Correcting the Bias Correction for the Bootstrap Confidence Interval in Mediation
Analysis

Tristan D. Tibbe, Amanda K. Montoya

University of California, Los Angeles

Abstract

The bias-corrected bootstrap confidence interval (BCBCI) was once the method of choice for conducting inference on the indirect effect in mediation analysis due to its high power in small samples, but now it is criticized by methodologists for its inflated type I error rates. In its place, the percentile bootstrap confidence interval (PBCI), which does not adjust for bias, is currently the recommended inferential method for indirect effects. This study proposes two alternative bias-corrected bootstrap methods for creating confidence intervals around the indirect effect. Using a Monte Carlo simulation, these methods were compared to the BCBCI, PBCI, and a bias-corrected method introduced by Chen and Fritz (2021). The results showed that the methods perform on a continuum, where the BCBCI has the best balance (i.e., having closest to an equal proportion of CIs falling above and below the true effect), highest power, and highest type I error rate; the PBCI has the worst balance, lowest power, and lowest type I error rate; and the alternative bias-corrected methods fall between these two methods on all three performance criteria. An extension of the original simulation that compared the bias-corrected methods to the PBCI after controlling for type I error rate inflation suggests that the increased power of these methods might only be due to their higher type I error rates. Thus, if control over the type I error rate is desired, the PBCI is still the recommended method for use with the indirect effect. Future research should examine the performance of these methods in the presence of missing data, confounding variables, and other real-world complications to enhance the generalizability of these results.

Correcting the Bias Correction for the Bootstrap Confidence Interval in Mediation Analysis

Mediation analysis is a statistical method that allows researchers to examine how one variable is able to influence another variable. It is a valuable tool in psychology research because it allows psychologists to expose the mechanism(s) through which psychological phenomena operate. As an example, Osberg and Eggert (2012) used mediation analysis to show that daily hassles increased bulimic symptoms in students by first increasing their irrational food beliefs. Thus, they were able to reveal not only that daily hassles affected students' bulimic symptoms, but also *how* they did so.

Unlike a simple linear regression equation that estimates only the effect of a predictor variable X on an outcome variable Y , a **simple mediation model** (like the one used by Osberg and Eggert, 2012) involves a system of regression equations that estimate the **indirect effect** of a predictor X on an outcome Y through a single mediator variable M . A significance test is then conducted to see if the mediator variable is the means through which X is able to affect Y (i.e., if X influences M and in turn M influences Y). In the example provided above, X would be daily hassles, Y would be bulimic symptoms, and M would be irrational food beliefs. The system of linear regression equations that constitutes a simple mediation model is as follows,

$$Y_i = d_1 + cX_i + e_{1i} \tag{1}$$

$$M_i = d_2 + aX_i + e_{2i} \tag{2}$$

$$Y_i = d_3 + c'X_i + bM_i + e_{3i} \tag{3}$$

where the ds represent the intercepts of the equations, the es represent the error terms (assumed to be independent, normally distributed random variables with means of zero and constant variances), and the subscript i indicates that the variable values belong to the i th individual. The equations above contain the parameters we estimate when we run a simple mediation analysis on a sample of data.

To help visualize the relationship between X , Y , and M , see the path diagram in Figure 1. In the diagram, the effect of X on M is found by tracing the a -path from X

to M , and the effect of M on Y controlling for X is found by tracing the b -path from M to Y . The indirect effect of X on Y through M is found by tracing from the a -path starting at X through the b -path ending at Y . Mathematically, the indirect effect is calculated by multiplying the a and b coefficients from Equations 2 and 3. Thus, in the context of the Osberg and Eggert (2012) study, ab would give the effect of daily hassles on students' bulimic symptoms through irrational food beliefs. In a sample, the true indirect effect ab is estimated by $\hat{a}\hat{b}$, the product of the ordinary least squares estimates of a and b : $\hat{a}$ and $\hat{b}$, respectively.

Even though both $\hat{a}$ and $\hat{b}$ are normally distributed under the assumptions of linear regression, their product $\hat{a}\hat{b}$ is *not* normally distributed (Aroian, 1947; Aroian, Taneja, & Cornwell, 1978; Craig, 1936). As a result, the use of normal theory tests of statistical significance with the indirect effect have been shown to have lower power to detect a true effect than methods that do not assume normality (MacKinnon, Lockwood, Hoffman, West, & Sheets, 2002). To circumvent this issue, researchers have employed many inferential methods that do not assume anything about the shape of the indirect effect's sampling distribution. One of the most popular options is the nonparametric bootstrap confidence interval, with one recent study finding that a bootstrap confidence interval was applied to the indirect effect in over half of all mediation articles published in *Psychological Science* from 2011 to 2012 (Hayes & Scharkow, 2013). Rather than making assumptions about what the sampling distribution of the indirect effect looks like, the bootstrap allows researchers to form an empirical approximation of this sampling distribution. Leveraging the power of modern computers, bootstrap procedures resample with replacement from the original sample of data many times (e.g., 1000 times). Each of these 1000 bootstrap samples is the same size as the original sample, and from each one a bootstrap indirect effect estimate, $\hat{a}\hat{b}^*$, is computed. These 1000 bootstrap indirect effect estimates form an observed bootstrap sampling distribution of the indirect effect. This observed bootstrap distribution containing 1000 bootstrap estimates is distinct from the true bootstrap sampling distribution, which contains all possible bootstrap estimates that can be drawn from the

sample. The observed bootstrap sampling distribution will be referred to simply as the bootstrap sampling distribution throughout the rest of this paper.

Based on the bootstrap sampling distribution, a bootstrap confidence interval for the indirect effect can be formed by first ordering the bootstrap estimates from smallest to largest (Efron & Tibshirani, 1993). Multiple methods can then be used to determine the endpoints of the bootstrap confidence interval. Of those available, two have been used the most in the mediation literature: the percentile bootstrap confidence interval (PBCI) and the bias-corrected bootstrap confidence interval (BCBCI).

Bootstrap Confidence Interval Techniques

The following two subsections detail the motivation and implementation of these two bootstrap confidence interval techniques, as well as describe type I error rate issues with the bias-corrected bootstrap confidence interval in mediation analysis that motivate the current simulation study. The manuscript then transitions to discussing the primary goal of this research: developing alternative bias-corrected bootstrap techniques that address these issues with the traditional bias-corrected bootstrap confidence interval when applied to the indirect effect.

Percentile Bootstrap Confidence Interval

After ordering the bootstrap indirect effect estimates from smallest to largest and determining the acceptable type I error rate for the analysis, α , the $\frac{\alpha}{2} \times 100$ th and $(1 - \frac{\alpha}{2}) \times 100$ th percentiles of the bootstrap sampling distribution are used to form the lower and upper limits of what is known as a $(1 - \alpha) \times 100\%$ PBCI (Efron & Tibshirani, 1993).

