**Selecting a Within- or Between-Subject Design for Mediation: Validity, Causality, and
Statistical Power**

Amanda K. Montoya

University of California, Los Angeles

In Press at *Multivariate Behavioral Research*

Author Note

Amanda Kay Montoya, Department of Psychology, UCLA.

Correspondence concerning this article should be addressed to Amanda Montoya,
Department of Psychology, University of California, Los Angeles, 502 Portola Plaza, Los
Angeles, CA 90095. Email: akmontoya@ucla.edu.

Within- or Between-Subject Design for Mediation

Selecting a Within- or Between-Subject Design for Mediation: Validity, Causality, and Statistical Power

Abstract

Researchers with mediation hypotheses must consider which design to use: within-subject or between-subject? In this paper, I argue that three factors should influence design choice: validity, causality, and statistical power. Threats to validity include carry-over effects, participant awareness, measurement, and more. Causality is a core element of mediation, and the assumptions required for causal inference differ between the two designs. Between-subject designs require more restrictive no-confounder assumptions, but within-subject designs require the assumption of no carry-over effects. Statistical power should be higher in within-subject designs, but the degree and conditions of this advantage are unknown for mediation analysis. A Monte Carlo simulation compares designs under a broad range of sample sizes, effect sizes, and correlations among repeated measurements. The results show within-subject designs require about half the sample size of between-subject designs to detect indirect effects of the same size, but this difference can vary with population parameters. I provide an empirical example and R script for conducting power analysis for within-subject mediation analysis. Researchers interested in conducting mediation analysis should not select within-subject designs merely because of higher power, but they should also consider validity and causality in their decision, both of which can favor between-subject designs.

Keywords: Mediation analysis, indirect effect, within-subject design, power, type I error, Monte Carlo simulation, power analysis, validity, causal inference.

Within- or Between-Subject Design for Mediation

Mediation analysis examines the indirect effect of some causal antecedent variable (X) on an outcome (Y) through a proposed intermediary variable (M , a mediator). Mediation analysis is exceedingly common in psychology research and other behavioral sciences (Hayes & Scharkow, 2013; Yzerbyt et al., 2018). This analysis is particularly useful for understanding the mechanisms by which a manipulation may influence an outcome. For example, clinical psychologists may want to understand the psychological mechanisms by which behavioral therapies impact psychological disorders. Education researchers may want to understand the mechanisms by which certain educational practices influence student learning. Understanding mechanisms is central to many research questions throughout psychological and behavioral research, and mediation analysis can provide insight into these mechanisms when appropriate assumptions with respect to causality are met.

The recent development of a path analytic approach to conducting mediation analysis for two-condition, within-subject designs (Montoya & Hayes, 2017), as well as the release of a free macro for SPSS and SAS to conduct these analyses, MEMORE (available at akmontoya.com), has resulted in a rapid rise in the use of these models in psychology and other behavioral sciences (e.g., Breitborde, Wastler, Pine, & Moe, 2020; Goodboy, Dillow, Knoster, & Howard, 2021; Magson, Freeman, Rapee, Richardson, Oar, & Fardouly, 2021; Nguyen, Carnevale, Scholer, Miele, & Fujita, 2019; Pazda & Thostenson, 2019; Thorstenson, Pazda, & Lichtenfeld, 2020). With this sudden increase in adoption, it is important to understand the impact of design on claims of mediation. Researchers interested in testing mediation hypotheses will need to consider specific factors in selecting their design. In this manuscript I describe three primary factors that should weigh on this decision: validity, causality, and statistical power.

Within- or Between-Subject Design for Mediation

This article focuses on two specific experimental designs called the *two-condition, between-subject design* and the *two-condition, within-subject design*. In the two-condition, between-subject design, participants are randomly assigned to levels of X (i.e., 50% of participants are assigned to Condition 1 and 50% are assigned to Condition 2), and for the purposes of mediation analysis each participant is observed once on the proposed mediator M and the proposed outcome Y . This design is very common in psychology and other behavioral sciences, and it is one of the most common designs used for mediation analysis. For comparison, consider the two-condition, within-subject design, which is the repeated-measures equivalent of the two-condition, between-subject design. In the two-condition, within-subject design all participants experience both of two experimental conditions and are observed on the mediator (M) and the outcome (Y) in both conditions. For brevity, the two-condition, between-subject design is referred to as a “between-subject design” and the two-condition, within-subject design is referred to as a “within-subject design” throughout the remainder of the manuscript.

