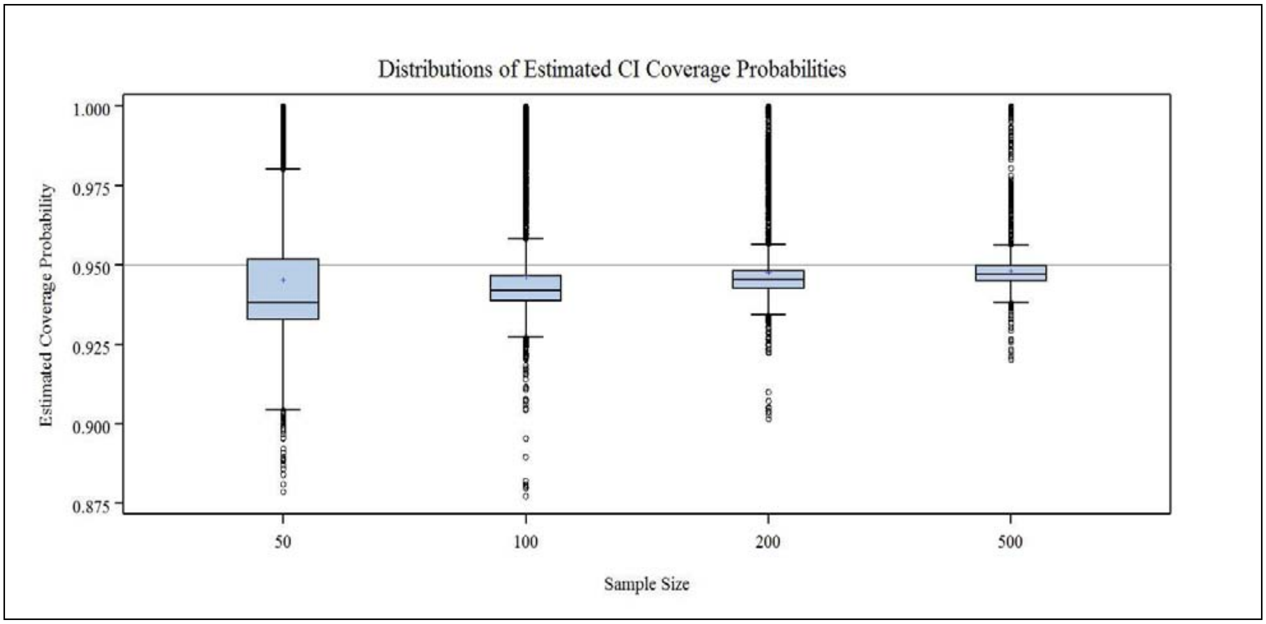

Figure 5. Distributions of Estimated Confidence Interval Coverage

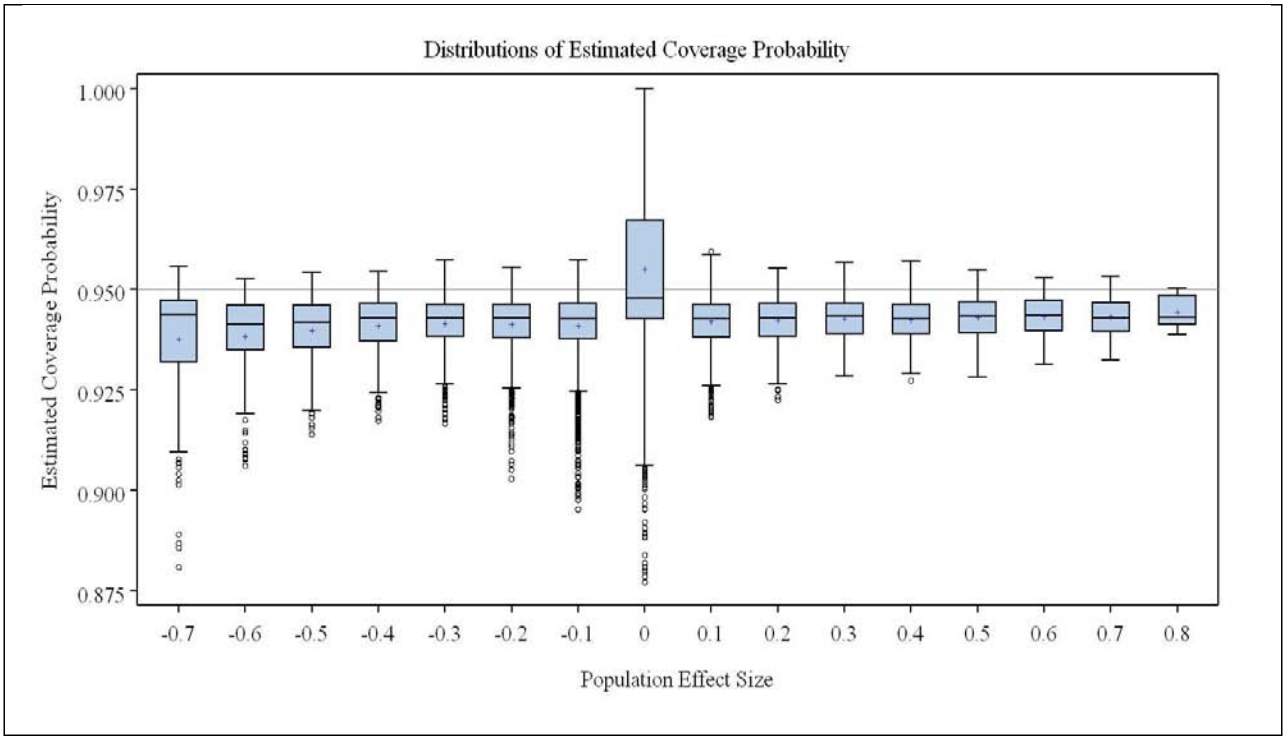

Figure 6. Distributions of Estimated Confidence Interval Coverage by Population Effect Size

The average confidence interval coverage probabilities were also associated with the magnitude of each of the population correlations used in the simulation (with η^2 ranging from .09 to .13). Because the pattern was similar across correlation coefficients, only the correlation between the independent variable and the moderating variable (r_{mx}) is displayed in Figure 7. The average interval coverage probabilities were all near the nominal confidence interval level (95%). The highest coverage was obtained with a population correlation of zero and the dispersion increased as correlation departed from zero.

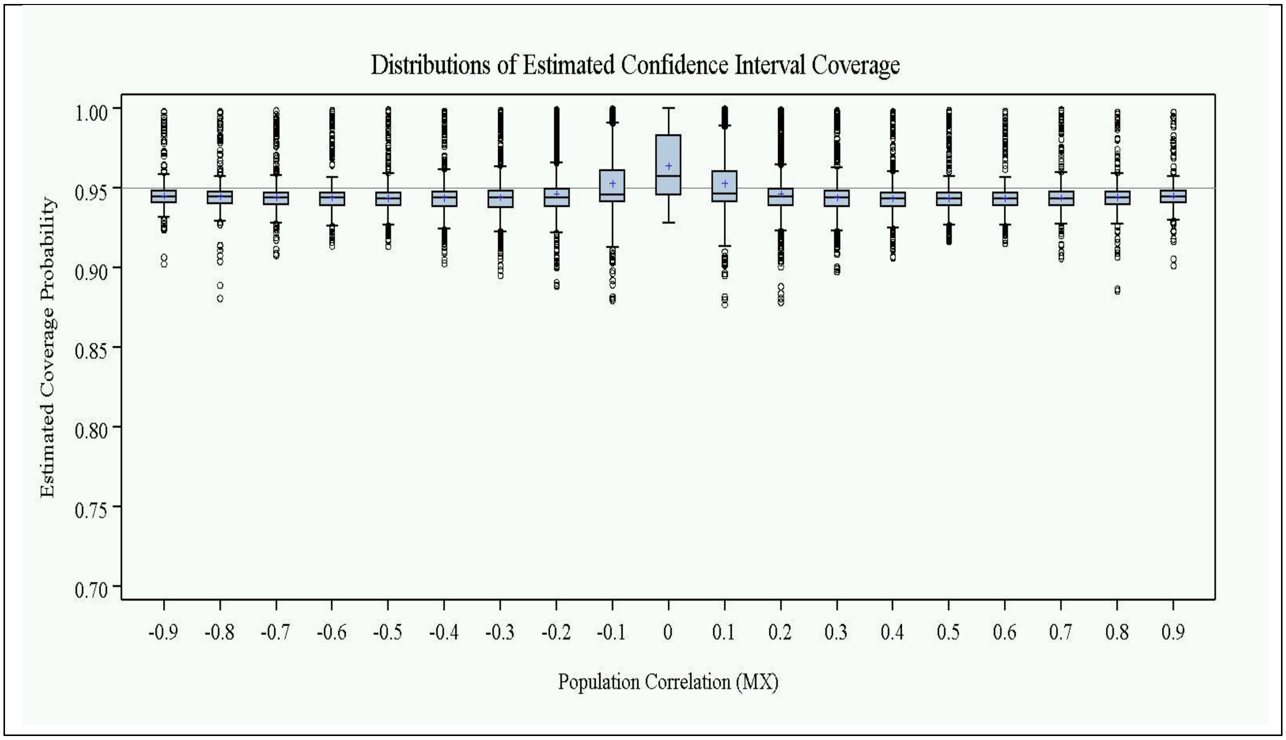

Figure 7. Distributions of Estimated Confidence Interval Coverage by Population Correlation (MX)

In order to further investigate the dispersion observed with a population effect size of zero, the accuracy of confidence interval coverage by three factors (sample size, population correlation, and population effect size) were computed. The accuracy of confidence interval coverage was assessed by using Bradley's (1978) liberal robustness criterion. By his method, the proportion of time the parameter was contained within the confidence intervals was compared with the nominal proportion (.95). It assesses if the proportions deviated significantly from the expected proportion. Proportions were considered robust if they fell within the range (1-α) ± α/2. For example, for 95% confidence intervals, a robust confidence interval method should contain the parameter between 92.5% and 97.5% of the time.