Both the PBCI and the BCBCI, described later, rely on the assumption that some monotonic increasing function, $g(\cdot)$, exists to transform the sampling distributions of the sample indirect effect and its bootstrap estimator to a known symmetric distribution with variance $\sigma_{g(\hat{a}\hat{b})}^2$, the variance of $g(\hat{a}\hat{b})$ (Efron, 1982a; Stine, 1989). The PBCI's and BCBCI's assumptions differ, however, in regards to the expected values of those transformed sampling distributions. The normal distribution is usually chosen to

be the known symmetric distribution—and it is used throughout the rest of this paper—because it is the asymptotic distribution of most statistics (Stine, 1989).

The $(1 - \alpha) \times 100\%$ PBCI created from the true bootstrap sampling distribution—the bootstrap sampling distribution containing all possible bootstrap samples—has correct coverage (i.e., has $\alpha \times 100\%$ of the true sampling distribution of the indirect effect outside its limits) if the following holds true:

$$E(g(\hat{a}\hat{b})) = g(ab) \quad (4)$$

$$E(g(\hat{a}\hat{b}^*)) = g(\hat{a}\hat{b}). \quad (5)$$

That is, the expected value of the transformed sample indirect effect $g(\hat{a}\hat{b})$ has to be equal to the transformed true indirect effect $g(ab)$, and the expected value of the transformed bootstrap indirect effect $g(\hat{a}\hat{b}^*)$ has to be equal to the transformed sample indirect effect $g(\hat{a}\hat{b})$ (Stine, 1989). If either Equation 4 or Equation 5 does not hold true, then there is bias present in either the transformed sample indirect effect or the transformed bootstrap indirect effect and the PBCI does not have correct coverage. The BCBCI is a modification of the PBCI that loosens this requirement in an attempt to form confidence intervals with correct coverage even in the presence of bias.

Bias-Corrected Bootstrap Confidence Interval

In place of Equations 4 and 5, the BCBCI created using the true bootstrap sampling distribution has correct coverage if

$$E(g(\hat{a}\hat{b})) = g(ab) - z_{adj} \times \sigma_{g(\hat{a}\hat{b})} \quad (6)$$

$$E(g(\hat{a}\hat{b}^*)) = g(\hat{a}\hat{b}) - z_{adj} \times \sigma_{g(\hat{a}\hat{b})} \quad (7)$$

where z_{adj} is a constant and $\sigma_{g(\hat{a}\hat{b})}$ is the standard error of $g(\hat{a}\hat{b})$ (Stine, 1989). As before with the PBCI, these transformed terms take $\hat{a}\hat{b}$ and $\hat{a}\hat{b}^*$ and reshape their sampling distributions so that they both have a known, symmetric sampling distribution (i.e., the normal distribution). The BCBCI then goes a step further than the PBCI by allowing the transformed sample and bootstrap estimators to have bias of the form

$-z_{adj} \times \sigma_{g(\hat{a}\hat{b})}$. This means that as the variability of the estimate decreases (i.e., as $\sigma_{g(\hat{a}\hat{b})}$ gets closer and closer to zero), so too does the bias of these transformed estimators.

Even if the transformed estimators are biased, however, this does not imply that there is bias of the untransformed sample indirect effect (**sample bias**) or bias of the untransformed bootstrap indirect effect (**bootstrap bias**). This is due to the fact that $g(\cdot)$ does not have to be a linear transformation, only monotonic increasing, and bias of the mean (referred to simply as bias throughout this paper) is not preserved under nonlinear transformations (e.g., Needham, 1993). Thus, the expected value of a statistic can be equal to the parameter it estimates, but this does not imply the expected value of a nonlinear transformation of that statistic will be equal to the nonlinear transformed value of its parameter. A good example of this is the sample variance, s^2 : Although s^2 is a mean unbiased estimator of the population variance σ^2 (i.e., $E(s^2) = \sigma^2$), the square root of s^2 —which is a nonlinear, monotonic increasing transformation of s^2 —is a biased estimator of the population standard deviation (i.e., $E(s) \neq \sigma$). Similarly, $\hat{a}\hat{b}$ and $\hat{a}\hat{b}^*$ could be unbiased even if $g(\hat{a}\hat{b})$ and $g(\hat{a}\hat{b}^*)$ are not, and vice versa.

Unlike bias of the mean, bias of the median *is* preserved under nonlinear transformations (e.g., Brown, 1947), and because $g(\cdot)$ transforms to a symmetric distribution, the mean of the transformed distribution is equal to its median. Thus, Equations 6 and 7 imply that

$$Med(g(\hat{a}\hat{b})) = g(ab) - z_{adj} \times \sigma_{g(\hat{a}\hat{b})} \quad (8)$$

$$Med(g(\hat{a}\hat{b}^*)) = g(\hat{a}\hat{b}) - z_{adj} \times \sigma_{g(\hat{a}\hat{b})} \quad (9)$$

where $Med(\cdot)$ signifies the median of the variable in parentheses. Using the fact that bias of the median is preserved under nonlinear transformations, then, we find that Equations 8 and 9 imply that

$$Med(\hat{a}\hat{b}) = g^{-1}(g(ab) - z_{adj} \times \sigma_{g(\hat{a}\hat{b})}) \quad (10)$$

$$Med(\hat{a}\hat{b}^*) = g^{-1}(g(\hat{a}\hat{b}) - z_{adj} \times \sigma_{g(\hat{a}\hat{b})}) \quad (11)$$

(see Efron, 1979). If $z_{adj} = 0$ in Equations 10 and 11, the $z_{adj} \times \sigma_{g(\hat{a}\hat{b})}$ terms drop out and so $g^{-1}(\cdot)$ and $g(\cdot)$ cancel, leaving ab as the median of $\hat{a}\hat{b}$'s sampling distribution and

$\hat{a}\hat{b}$ as the median of $\hat{a}\hat{b}^*$'s sampling distribution. The farther from zero z_{adj} is, the farther the medians of the untransformed sampling distributions are from the corresponding true parameter values.

Thus, the bias constant z_{adj} is a measure of the median bias of the untransformed sample and bootstrap estimators. It is found using the equation

$$z_{adj} = \Phi^{-1} \left(G^*(\hat{a}\hat{b}) \right) \quad (12)$$

where Φ^{-1} is the inverse of the normal cumulative distribution function and $G^*(\hat{a}\hat{b})$ is the cumulative distribution function of the true bootstrap sampling distribution evaluated at the sample indirect effect (Stine, 1989). In other words, z_{adj} is the z -score corresponding to the probability that a bootstrap estimate randomly selected from the true bootstrap sampling distribution of the indirect effect will be less than or equal to the original sample estimate. If this probability is .50, and thus the original estimate is the median of the true bootstrap sampling distribution, the corresponding z -score will be zero and so the bias terms in Equations 6 through 11 will be equal to zero as well.