This paper describes the important factors researchers should use to select between a two-condition, between- and within-subject design. These three primary factors are validity, causality, and statistical power. While statistical power is often the focus of the comparison between experimental designs, I argue that validity and causality are also important factors to evaluate in the selection of a design. I begin by describing how to estimate and conduct inference on indirect effects in each design. Next, I summarize common threats to validity that would be particularly important for evaluating between- versus within-subject designs. Depending on the research context, one design may be more valid than another, or in some cases both designs could be equally valid. Next, I describe the assumptions required for causal mediation analysis in each design and compare these assumptions to help researchers evaluate which set of

Within- or Between-Subject Design for Mediation

assumptions are more reasonable for their research case. To evaluate statistical power, a Monte Carlo simulation compares the statistical power of between- and within-subject designs under a variety of realistic conditions. Next, the empirical example shows how to plan sample size for a within-subject mediation analysis and compares it to similar planning for a between-subject design, providing an R script for this process. While statistical power is often front and center in any comparison of between- and within-subject designs, I argue statistical power should be considered secondarily, only after evaluating validity and causality. This paper concludes with recommendations for applied researchers using these methods, methodological contributions of the current research, limitations of the simulation study, and suggestions for future directions.

Estimating Indirect Effects

Between-Subject Designs

This article focuses on studies with only two possible conditions, and therefore X is a dichotomous variable (assumed to be coded with a 1-unit difference between the two conditions, e.g., 0,1 or -0.5, 0.5). M and Y are continuous variables with normally distributed error terms. It is possible for M and Y to be categorical or non-normal variables; however, alternative estimation procedures would be needed, see Preacher (2015) for a review.

Indirect effects in between-subject designs are typically estimated using either linear regression (Ordinary Least Squares; OLS) or structural equation modeling (Maximum Likelihood; ML). The descriptions in this paper will rely on the OLS approach; however, the extension to structural equation models is straightforward, and for observed variable models, the coefficient estimates are the same for OLS and ML approaches (Ledgerwood & Shrout, 2011; Hayes, Montoya, & Rockwood, 2017). A mediation model for a between-subject design can be estimated using two equations: one that estimates the effect of X on M (Equation 1) and one that

Within- or Between-Subject Design for Mediation

estimates the effect of X and M on Y (Equation 2a or 2b). Equation 2a does not allow for an interaction between M and X in the model for Y . If the researchers want to allow for that interaction, they would use Equation 2b.

$$M_i = a_0 + aX_i + e_{M_i} \quad (1)$$

$$Y_i = c'_0 + c'X_i + bM_i + e_{Y_i} \quad (2a)$$

$$Y_i = c'^*_0 + c'^*X_i + b^*M_i + d^*X_iM_i + e_{Y_i^*} \quad (2b)$$

Here X_i , M_i , and Y_i are the observations for the experimental manipulation, the mediator, and the outcome, respectively, for individual i . The intercepts are represented using a_0 and c'_0 or c'^*_0 .

The errors are represented using e_{M_i} and e_{Y_i} or $e_{Y_i^*}$, where e_{M_i} is assumed to be independent of e_{Y_i} or $e_{Y_i^*}$, whichever is used in the estimation of the model. All residuals are assumed to be normally distributed with a mean of zero and unique variance. The effect of the experimental manipulation on the mediator is represented by the coefficient a . In Equation 2a, the effect of the experimental manipulation on the outcome controlling for the mediator is represented by the coefficient c' (the direct effect), and the effect of the mediator on the outcome controlling for the experimental manipulation is represented by the coefficient b . In Equation 2b, c'^* denotes the effect of the experimental manipulation on the outcome when the mediator is zero, b^* denotes the effect of the mediator on the outcome when the experimental manipulation is zero, and d^* denotes the degree to which the relationship between X and Y depends on M and vice versa (i.e., how much the relationship between M and Y depends on X). If X and M are mean centered and X is balanced, then the interpretations of c'^* and b^* should be very close to those of c' and b , since they are conditioned on the average of the other variable. These equations are represented as a single path diagram in Figure 1. Based on the diagram, there are two pathways by which X can influence Y : through the direct effect, c' , and through the indirect effect where X influences M