Table 1 illustrates the proportion of conditions that meets Bradley's liberal criterion by sample size. As seen in the Table 1, when sample size increased, the proportion that meets the criterion is increased. With a sample size of 500, 98.5% of the conditions provided accurate confidence intervals (in contrast with 77.5% of the conditions with a sample size of 50). This result is consistent with Figure 6.

Table 1
Confidence interval coverage proportions by sample size with Bradley Liberal Criterion

Sample size	N of Out range (%)	N of Criteria range (%)	Total (%)
50	1022(22.5)	3523(77.5)	4545 (25.0)
100	407 (9.0)	4124 (91.0)	4531 (25.0)
200	260 (5.8)	4258 (94.2)	4518 (25.0)
500	67 (1.5)	4444 (98.5)	4511 (25.0)
Total (%)	1756 (9.7)	16349 (90.3)	18105 (100)

Table 2 presents the proportion of conditions with confidence interval coverage that meet Bradley's liberal criterion by population effect size. The accuracy of confidence interval coverage was high and consistent across various population effect size values (ranging from 82.9 to 100), except the population effect size of zero. This effect size has the smallest percentage of conditions with accurate intervals (79.7%) and it has the largest number of conditions (*N*=6985) among various population effect sizes.

Table 2
Confidence interval coverage proportions by population effect size with Bradley Liberal Criterion

Population effect size	N of Out range (%)	N of Criteria range (%)	Total (%)
-.7	27 (17.1)	131 (82.9)	158 (0.9)
-.6	14 (11.1)	112 (88.9)	126 (0.7)
-.5	18 (8.8)	187 (91.2)	205 (1.1)
-.4	19 (4.9)	368 (95.1)	387 (2.1)
-.3	24 (3.2)	718 (96.8)	742 (4.1)
-.2	59 (4.1)	1389 (95.9)	1448 (8.0)
-.1	153 (5.9)	2451 (94.1)	2604 (14.4)
0	**1416 (20.27)**	**5569 (79.7)**	**6985 (38.6)**
.1	23 (0.9)	1509 (99.1)	2532 (99.1)
.2	2 (0.1)	1486 (99.9)	1488 (99.9)
.3	0 (0.0)	483 (100)	483 (100)
.4	0 (0.0)	436 (100)	436 (100)
.5	0 (0.0)	319 (100)	319 (100)
.6	0 (0.0)	134 (100)	134 (100)
.7	0 (0.0)	38 (100)	38 (100)
.8	0 (0.0)	19 (100)	19 (100)
Total (%)	1755 (9.7)	16349 (90.3)	18104 (100)

Table 3 presents the proportion of conditions with confidence interval coverage that meet Bradley's liberal criterion by population correlation (XM). Consistent with Figure 8, the accuracy of confidence interval coverage was similar across various population correlation values and the smallest percentage of conditions with accurate intervals (66%) occurs with a population correlation of zero.

Table 3
Confidence interval coverage proportions by population correlation (XM) with Bradley Liberal Criterion

Population Correlation	N of Out range (%)	N of Criteria range (%)	Total (%)
-.9	21(3.9)	515 (96.1)	536 (3.0)
-.8	32 (4.5)	683 (95.5)	715 (4.0)
-.7	38 (4.5)	808 (95.5)	846 (4.7)
-.6	42 (4.6)	881 (95.4)	923 (5.1)
-.5	54 (5.4)	951 (94.6)	1005 (5.6)
-.4	74 (7.0)	987 (93.0)	1061 (5.9)
-.3	96 (8.5)	1038 (91.5)	1134 (6.3)
-.2	142 (12.2)	1021 (87.8)	1163 (6.4)
-.1	248 (21.0)	935 (79.0)	1183 (6.5)
0	**276 (34.0)**	**535 (66.0)**	811 (4.5)
.1	234 (19.8)	949 (80.2)	1183 (6.5)
.2	140 (11.8)	1045 (88.2)	1185 (6.6)
.3	97 (8.4)	1053 (91.6)	1150 (6.4)
.4	72 (6.6)	1024 (93.4)	1096 (6.1)
.5	50 (4.9)	977 (95.1)	1027 (5.7)
.6	49 (5.2)	903 (94.8)	952 (5.3)
.7	34 (3.9)	833 (96.1)	867 (4.8)
.8	34 (4.6)	701 (95.4)	735 (4.1)
.9	22 (4.1)	510 (95.9)	532 (3.0)
Total (%)	1755 (9.7)	16349 (90.3)	18104 (100)