In practice, G^* is unknown, and thus z_{adj} is estimated using

$$\hat{z}_{adj} = \Phi^{-1} \left(\frac{\#\hat{a}\hat{b}^* < \hat{a}\hat{b}}{B} \right) \quad (13)$$

where $\#\hat{a}\hat{b}^* < \hat{a}\hat{b}$ is the number of bootstrap indirect effect estimates in the (observed) bootstrap sampling distribution that are less than the original sample estimate and B is the total number of bootstrap indirect effect estimates collected (e.g., 1000) (Efron & Tibshirani, 1993). Thus, similar to z_{adj} , $\hat{z}_{adj}$ is the z -score corresponding to the proportion of bootstrap estimates less than the original sample estimate in the bootstrap sampling distribution.

A $(1 - \alpha) \times 100\%$ BCBCI is formed by first calculating the z -scores corresponding to the percentiles at which to place the confidence interval's lower and upper limits using the following equations:

$$\text{Lower: } 2 \times \hat{z}_{adj} + z_{\frac{\alpha}{2}} \quad (14)$$

$$\text{Upper: } 2 \times \hat{z}_{adj} + z_{1-\frac{\alpha}{2}} \quad (15)$$

where $z_{\frac{\alpha}{2}}$ is the z -score corresponding to the $\frac{\alpha}{2} \times 100$ th percentile of the standard normal distribution and $z_{1-\frac{\alpha}{2}}$ is the z -score corresponding to the $(1 - \frac{\alpha}{2}) \times 100$ th percentile of the standard normal distribution (Efron & Tibshirani, 1993). The $\hat{z}_{adj}$ term is added to both z -scores twice: once to account for the estimated median bias of the sample indirect effect and once to account for the estimated median bias of the bootstrap indirect effect.

The final step in forming the BCBCI is converting the z -scores from Equations 14 and 15 to the percentiles of the bootstrap sampling distribution at which to place the lower and upper confidence limits using the following equations:

$$\text{Lower: } \Phi \left(2 \times \hat{z}_{adj} + z_{\frac{\alpha}{2}} \right) \quad (16)$$

$$\text{Upper: } \Phi \left(2 \times \hat{z}_{adj} + z_{1-\frac{\alpha}{2}} \right) \quad (17)$$

where Φ is the normal cumulative distribution function. Note that if $\hat{z}_{adj} = 0$, Equations 16 and 17 will produce the $\frac{\alpha}{2} \times 100$ th and $(1 - \frac{\alpha}{2}) \times 100$ th percentiles of the bootstrap sampling distribution, resulting in a $(1 - \alpha) \times 100\%$ BCBCI that is equal to the corresponding $(1 - \alpha) \times 100\%$ PBCI.

Although the BCBCI is designed to have less strict assumptions than the PBCI—and thus be appropriate in more situations—there is a growing body of research indicating that confidence intervals for the indirect effect created by the bias-corrected bootstrap have inflated type I error rates compared to the PBCI and other confidence interval methods, particularly with smaller sample sizes ($n < 500$) (Biesanz, Falk, & Savalei, 2010; Chen & Fritz, 2021; Fritz, Taylor, & MacKinnon, 2012; MacKinnon, Lockwood, & Williams, 2004). Fritz et al. (2012) found that increasing the number of bootstrap samples used to create the BCBCI had no effect on these elevated type I error rates. The issue is perpetuated by the BCBCI's continued popularity in published mediation articles, recently demonstrated by Götz, O'Boyle, Gonzalez-Mulé, Banks, and Bollmann (2021). They collected all mediation articles that applied resampling techniques to the indirect effect published from 2018 to 2019 in five of the APA's most prominent journals and found that about 25% still used the BCBCI in at least one mediation analysis.

Researchers prefer the BCBCI to competing inferential methods because of its higher power to detect a true indirect effect (MacKinnon et al., 2004; Preacher & Hayes, 2008; Williams & MacKinnon, 2008). It may offer advantages in terms of balance too. Balance refers to how evenly confidence intervals fall below and above the true effect. For example, we would expect a 95% confidence interval to capture the true effect 95% of the time and exclude the true effect 5% of the time (meaning it has correct coverage). If the confidence interval were perfectly balanced as well, it would fall below the true effect 2.5% of the time and it would fall above the true effect 2.5% of the time. Previous research has found evidence that the BCBCI may offer better balance than the PBCI with nonzero indirect effects (Williams & MacKinnon, 2008). Although it was also found that the PBCI may be more balanced when the indirect effect is equal to zero, an examination of the tabled data from Williams and MacKinnon (2008) revealed that the performance of these two methods were quite similar overall. Additionally, a later study by Fritz et al. (2012) failed to find a significant difference in the balance of the PBCI and BCBCI. Still more recent, supplemental material from Chen and Fritz (2021) suggests that the BCBCI may offer better overall balance across both zero and nonzero indirect effects. Clearly, further comparison of the balance of these two methods is warranted, and it would be beneficial to find a balanced alternative to the BCBCI that maintains its high power and controls type I error rates for the indirect effect.

Chen and Fritz (2021) attempted to identify such an alternative in a recent simulation study. They proposed several bias-corrected bootstrap methods that replaced $\hat{a}\hat{b}$ in Equation 13 with different measures of central tendency, including the Winsorized mean of the bootstrap sampling distribution set at different percentages of Winsorization. For example, their 30% Winsorized bias-corrected bootstrap confidence interval (WBCBCI) involved first finding the 30% Winsorized mean of the bootstrap sampling distribution of the indirect effect, and then calculating the bias correction by replacing $\hat{a}\hat{b}$ with the Winsorized mean in Equation 13. Equations 14 through 17 were then used to calculate the bias-corrected confidence interval just as they were for the BCBCI. Ultimately, Chen and Fritz (2021) found that there was always a tradeoff

between type I error rate and power, with power and type I error rate both decreasing as percentage of Winsorization increased. No method exceeded the BCBCI in terms of power, no method had better control over the type I error rate than the PBCI, and all methods had balance levels that fell between those of the BCBCI and the PBCI (the BCBCI with the best balance overall and PBCI with the worst).

Instead of changing the measure of central tendency used in the BCBCI's bias estimation as Chen and Fritz (2021) did, the present study proposes two bootstrap confidence interval alternatives to the BCBCI based on the assumptions made to appropriately apply the BCBCI to the indirect effect: one introduced by Stine (1989) adapted here for use with the indirect effect, and another original method developed by the research team. The performance of these methods are then compared to the BCBCI, the PBCI, and the WBCBCI proposed by Chen and Fritz (2021) in a simulation study.