Within- or Between-Subject Design for Mediation

which then influences Y . The indirect effect is quantified as the product of the a and b paths: ab . In the case of Equation 2b, two separate indirect effects can be estimated, for each condition (corresponding to the *pure natural indirect effect* and the *total natural indirect effect*, See Causality Section). Prior literature has not focused on including the XM interaction while still estimating the indirect effect, so in this study the focus is on the average indirect effect, which is estimated by ab^* when both X and M are mean centered. For the remainder of the manuscript the notation ab for the indirect effect is used, acknowledging that if the researcher estimates the model using Equation 2b their estimate of the indirect effect would be ab^* .

[Figure 1 Here]

An important property of mediation analysis with a continuous mediator and outcome is that the two pathways by which X affects Y are a partition of the total effect, c , which is a quantification of the overall effect of X on Y (without controlling for the mediator). The total effect, c , can be estimated in its own regression equation:

$$Y_i = c_0 + cX_i + e_{Y_i}^\dagger$$

Terms in this equation are defined similarly to Equations 1 and 2. The sum of the indirect and direct effects equals the total effect (at both the sample and population level):

$$c = c' + ab \tag{3}$$

Notably, this is only true when using Equation 2a, and not 2b, though the sums would be close, and this is true at the parameter level for Equation 2b.

The focus of mediation analysis is typically on the presence/absence and magnitude of the indirect effect. There are many inferential methods developed for indirect effects; however, in the main article I focus on the percentile bootstrap confidence interval. Additional inferential

Within- or Between-Subject Design for Mediation

methods are discussed in the Supplementary Materials. Inference for the indirect effect is typically done using a bootstrap confidence interval, due to unusually shaped sampling distribution of the indirect effect (Hayes & Scharkow, 2013). Bootstrap confidence intervals are generated by creating K bootstrap estimates of the indirect effect by sampling from the observed data with replacement, generating K bootstrap samples of size N , where N is the number of cases in the original data and K is recommended to be a large number (e.g. 5,000; Preacher & Hayes, 2004; Shrout & Bolger, 2002). An estimate of the indirect effect is calculated from each bootstrap sample using the regression methods in Equations 1 and 2, generating K estimates of the indirect effect. The distribution of the bootstrapped estimates provides an estimated sampling distribution for the indirect effect (ab). The percentile bootstrap confidence interval is formed by sorting the estimates from smallest to largest and selects the upper and lower bounds for the confidence interval such that $100(\frac{\alpha}{2})\%$ of the bootstrapped estimates lie below the lower bound and $100(1 - \frac{\alpha}{2})\%$ of the bootstrapped estimates lie below the upper bound, where α is the level of the inferential test (e.g., 0.05, Efron & Tibshirani, 1994).

Within-Subject Designs

An equivalent experiment can be run where instead of randomly assigning individuals to one condition, each individual participates in both conditions and is measured on the mediator and the outcome in both conditions. Judd and McClelland (1989) and Judd, Kenny, and McClelland (2001) outlined a procedure for evaluating mediation in such models, and Montoya and Hayes (2017) showed how to represent these models as path models, estimate the indirect effect, and conduct inference using this estimate.