CONCLUSION

Effect sizes are used to describe the strength of a relationship. When researchers use effect sizes together with the results of hypothesis tests, they can make statements about the statistical significance as well as the practical significance of their findings. Effect sizes for mediation analysis have been investigated for many years. Fairchild et al. (2009) suggested an approach to estimate the effect sizes for mediation analysis and presented a SAS macro to estimate these effect sizes. However, a method to estimate confidence intervals for these effect sizes was not developed. The CI_MEDIATE macro provides a helpful tool for researchers to estimate confidence intervals for the effect size of the mediated effect. By using both R^2_{med} and a confidence interval, not only can researchers report on the practical significance of their findings, they also can describe plausible values for the population effect size. The macro may be modified easily to provide additional information. For example, the set of effect sizes from the bootstrap samples may be output for further analysis or for graphical display.

Simulation work demonstrated sample size was most strongly associated with both of the outcomes (statistical bias and confidence interval coverage). As sample size increased, the parameter estimates converged to the parameter values and the variability of the estimates tended to decrease. The variability tended to increase as the effect size moved away from zero for confidence interval coverage.. However, for statistical bias, the variability increased as population effect size became more extreme and negative and decreased as the effect size became more positive. The confidence interval coverage probabilities were near the nominal level for the majority of conditions examined.

REFERENCES

Alwin, D. F. & Hauser, R. M. (1975). The decomposition of effects in path analysis. *American Sociological Review*, *40*, 37 – 47.

Baek, E. K., Petit-Bois, M., Pham, T. V., & Kromrey, J. D. (2012, February). *Point and interval estimates of effect size for mediation analysis: An empirical investigation of accuracy and precision*. Paper presented at the annual meeting of the Eastern Educational Research Association, Hilton Head, SC.

Bradley, J. V. (1978). Robustness? *British Journal of Mathematical and Statistical Psychology, 31,* 144-152.

Cumming G. & Finch, S. (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*, 532-74.

Efron, B. & Gong, G. (1983). *A leisurely look at the bootstrap, the jackknife, and cross-validation. The American Statistician, 37(1),* pg 36-49.

Efron, B. & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science, 1(1),* p. 54-77.

Fairchild, A. J., MacKinnon, D. P., Taborga, M. P., & Taylor, A. B. (2009). R^2 effect-size measures for mediation analysis. *Behavior Research Methods*, *41*, 486-498.

Fidler F. & Thompson B. (2001). Computing correct confidence intervals for ANOVA fixed and random-effects effect sizes. *Educational and Psychological Measurement, 61*, 575-604

Finch, W. H. & French, B. F. (2010, May). *A comparison of methods for estimating confidence intervals for omega squared effect size*. Paper presented at the annual meeting of the American Education Research Association, Denver.

Hayes, A. F. (2009). Beyond Baron and Kenney: Statistical mediation analysis in the new millennium. *Communication Monographs*, *76*, 408 - 420.

MacKinnon, D. P. & Dwyer, J. H. (1993). Estimating mediated effects in prevention studies. *Evaluation Review*, *17*, 144-158.

MacKinnon, D. P., Fairchild, A. J., & Fritz, M. S., (2007). Mediation analysis. *Annual Review of Psychology*, *58*, 593 - 614.

Preacher K. J. & Hayes A. F. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, & Computer*, 36 (4), 717-731

Smithson, M. (2001). Correct confidence intervals for various regression effect sizes and parameters. *Educational and Psychological Measurement, 61,* 605-632.

Sobel, M. E. (1982). Asymptotic confidence intervals for indirect effects in structural equation models. *Sociological Methodology*, *13*, 290 - 312.

Stine, R. (1990). An introduction to bootstrap methods. *Sociological Methods and Research, 18 (2&3),* p. 243-291.

Thompson, B. (2002). What future quantitative social science research could look like: Confidence intervals for effect sizes. *Educational Researcher*, 31, 25-32.

Thompson, B. (2007). Effect sizes, confidence intervals, and confidence intervals for effect sizes. *Psychology in the Schools*, 44, 423-432.

Wilkinson & APA Task Force on Statistical Inference. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54, 594-604.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the first author at:

Merlande Petit-Bois
University of South Florida
4202 East Fowler Ave., EDU 105
Tampa, FL 33620

Work Phone: (813) 974-4933
Fax: (813) 974-6126
Email: mpetitbo@usf.edu