Alternatives to the BCBCI for the Indirect Effect

The alternative bias-corrected bootstrap methods discussed in this section target assumptions underlying the BCBCI's correction that may prove fallible when it is applied to the indirect effect. The first method, the significance-tested bias-corrected bootstrap confidence interval (stBCBCI), implements a significance test to determine whether differences observed between the sample indirect effect and the median of the bootstrap sampling distribution are large enough to warrant the use of a bias correction. The second method, the reduced bias-corrected bootstrap confidence interval (rBCBCI), takes bias properties of the sample indirect effect into account when its bias correction is applied, resulting in a bias-corrected confidence interval tailored specifically to the statistical properties of the sample indirect effect.

Significance-Tested Bias-Corrected Bootstrap Confidence Interval

Since $\hat{z}_{adj}$ is only an estimate of z_{adj} , it is always possible when $z_{adj} = 0$ to still observe a nonzero value of $\hat{z}_{adj}$ in a sample due to random variability. If $z_{adj} = 0$ and $\hat{z}_{adj}$ is nonzero, due solely to random error, the BCBCI will inappropriately correct for bias that does not actually exist. The result is a BCBCI that inaccurately shifts the

confidence limits of a PBCI that was already well-centered around the sample indirect effect and the true indirect effect, moving these effects closer to one of the two confidence limits. If the true indirect effect equals zero, zero would be closer to one of the two limits, resulting in a greater chance of the BCBCI excluding the true parameter value of zero and committing a type I error than the unadjusted PBCI.

To protect against such scenarios as the one described above, Stine (1989) recommended performing a test to ensure the median bias in the bootstrap sampling distribution was statistically significant, rather than just due to random variability, before applying the BCBCI's bias correction. In his method, the outcome of a binomial test with null hypothesis

$$H_0 : \frac{\#\hat{a}\hat{b}^* < \hat{a}\hat{b}}{B} = .5$$

and two-tailed alternative hypothesis

$$H_a : \frac{\#\hat{a}\hat{b}^* < \hat{a}\hat{b}}{B} \neq .5$$

determines what bootstrap confidence interval method to employ: If the test is statistically significant at an α -level of .05, the BCBCI is used because the proportion of bootstrap indirect effect estimates less than the original sample estimate is significantly different from .50. On the other hand, if the test is not significant, the unadjusted PBCI is used instead because it cannot be concluded that $\hat{a}\hat{b}$ significantly differs from the median of the bootstrap sampling distribution. The α -level used to determine the significance of the median bias here can be different from the α -level used to determine the confidence level of the bootstrap confidence intervals. Stine (1989) applied this stBCBCI to the sample correlation, while in this study it is applied to the sample indirect effect.

Reduced Bias-Corrected Bootstrap Confidence Interval

While the stBCBCI is meant to protect against applying a bias correction if no bias is present, what about if bias *is* present but the BCBCI's bias correction is too large? As an asymptotically unbiased estimator, the sample indirect effect is guaranteed

to be unbiased when the sample size grows large (e.g., Sobel, 1982). Even in small samples, MacKinnon, Warsi, and Dwyer (1995) empirically showed that the bias of the indirect effect is small. Since the mean of the sample indirect effect equals its median when the true indirect effect is zero, median bias should be similarly small when there is no indirect effect to detect (Craig, 1936). Thus, when the true indirect effect is zero, it is possible that the BCBCI overcorrects for bias and shifts the PBCI's confidence limits so much that they exclude zero, resulting in the inflated type I error rates seen for the BCBCI. To prevent such overcorrections, it may be more appropriate to modify Equations 16 and 17:

$$\begin{aligned}\text{Lower: } & \Phi\left(\hat{z}_{adj} + z_{\frac{\alpha}{2}}\right) \\ \text{Upper: } & \Phi\left(\hat{z}_{adj} + z_{1-\frac{\alpha}{2}}\right).\end{aligned}$$

Thus, $\hat{z}_{adj}$ is added to the z -scores corresponding to the $\frac{\alpha}{2} \times 100$ th and $(1 - \frac{\alpha}{2}) \times 100$ th percentiles of the standard normal distribution only once instead of twice, reducing how much the corrected limits are shifted. These modified limits form the $(1 - \alpha) \times 100\%$ rBCBCI for the indirect effect.

The following section describes a simulation study that compared the stBCBCI and the rBCBCI to the PBCI, the BCBCI, and the 30% Winsorized bias-corrected bootstrap confidence interval (WBCBCI) proposed by Chen and Fritz (2021).

Simulation

Manipulated Factors

To compare the performance of the five bootstrap confidence interval methods listed above, a Monte Carlo simulation was conducted. In addition to the bootstrap method used, three other factors were varied in the simulation: effect size of the a -path, effect size of the b -path, and sample size. First, to see how the bootstrap methods performed in the presence of a variety of different effect sizes, the sizes of the a - and b -paths were varied across conditions to be 0.00 (null effect), 0.16 (small effect), 0.39 (medium effect), and 0.59 (large effect) in accordance with effect size conventions

established by Cohen (1988). These four different a -path sizes and four different b -path sizes then multiplied to form a total of 16 different path-size combinations.

Second, since previous studies have shown that the BCBCI has inflated type I error rates with samples of $n < 500$, the bootstrap methods were tested with five different sample sizes: 25, 50, 75, 100, and 500. This resulted in a total of 80 conditions. Each condition was run 1000 times, resulting in a total of 80,000 iterations run in the simulation.

Data Generation

The Monte Carlo simulation was coded and run in R Version 4.0.2 (R Core Team, 2020). For each iteration, values of X were randomly generated from a standard normal distribution (mean of zero and standard deviation of one), with the number of X values drawn determined by the sample size factor. These X values were then plugged into Equation 2 and summed with error terms randomly drawn from the standard normal distribution to generate corresponding values of M . The intercept of the equation was set to zero, and the a -path was set to one of the four effect size levels listed above. After generating values of M using Equation 2, Equation 3 was used with the M and X values to generate the corresponding Y values. Once again, the equation's error term was added by randomly selecting a value from the standard normal distribution and the intercept was set to zero. The size of the b -path was determined by the size of the b -path for that condition, and the c' -path was set to zero since this path has not been found to have any effect on bootstrap confidence interval approaches for the indirect effect (Fritz & MacKinnon, 2007; MacKinnon et al., 2004).

Once values for X , M , and Y were obtained, ordinary least squares regression was applied to solve Equations 2 and 3 for $\hat{a}$ and $\hat{b}$, the least squares estimators of a and b . These estimates were then multiplied together to get $\hat{a}\hat{b}$, the original sample estimate of the indirect effect. The simulation then drew samples of the same size as the original sample with replacement from the generated X , M , and Y values 5000 times, calculating the indirect effect in each resample to form 5000 bootstrap estimates of the

indirect effect. These estimates were then ordered from smallest to largest, and the PBCI, BCBCI, stBCBCI, rBCBCI, and WBCBCI procedures were used to calculate 95% confidence intervals for the indirect effect to compare the methods in terms of type I error rate, power, and balance.