Within- or Between-Subject Design for Mediation

Each individual has an observation on the mediator in both of the two experimental conditions (M_{1i} and M_{2i}) and an observation on the outcome in both of the two experimental conditions (Y_{1i} and Y_{2i}). The base model involves the following equations:

$$M_{1i} = a_1 + e_{M_{1i}} \quad (4)$$

$$M_{2i} = a_2 + e_{M_{2i}} \quad (5)$$

$$Y_{1i} = g_{10} + g_{11}M_{1i} + e_{Y_{1i}} \quad (6)$$

$$Y_{2i} = g_{20} + g_{21}M_{2i} + e_{Y_{2i}} \quad (7)$$

The intercepts a_1 and a_2 are the means of M_1 and M_2 respectively. The intercepts g_{10} and g_{20} are the predicted Y_1 and Y_2 when their respective mediator is zero. The coefficients g_{11} and g_{21} represent the relationship between the mediator and the outcome in each of the two conditions. Equations 4 – 7 assume temporal stability and no carry over, so only the measures from the same condition are used to predict each other and no Y s are used to predict M s (Judd, Kenny, & McClelland, 2001). The error terms ($e_{M_{1i}}, e_{M_{2i}}, e_{Y_{1i}}, e_{Y_{2i}}$) are a combination of both random error and person-level variation. The person-level variability is captured in the covariance between the two repeated-measures (e.g., e_{Y_1} and e_{Y_2}). The residuals are assumed to be normally distributed with a mean of zero and a covariance matrix with the following structure:

$$\Sigma = \begin{bmatrix} \sigma_M & cov_M & 0 & 0 \\ cov_M & \sigma_M & 0 & 0 \\ 0 & 0 & \sigma_Y & cov_Y \\ 0 & 0 & cov_Y & \sigma_Y \end{bmatrix}.$$

Within- or Between-Subject Design for Mediation

An intercept-only model of the difference between the mediators will estimate the influence of the experimental manipulation on the mediators. The intercept, a , is the estimate of the average difference between M_{1i} and M_{2i} .

$$M_{2i} - M_{1i} = a_2 - a_1 + e_{M_{2i}} - e_{M_{1i}} = a + e_{M_i} \quad (8)$$

Similarly, a model for the difference in the outcomes estimates the influence of the mediator on the outcome controlling for experimental condition (b), and the intercept (c') estimates the influence of the experimental manipulation on the outcome controlling for the mediator. Just as in the between-subject case, there are two ways to estimate this model. If the researcher is comfortable assuming the relationship between M and Y is equal across conditions (i.e., there is no condition by mediator interaction, $g_{11} = g_{21}$) it can be estimated using Equation 9a, and b is equal to this relationship ($b = g_{11} = g_{21}$). If this is not a safe assumption to make, Equation 9b should be used, and the b^* coefficient corresponds to the average relationship between M and Y (i.e., $\frac{g_{11} + g_{21}}{2}$).

$$Y_{2i} - Y_{1i} = c' + b(M_{2i} - M_{1i}) + e_{Y_i} \quad (9a)$$

$$Y_{2i} - Y_{1i} = c'^* + b^*(M_{2i} - M_{1i}) + d^*(M_{2i} + M_{1i})^C + e_{Y_i}^* \quad (9b)$$

In the above equation the C indicates that the sum of the mediators has been grand mean centered (i.e., $(M_{2i} + M_{1i})^C = (M_{2i} + M_{1i}) - \overline{(M_{2i} + M_{1i})}$, where $\overline{(M_{2i} + M_{1i})} = \sum_{i=1}^N (M_{2i} + M_{1i}) / N$ is the sample mean of the sum of the mediators). This grand mean centering is what allows the intercept (c'^*) to be interpreted as the direct effect (Judd, Kenny, & McClelland, 2001; Montoya & Hayes, 2017). The coefficient d^* represents the degree to which the relationship between M_1 and Y_1 differs from the relationship between M_2 and Y_2 (i.e.,

Within- or Between-Subject Design for Mediation

$\frac{g_{11}-g_{21}}{2}$)¹. This coefficient is described as a moderation parameter because it reflects the degree to which the M - Y relationship differs across conditions. Similar to the between-subject case, these equations can be represented as a path model (See Figure 2).