Measured Outcomes

For each condition, **sample bias** was calculated as

$$\frac{\sum_{j=1}^{1000} \hat{ab}_j}{1000} - ab,$$

where $\hat{ab}_j$ is the sample indirect effect estimate in iteration j and 1000 is the total number of iterations. Also, **bootstrap bias** was calculated for each condition as

$$\frac{\sum_{j=1}^{1000} \left(\frac{\sum_{k=1}^{5000} \hat{ab}_{kj}^*}{5000} - \hat{ab}_j \right)}{1000},$$

where $\hat{ab}_{kj}^*$ is the bootstrap indirect effect estimate from bootstrap sample k in iteration j , and 5000 is the total number of bootstrap indirect effect samples drawn in iteration j . These measures were collected to monitor the actual mean bias present across conditions.

The performance of the methods included in this simulation study was measured using the type I error rates, power, and balance of the bootstrap confidence intervals. In the seven conditions where either the a -path, the b -path, or both were equal to zero (meaning the true indirect effect was zero), the number of times a confidence interval excluded zero was tallied up and divided by the total number of iterations to get the percentage of times zero was excluded by each method (i.e., the method's **type I error rate**). The sizes of the type I error rates were assessed using Bradley's liberal robustness criterion which, with $\alpha = .05$, resulted in an interval from .025 to .075 (Bradley, 1978). If the type I error rate of a confidence interval fell outside of this interval, it was deemed to have either an inflated (if it fell above the interval) or a conservative (if it fell below the interval) type I error rate. The **power** of each method was calculated exactly the same way as the type I error rate except in the conditions where both the a -path and the b -path were nonzero.

For each bootstrap method within each condition, **balance** was calculated by recording the number of times the confidence interval was either above or below the true indirect effect (i.e., did not capture the true indirect effect). These tallies were then counted across iterations to give the total number of confidence intervals that fell above the true effect and the total number that fell below the true effect in each condition. Finally, the total number of confidence intervals that fell above the true effect was divided by the total number of confidence intervals that excluded the true effect (i.e., the number above the true effect plus the number below it). If this value was equal to .50, the method was perfectly balanced, and the farther the value was from .50 the more imbalance was present. Values greater than .50 indicated that the confidence interval fell above the true effect more often than it fell below, and values smaller than .50 indicated the confidence interval fell below the true effect more often than it fell above. A binomial significance test was also applied to the balance values to see which were significantly different from .50 at an α -level of .05.

Results

The following subsections summarize the results of the simulation study in terms of the measured outcome variables using a series of figures. Each figure contains multiple graphs, and each graph pertains to an *a*-path-*b*-path combination. Within each graph, the outcome variable is plotted on the *y*-axis and the natural logarithm of the sample size is plotted on the *x*-axis. The outcome variables are discussed in the order bias, type I error rate, power, and then balance.

Bias

Figure 2 displays the sample bias and bootstrap bias in each condition. The largest sample biases were observed with a sample size of 25 and a nonzero true indirect effect, with the largest positive bias of 0.013 obtained when $a = b = 0.59$, and the largest negative bias of -0.009 obtained when $a = b = 0.39$. With nonzero true indirect effects, sample bias appeared to be the worst at the smallest two or three sample sizes when the size of one or both the paths was larger (0.39 or 0.59). A notable exception to

this pattern occurred when $a = .59$ and $b = .14$, where the sample bias remained close to zero across all sample sizes. When $a = 0$, the sample bias remained near zero regardless of the size of b or the sample size, and when $b = 0$ and the sample size was 25, the sample bias grew as the size of a increased. Overall, the sample bias shrank towards zero in all effect size conditions as sample size increased. Bootstrap bias, on the other hand, remained near zero across all conditions.

Type I Error Rate

Figures 3 and 4 display the type I error rates of the methods when $a = 0$ and $b = 0$, respectively. The black horizontal line at .05 on these graphs represents the target type I error rate determined by the α -level of .05, and the grey shaded region indicates Bradley's liberal robustness criterion. The first notable feature of these graphs is that, regardless of condition, the order of the methods in terms of type I error rate almost always remained the same, with the BCBCI having the highest type I error rate followed by the stBCBCI, the rBCBCI, the WBCBCI, and finally the PBCI with the lowest type I error rate. The only time this order ever changed occurred when $b = 0.59$ and the sample size was set at 500. In this condition, the rBCBCI's type I error rate of .051 exceeded the stBCBCI's type I error rate of .048.

When $a = b = 0$, all methods were too conservative regardless of sample size, and the WBCBCI and PBCI attained the lowest observed type I error rate of 0 when $n = 25$. All other instances of the methods' type I error rates falling below the lower limit of Bradley's liberal robustness criterion occurred when either the a -path or the b -path was equal to 0.14 and the sample size was less than or equal to 100. The WBCBCI and the PBCI had type I error rates that were too conservative in the highest number of conditions, with their type I error rates falling below the lower limit of Bradley's liberal robustness criterion in 13 of the 35 conditions where $ab = 0$. On the other hand, the BCBCI and the stBCBCI had type I error rates that were too liberal in the highest number of conditions, with their type I error rates falling above the upper limit of Bradley's liberal robustness criterion in 8 of the 35 conditions. The only other

method with a type I error rate that exceeded the upper limit of the criterion was the rBCBCI, which had a type I error rate of 0.077 when $b = 0.59$ and $n = 100$. The maximum observed type I error rate of .088 was exhibited by both the BCBCI and the stBCBCI when $b = 0.39$ and the sample size was 100 and again by the BCBCI when $a = 0.59$ and $n = 25$.

To summarize type I error rate trends, the type I error rates of all methods were too conservative when $a = b = 0$ in every sample size condition. When the size of the a -path or b -path was .14, type I error rates increased as sample size increased overall. The relationship between type I error rate and sample size was noticeably nonmonotonic when the nonzero path was .39 or .59, with type I error rates increasing until about $n = 75$ or $n = 100$ before they decreased to near the target type I error rate of .05 by $n = 500$. Regardless of the condition, however, the order of the methods almost always remained the same, with the order from highest type I error rate to lowest being: BCBCI, stBCBCI, rBCBCI, WBCBCI, and PBCI.

Power

Figure 5 displays the power of all the methods. As expected, empirical power increased as both sample size and the sizes of the a - and b -paths increased, reaching one when a and/or b was 0.39 or 0.59 and the sample size was 100 or 500. With rare exceptions (where the difference in power was at most .03), the order of the methods remained the same across all conditions, with the BCBCI having the highest power followed by the stBCBCI, rBCBCI, WBCBCI, and finally the PBCI with the lowest power.

Balance

Figure 6 displays the balance of the 95% bootstrap confidence interval methods used. Recall that a value of .50 indicates that exactly half of the true indirect effects not captured by the confidence interval were below the lower limit, and so the confidence interval was perfectly balanced. Thus, the black horizontal line at .50 on the graphs indicates perfect balance. A point above the line indicates a condition in which

the confidence interval fell above the true effect more often than it fell below, and a point below the line indicates a condition in which the confidence interval fell below the true effect more often than it fell above. Note that there is no data point for either the PBCI or the WBCBCI in the $a = b = 0$ graph at the sample size of 25 because these methods captured every true indirect effect in this condition.

With the exception of the conditions in which $a = b = 0$ and $a = b = .14$, the balance of all methods seemed to converge to the same value as sample size grew to 500. The worst balance occurred when $a = b = 0$, particularly when the sample size was 50 or 100 for the rBCBCI and the WBCBCI or when the sample size was 50, 75, or 100 for the PBCI. At these sample sizes, every time these confidence intervals failed to capture the true indirect effect, they fell above the true effect, but this was largely due to the low number of true indirect effects excluded at these sample sizes. For example, the rBCBCI, WBCBCI, and PBCI only excluded zero a single time when the sample size was set to 100. In fact, the WBCBCI and PBCI never excluded zero more than four times for a given sample size when $a = b = 0$.

The BCBCI and stBCBCI were more balanced than the other three methods when $a = b = 0$ and the sample size was 50 or 100, and when the sample size was 100 the BCBCI and the stBCBCI were perfectly balanced (falling above zero three times each and falling below zero three times each). In fact, the BCBCI and the stBCBCI were the best balanced methods overall, achieving the proportions closest to .50 in a combined total of 49 of the 80 conditions (one time both tying with the PBCI). The rBCBCI and WBCBCI, which tied in two conditions, had the best balance in 14 and 8 conditions, respectively. Finally, the PBCI had the best balance in 12 conditions, tying once with both the BCBCI and stBCBCI and one other time with the WBCBCI.

Besides the conditions in which $a = b = 0$, the patterns of balance were similar across methods, with the BCBCI and stBCBCI most often closest to perfectly balanced followed by the rBCBCI, PBCI, and WBCBCI. Examining only the proportions that were significantly different from .50 in Figure 7 reveals the same pattern of balance, with the BCBCI significantly different in only 16 conditions, the stBCBCI in 17, the

rBCBCI in 18, the PBCI in 25, and the WBCBCI in 26.

Discussion

The results of the simulation study showed that, for each performance measure, the examined bootstrap methods fell on a continuum. From lowest type I error rate and power to highest type I error rate and power, the order went PBCI, WBCBCI, rBCBCI, stBCBCI, and BCBCI. The order from worst balance to best balance was the same with the exception that the positions of the WBCBCI and PBCI were switched. One notable feature of the simulation was that the type I error rates of all methods often surpassed the $\alpha = .05$ level that corresponds to a 95% confidence interval. Since power is a function of α , increasing as the α -level increases, power may be artificially inflated by these liberal type I error rates, and thus comparing the power of these methods without consideration for their type I error rates may be misleading. To address this issue, a supplementary simulation was run that compared the power of each bias-corrected method to a corresponding PBCI set to the maximum type I error rate achieved by the bias-corrected method.

Supplementary Power Simulation

After the primary simulation was complete, the maximum type I error rate found for each of the four 95% bias-corrected confidence intervals was recorded. In this supplementary simulation, the maximum type I error rate achieved by each bias-corrected method was used to determine the α -level at which to set a $(1 - \alpha) \times 100\%$ PBCI for comparison. For example, the maximum type I error rate found for the rBCBCI in the primary simulation was .077, and so a $(1 - .077) \times 100 = 92.3\%$ PBCI was calculated to compare statistical power with the rBCBCI in this simulation. Setting the PBCI (the method with the lowest observed power in the primary simulation) to a confidence level corresponding to the maximum type I error rate achieved by each bias-corrected method should offer a worst-case-scenario comparison, ensuring that any power differences due to type I error rate inflation are removed even if each method is at its most liberal. Using the

maximum type I error rate of each method should be a more informative option than using the average type I error rate of each method; the overly conservative type I error rates the methods achieved in the primary simulation's $a = b = 0$ conditions brought their mean error rates down to almost nominal levels, resulting in comparison PBCIs that would be set near a 95% confidence level like the PBCI in Figure 5.

Manipulated Factors

The manipulated factors in this simulation were exactly the same as those in the primary simulation, but the conditions in which $ab = 0$ are not focused on since the outcome of interest in this simulation was power.

Data Generation

The exact same procedure was followed to generate data in this simulation as was used in the primary simulation, except four different PBCIs were formed to be compared with the four 95% bias-corrected bootstrap confidence interval methods. The confidence level of each PBCI corresponded to the maximum type I error rate its corresponding bias-corrected method(s) reached during the first simulation. These type I error rates were .088 for the BCBCI and stBCBCI (resulting in a 91.2% comparison PBCI), .077 for the rBCBCI (resulting in a 92.3% comparison PBCI), and 0.073 for the WBCBCI (resulting in a 92.7% comparison PBCI).

Measured Outcome

The power of the bootstrap methods was calculated the same way it was in the primary simulation.

Supplementary Results

Figures 8 presents the BCBCI, stBCBCI, rBCBCI, and WBCBCI plotted with their comparison PBCIs that control for the inflated type I error rates of each method. As can be seen in the figure, the advantages in terms of power of the bias-corrected methods over the PBCI that were evident in Figure 5 are no longer present, with the

adjusted PBCIs achieving higher power than their corresponding bias-corrected methods in all but six conditions (in which both the BCBCI and stBCBCI obtained higher power, but the difference in power was never more than .005). Thus, controlling for type I error rate, the bias-corrected methods did not seem to perform any better than the PBCI.

General Discussion

The present study compared the PBCI, BCBCI, and WBCBCI to two new alternative bias-corrected bootstrap techniques for the indirect effect: the stBCBCI and the rBCBCI. Performance measures included type I error rate, power, and balance. In line with previous research, the BCBCI exhibited the most elevated type I error rates and the highest power, and the PBCI had the most control over type I error rates and the lowest overall power. In terms of balance, the results of this study agree with those of Chen and Fritz (2021) who also found that the BCBCI offered overall better balance than other bootstrap methods. Across all conditions, the stBCBCI performed very similarly to the BCBCI on all performance criteria, and the WBCBCI performed quite similarly to the PBCI. Increasing the percentage of trimming should make the WBCBCI more and more similar to the PBCI, with the 50% Winsorized BCBCI producing the exact same confidence intervals as the PBCI (Chen & Fritz, 2021). The rBCBCI, on the other hand, offered a balance of the benefits and shortcomings of the BCBCI and the PBCI, falling somewhere in the middle on all performance criteria.