[Figure 2 Here]

In the within-subject case, the total effect is estimated by using an intercept-only model to predict the difference between the outcomes:

$$Y_{2i} - Y_{1i} = c + e_{Y_i^\dagger} \quad (10)$$

The total effect as estimated above partitions into the direct effect and indirect effect such that Equation 3 also holds for the within-subject case for both b and b^* .

The focus of inference for mediation is, once again, on the indirect effect, ab (using either b or b^* depending on which version of Equation 9 is used). Estimating bootstrap confidence intervals to a within-subject design uses the same procedure for generating K bootstrap samples, where the observed data are the difference in outcomes, difference in mediators, and sum of the mediators. Equations 8 and 9 are used to estimate the a and b paths, then the product of these paths are calculated to estimate the indirect effect. This procedure will generate K estimates of the indirect effect, and the percentile method to select the upper and lower bounds of the confidence interval can be used to calculate a confidence interval for the indirect effect in a within-subject design.

Evaluating Validity of Between- and Within-Subject Designs

¹ Some presentations of this model use the average of the mediators rather than the sum (e.g., Montoya & Hayes, 2017) in which case the d parameter is equal to the raw difference between the two coefficients $g_{11} - g_{21}$. This is also how MEMORE estimates the model. Inference is not changed by using the sum or the average, it merely scales the coefficient. However, because the simulation uses the “sum” parameterization presented, it is presented in this form here.

Within- or Between-Subject Design for Mediation

Shadish, Cook, and Campbell (2002) define validity as “the truth of, correctness of, or degree of support for an inference.” The concept of validity should underlie all planning of research, but is particularly important when selecting between experimental designs. The inference from a between- or within-subject mediation analysis will be the same (e.g., X indirectly effects Y through M), and so a researcher should select the design that supports this inference most directly. While both within- and between-subject designs have relative pros and cons with respect to validity, within-subject designs seem to introduce more threats to validity. In this section, I overview five important considerations when evaluating the validity of between- and within-subject designs: carry-over effects, participant awareness, ecological validity, level of effect of interest, and changes in measurement properties. Researchers should closely consider these factors when selecting a design. While these issues are by no means new, I focus on important considerations in selecting a design for mediation analysis. Excellent resources that expand on the comparison of validity among between- and within-subject designs include Greenwald (1976), Maxwell, Delaney, & Kelley (2018), Senn (1992), Grice (1966), Rothstein (1974), and Poulton (1973).

Carry-Over Effects

Carry-over effects occur when previous levels of treatments influence variables measured for future levels of treatments. For example, if a researcher were studying how exercise suppresses appetite in the short term, they may randomly assign participants to either exercise or not exercise on Day 1 then immediately report their appetite level. Then on Day 2 participants would complete the other condition (exercise or no exercise) and report appetite immediately after. However, those who exercise on Day 1 may have increased appetite on Day 2 due to their previous exercise. Carry-over effects may impact the validity of within-subject designs by

Within- or Between-Subject Design for Mediation

invalidating the use of a subject as their own control and biasing parameter estimates (Grice, 1966). If these effects cannot be mitigated through counterbalancing (but see Maxwell, Delaney, & Kelley, 2018) or wash-out periods (Fisher & Wallenstein, 1981), then researchers who believe there will be carry-over effects should not use within-subject designs. Unique to mediation analysis for a within-subject design, the researcher must feel confident that there are no carry-over effects for both the mediator and the outcome, rather than just focusing on a single outcome variable. Similarly, if the intended impact of the intervention is meant to be long lasting (e.g., educational interventions), within-subject designs would not be appropriate or feasible.

Participant Awareness

It may be important for participants to be relatively unaware of the hypotheses and goals of the study. Otherwise, demand characteristics, where participants purposefully attempt to align their behavior or responses with hypotheses, threaten the validity of the study (Allen, 2017). Similar effects can occur through sensitization, where, by manipulating a variable, the participants become hyperaware of it, and thus they are more reactive than if they had only experienced one level of the treatment (Greenwald, 1976). Participant awareness can be problematic for within- and between-subject designs, but may be particularly difficult for within-subject designs as participants experience both levels of the treatment. Studies where experiencing both conditions would make participants aware of the manipulation should implement between-subject designs rather than within-subject designs.