Although the goal of this research was to develop a method that maintained the high power and balance of the BCBCI while still controlling the type I error rate, the primary simulation revealed that there was always a tradeoff between the three performance measures. Across conditions, the methods examined fell on a continuum, which in order of increasing type I error rate, power, and balance went: PBCI, WBCBCI, rBCBCI, stBCBCI, and BCBCI (with the exception that the WBCBCI had slightly worse balance than the PBCI). Thus, a method with overall better balance and higher power also had higher type I error rates than its competitors. Decreasing the

confidence level of the PBCI in the supplementary simulation seemed to provide the same power benefit using any of the bias-corrected methods did, however, and thus the only clear advantage of the bias-corrected methods is the better balance they provide. As such, if controlling for false positives is more important than finding an effect if it exists, the PBCI is still recommended. If detecting an existing effect is most important but control over the type I error rate is still a concern, the rBCBCI is a good compromise between the PBCI and the BCBCI that has the benefit of being better balanced than the former method. These decisions are only impactful when the sample size is small, however; in almost all effect size conditions in the primary simulation, differences between methods on all three performance measures decreased with sample size to the point that, with $n = 500$, the performance of each method had converged to nearly the same value.

Regardless of condition, the performance of the stBCBCI closely resembled the performance of the BCBCI in terms of type I error rate, power, and balance. This was due to how often the stBCBCI's significance test indicated a significant difference between the median of the bootstrap sampling distribution and the corresponding sample indirect effect estimate. On average, this significance test rejected the null hypothesis in over 72% of the iterations in each condition: about 71% of the time when the true indirect effect was zero and about 74% of the time when the true indirect effect was nonzero. Thus, the stBCBCI used the bias-corrected bootstrap for its confidence interval most of the time. Since the BCBCI was used almost three times more often than the PBCI, it makes sense that the stBCBCI performed very similarly to the former test. In fact, with 5000 bootstrap replications representing the sample size of each trial, the binomial test of the stBCBCI would reach significance at an α -level of .05 if the observed proportion was greater than 0.0138 away from 0.50, and so even minor deviations away from the median of the bootstrap distribution resulted in the BCBCI being employed. Using a more conservative α -level (e.g., $\alpha = 0.05$) may be worthwhile to increase the significance threshold.

Still, the frequent significance of the stBCBCI's test of median bias highlights the

fact that median and mean bias are distinct quantities. While the discrepancy between the sample indirect effect and the median of the bootstrap sampling distribution was often large enough to warrant the use of the stBCBCI, bootstrap bias (i.e., mean bias of the bootstrap indirect effect) recorded during the simulation remained near zero across all conditions. In 76 of the 80 conditions observed in this study, bootstrap bias was smaller than sample bias (i.e., mean bias of the sample indirect effect). Further exploration of the differences between the mean and median bias properties of the indirect effect, and how well they agree with the bias properties assumed under the bias correction of the BCBCI, could advance our understanding of the BCBCI's appropriateness when applied to the indirect effect. Diagnostic functions like those discussed in Efron (1982b) could also help assess whether the existence of the function $g(\cdot)$ is a reasonable assumption for the indirect effect.

Limitations and Future Directions

The following section discusses several limitations with the current study, including the conclusions made regarding balance, the lack of missing data, the lack of confounding variables, the use of a simple mediation model within the ordinary least squares (OLS) regression framework, and the use of 95% confidence intervals. Potential future directions are also described to address these issues and other questions that remain regarding the BCBCI.

The worst balance among the methods examined in this study was observed in conditions in which $a = b = 0$. However, these were also the conditions in which the fewest number of true indirect effects were excluded by the confidence intervals. For example, in the $a = b = 0$ and $n = 100$ condition, the rBCBCI, PBCI, and WBCBCI never had any confidence intervals fall below the true indirect effect because only one true effect was ever excluded by these bootstrap methods across all 1000 iterations of the condition. Thus, the confidence intervals never fell below the true effect, but they also only fell above it once. Due to very small numbers of indirect effects being excluded by confidence intervals in this condition and others in which $ab = 0$, more simulations

should be run to verify the accuracy of the balance values found in this study.

Future simulation studies should also examine the impact of missing data on these methods. In real-world scenarios, it is quite common for some data values to be missing, and thus methods of handling missing data must be applied in conjunction with bootstrapping to draw accurate inferences about the data. Future research should combine the bootstrap methods discussed here with popular missing data methods in the presence of different missing data mechanisms to see how their performances change and whether they offer any performance advantages over the PBCI in such situations. For existing work on missing data handling procedures combined with bootstrapping for the indirect effect, see Wu and Jai (2013) and Zhang and Wang (2013).

Another real-world complication ignored in this study is the presence of confounding variables. Every mediation model used in the simulations was correctly specified, meaning that the form of the model fit to the data matched the underlying relationships between the variables in the population. These are conditions in which any bias present should be at a minimum. With an unknown third variable present, the indirect effect produced by the mediation model may be more biased (see Valente, Pehlam III, Smyth, and MacKinnon, 2017 for an overview of confounding and how to address it in mediation analyses), and thus the bias-corrected methods may perform differently. It would be informative to see how the BCBCI and the other bias-correction variants perform in relation to the PBCI when the indirect effect estimate is (median) biased by a confounding variable in future work.

In addition to being correctly specified, all models included in the simulation studies were also simple mediation models estimated using OLS regression rather than the also-popular structural equation modeling (SEM) framework. Because OLS regression mediation models are mathematically equivalent to SEM mediation models when the mediator and outcome variables are continuous and observed, however, these results are expected to generalize to SEM mediation models (e.g., see Rijnhart, Twisk, Chinapaw, de Boer, and Heymans, 2017). For a general introduction to SEM in mediation analysis, see Gunzler, Chen, Wu, and Zhang, 2013.

The confidence intervals constructed using the OLS framework in this study were all set at the 95% confidence level. This raises the question of whether these results would generalize to other confidence levels, such as 80% or 90%. Previous research has already compared methods for constructing confidence intervals for the indirect effect set at different confidence levels, however, and the findings show that the methods perform similarly relative to each other regardless of the confidence level (MacKinnon et al., 2004). Thus, these results should generalize to other confidence levels.

The inflated type I error rates of the BCBCI applied to the indirect effect call into question its ability to perform in areas of statistics outside of mediation analysis as well. For example, Karlsson (2009) found evidence that the BCBCI applied to a weighted nonlinear quantile regression estimator for longitudinal data resulted in significant undercoverage—i.e., the true value being captured by the interval less often than it should according to the set confidence level—and recommended using the PBCI instead. Identifying and examining other BCBCI problem areas such as this in future studies can help further understanding of the BCBCI's issues and lead to the development of improved bias-corrections for the sample indirect effect and other statistics. Applying the new bias-corrected methods discussed in this study to other statistics would be worthwhile as well to see how they perform in these other areas.

Although—in correctly specified, complete data mediation analyses at least—the benefits of the bias correction are still accompanied by inflated type I error rates after the completion of this study, it is possible that there is still an alteration to the BCBCI's bias correction that can wed its increased power with the PBCI's control over the type I error rate. Thus, future research should continue examining ways in which the BCBCI can be tailored to the indirect effect in small samples to make it the robust alternative to the PBCI it was once believed to be.

Funding

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2034835. Any opinions,

findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Acknowledgements

The authors would like to thank all members of the QRClab for providing feedback on early drafts of this manuscript. Additional thanks to Dr. Han Du for her valuable insights that aided in the development of the alternative bias corrections discussed in this study and for her feedback on a later draft of the manuscript.

References

- Aroian, L. A. (1947). The probability function of the product of two normally distributed variables. *Annals of Mathematical Statistics*, 10, 265–271.
- Aroian, L. A., Taneja, V. S., & Cornwell, L. W. (1978). Mathematical forms of the distribution of the product of two normal variables. *Communications in Statistics: Theory and Methods*, A7, 165–172.
- Biesanz, J. C., Falk, C. F., & Savalei, V. (2010). Assessing mediational models: Testing and interval estimation for indirect effects. *Multivariate Behavioral Research*, 45, 661–701.
- Bradley, J. V. (1978). Robustness? *British Journal of Mathematical and Statistical Psychology*, 31, 144–152.
- Brown, G. W. (1947). On small-sample estimation. *The Annals of Mathematical Statistics*, 18(4), 582–585.
- Chen, D., & Fritz, M. S. (2021). Comparing alternative corrections for bias in the bias-corrected bootstrap test of mediation. *Evaluation and the Health Professions*.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd). Hillsdale: NJ:NJ: Lawrence Erlbaum Associates.
- Craig, C. C. (1936). On the frequency function of xy . *The Annals of Mathematical Statistics*, 7(1), 1–15.
- Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1), 1–26.
- Efron, B. (1982a). Nonparametric confidence intervals. In *The jackknife, the bootstrap and other resampling plans* (pp. 75–90). SIAM.
- Efron, B. (1982b). Transformation theory: How normal is a family of distributions? *The Annals of Statistics*, 10(2), 323–339.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Boca Raton, FL: Chapman & Hall.
- Fritz, M. S., & MacKinnon, D. P. (2007). Required sample size to detect the mediated effect. *Psychological Science*, 18(3), 233–239.

- Fritz, M. S., Taylor, A. B., & MacKinnon, D. P. (2012). Explanation of two anomalous results in statistical mediation analysis. *Multivariate Behavioral Research*, 47(1), 61–87.
- Götz, M., O’Boyle, E. H., Gonzalez-Mulé, E., Banks, G. C., & Bollmann, S. S. (2021). The “goldilocks zone”: (too) many confidence intervals in tests of mediation just exclude zero. *Psychological Bulletin*, 147(1), 95–114.
- Gunzler, D., Chen, T., Wu, P., & Zhang, H. (2013). Introduction to mediation analysis with structural equation modeling. *Shanghai Archives of Psychiatry*, 25(6), 390–394.
- Hayes, A. F., & Scharkow, M. (2013). The relative trustworthiness of inferential tests of the indirect effect in statistical mediation analysis: Does method really matter? *Psychological Science*, 24(10), 1918–1927.
- Karlsson, A. (2009). Bootstrap methods for bias correction and confidence interval estimation for nonlinear quantile regression of longitudinal data. *Journal of Statistical Computation and Simulation*, 79(10), 1205–1218.
- MacKinnon, D. P., Lockwood, C. M., Hoffman, J. M., West, S. G., & Sheets, V. (2002). A comparison of methods to test mediation and other intervening variable effects. *Psychological Methods*, 7(1), 83–104.
- MacKinnon, D. P., Lockwood, C. M., & Williams, J. (2004). Confidence limits for the indirect effect: Distribution of the product and resampling methods. *Multivariate Behavioral Research*, 39, 99–128.
- MacKinnon, D. P., Warsi, G., & Dwyer, J. H. (1995). A simulation study of mediated effect measures. *Multivariate Behavioral Research*, 30(1), 41.
- Needham, T. (1993). A visual explanation of Jensen’s inequality. *American Mathematical Monthly*, 100(8), 768–771.
- Osberg, T. M., & Eggert, M. (2012). Direct and indirect effects of stress on bulimic symptoms and bmi: The mediating role of irrational food beliefs. *Eating Behaviors*, 13(1), 54–57.

- Preacher, K. J., & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, *40*, 879–891.
- R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Rijnhart, J. J., Twisk, J. W., Chinapaw, M. J., de Boer, M. R., & Heymans, M. W. (2017). Comparison of methods for the analysis of relatively simple mediation models. *Contemporary Clinical Trials Communications*, *7*, 130–135.
doi:<https://doi.org/10.1016/j.conctc.2017.06.005>
- Sobel, M. E. (1982). Asymptotic confidence intervals for indirect effects in structural equation models. *Sociological Methodology*, *13*, 290–312.
- Stine, R. (1989). An introduction to bootstrap methods: Examples and ideas. *Sociological Methods and Research*, *18*(2), 243–291.
- Valente, M. J., Pehlam III, W. E., Smyth, H., & MacKinnon, D. P. (2017). Confounding in statistical mediation analysis: What it is and how to address it. *J Couns Psychol*, *64*(6), 659–671.
- Williams, J., & MacKinnon, D. P. (2008). Resampling and distribution of the product methods for testing indirect effects in complex models. *Structural Equation Modeling*, *15*, 23–51.
- Wu, W., & Jai, F. (2013). A new procedure to test mediation with missing data through nonparametric bootstrapping and multiple imputation. *Multivariate Behavioral Research*, *48*(5), 663–691.
- Zhang, Z., & Wang, L. (2013). Methods for mediation analysis with missing data. *Psychometrika*, *78*(1), 154–184.

Figure 1. Simple Mediation Model

Figure 2. Sample bias and average bootstrap sample bias across the range of a -path sizes, b -path sizes, and sample sizes.

Figure 3. Type I error rate of all methods when a -path is zero across the range of b -path sizes and sample sizes.

Figure 4. Type I error rate of all methods when b -path is zero across the range of a -path sizes and sample sizes.

Figure 5. Power of all methods across the range of a -path sizes, b -path sizes, and sample sizes.

Figure 6. Balance of all methods across the range of a -path sizes, b -path sizes, and sample sizes.

Figure 7. Balance of all methods that are significant at $\alpha = 0.05$ across the range of a -path sizes, b -path sizes, and sample sizes.

Figure 8. Power of all methods compared to PBCI controlling for type I error rate, across the range of a -path sizes, b -path sizes, and sample sizes.