Ecological Validity

In some cases, between- or within-subject designs may reduce the ecological validity of a finding. For example, if an effect is only present when participants are exposed to both stimuli, but in a real world setting it would be unrealistic for individuals to encounter both stimuli, then

Within- or Between-Subject Design for Mediation

the findings of the study have limited implications for real-world settings. Alternatively, there may be cases where being exposed to multiple stimuli is the norm (e.g., selecting among options for meals at a restaurant), and thus a between-subject design, where a participant only considers one option, would lack ecological validity. Within-subject designs are only recommended if they are equally or more ecologically valid than a between-subject design for a given case.

Level of Effect of Interest

There are multiple effects involved in an indirect effect, and researchers should consider whether each effect is measured at the desired level. In particular, there are two effects of interest when doing simple mediation analysis: the relationship between the treatment and the mediator (*a*) and the relationship between the mediator and the outcome controlling for treatment (*b*).

Indirect effects are meant to describe a causal mechanism; however, even when an experimental design is used, the relationship between the mediator and the outcome is still correlational. This is particularly important in considering the correspondence between how these relationships would be captured in within- vs. between-subject designs. Though it is reasonable to assume that the *a*-path is the same for between- and within-subject designs (assuming no carry-over effects), the same may not be true for the *b*-path. Researchers should reflect on which design they feel would best capture the type of mediator–outcome relationship they are interested in. For example, consider a researcher in developmental psychology who is interested in the effect of years of education on critical thinking, focusing on a within-person effect of how critical thinking changes as years of education changes over the lifespan. Compare this to a second researcher, who is in sociology and is focused on the effect of years of education on critical thinking from an individual-differences perspective. The first researcher should prefer a within-subject design (comparing individuals to themselves), whereas the second should prefer a

Within- or Between-Subject Design for Mediation

between-subject design (comparing across individuals). Before selecting a design, consider all the relationships involved in an indirect effect and select a design that measures these relationships at the level of interest.

Changes in Measurement Properties

An important difference between within-subject and between-subject designs is exposure to the same measurement process multiple times. Participants who are asked to complete the same measures multiple times can show practice effects or anchoring effects, where participants either move more or move less on the measure because they have seen it previously (Greenwald, 1976; Senn, 1992). Poulton (1973) showed that scales become less effective at tapping into the latent trait upon repeated administrations. Additional effort may be needed in a within-subject design to mask the purpose of measurement, which is related to the issue of participant awareness. If the purpose of measurement is known by the participant, it could bias estimates or add error in an unintended way.

Validity Summary

When considering designs for a research study, validity should guide researchers in their decision to use a between- or within-subject design. Many of these validity threats are more salient for within- compared to between-subject designs. In addition to validity, when researchers undertake mediation analysis, they must also consider the assumptions required for causal inference. This next section walks through these assumptions for both types of designs discussed in this paper and compares them to assist in the decision of which design to use.

Causality for Indirect Effects

Mediation is an inherently causal model, where the causal antecedent (X) causes the mediator (M), which then causes the outcome variable (Y). An indirect effect is meant to quantify

Within- or Between-Subject Design for Mediation

the causal effect that X has on Y through M . While random assignment to condition is sufficient for causal inference in examining the effect of X on Y , random assignment on X is not sufficient to know that all of the effects in a mediation analysis are themselves causal. Importantly, because the mediator is not randomly assigned, it is unknown whether the estimate of its effect on the outcome is an estimate of a causal effect. However, when certain assumptions hold, it can be shown that this effect is causal. These assumptions differ for between- and within-subject designs, and the feasibility of these assumptions should weigh on the decision of which design to use in a research study. I use the potential outcomes framework (Rubin, 2005) to discuss the assumptions required in both the between-subject and within-subject designs to “identify” indirect effects (i.e., to know that these are unbiased estimates of the causal effects). Though this discussion is relatively brief, this is a very active area of research and other resources cover these topics more in-depth (e.g., Bind, VanderWeele, Coull, & Schwartz, 2016; Imai, Keele, & Tingley, 2010; Josephy et al., 2015; VanderWeele, 2010, 2015).

Two Causal Indirect Effects

In the potential outcomes framework there are two indirect effects estimated: *pure natural* and *total natural* indirect effects. These are the effects of most direct interest in mediation analysis, and this section describes the assumptions required to identify these effects. Before exploring these assumptions, it is worth explaining why there are two effects in the potential outcomes approach, while in the regression-based approaches there is typically only one indirect effect. The potential-outcomes framework allows for an interaction between X and M in the model for Y , which defines the differences between Equations 2a and 2b as well as Equations 9a and 9b (MacKinnon et al., 2020). Because the effect of M on Y can differ across conditions, these two indirect effects are not necessarily equal across conditions. One effect will

Within- or Between-Subject Design for Mediation

involve the effect of M on Y in the control condition (pure natural indirect effect), and one will involve the effect of M on Y in the treatment condition (total natural indirect effect). If the interaction between M and Y is zero, then the two indirect effects are equal. This interaction is not included in most regression approaches for mediation analysis in between-subject designs (Baron & Kenny, 1986; Preacher & Hayes, 2004; MacKinnon et al., 2020), but it is included in most common uses of within-subject mediation (Judd, Kenny, & McClelland, 2001; Montoya & Hayes, 2017). Except, in within-subject designs, the average between the two effects (pure natural and total natural indirect effects) is used as a single estimate of the indirect effect rather than differentiating the two (i.e., using b^* to estimate the average effect of M on Y). This difference is explored more in the simulation study. Next, I turn to the assumptions required to identify indirect effects in each kind of design. While the mathematics underlying the potential-outcomes approach are omitted for simplicity, I discuss these assumptions with respect to the Applied Example discussed later in the paper: Whillans, Dunn, Smeets, Bekkers, and Norton (2017) examined whether a time-saving purchase as compared to a material purchase influences positive affect through the mechanism of reduced subjective time pressure.

Between-Subject Designs

Two potential outcomes to be considered are 1) what would a participant's positive affect have been if they were in the time-saving purchase condition (treatment), and 2) what would it have been if they were in the material purchase condition (control). The difference between these two potential outcomes would be the causal effect of purchase type on positive affect for that individual, and averaging over all individuals would give the average treatment effect on the outcome. In the between-subject case, it is not possible to observe both outcomes, and so information that is observed is used to estimate information that is not observed. Information

Within- or Between-Subject Design for Mediation

about those in the time-saving purchase condition is used to estimate information about the potential outcome for those in the material purchase condition (and vice versa). This process can be used to estimate the effect of X on both M and Y .

The potential outcomes framework is very useful for understanding causal effects; however, just calculating these effects does not necessarily make them causal. One of the great advantages of the potential outcomes framework is that it articulates what assumptions are required for this process to result in causal effects (Pearl, 2001; Robins & Greenland, 1992).

First, temporal precedence: the cause must precede the effect. For indirect effects this means X must precede M and Y , and M must precede Y . When X is randomly assigned, the first point is supported, but it is unknown whether the mediator precedes the outcome. In the above example, it would fall on the researchers to support the claim that time-pressure occurs before positive affect. This task can be incredibly difficult because psychological phenomena are often dynamic and ever-present (e.g., each person has a level of positive affect that fluctuates over time), and it is very difficult to know the exact order of events. Measuring constructs in a specific order can assist with this issue, but measurement order itself cannot be used to claim causal order.

A second assumption is the *no-omitted confounder assumption*: all confounders of the X - Y , X - M , and M - Y relationships are measured and included correctly in the model. For example, if sleep quality causally affects both time pressure and positive affect, then time pressure and positive affect will be correlated due to this shared cause (confounder). If sleep quality is not included in the model, then the correlation between time pressure and positive affect that is attributable to sleep quality will instead be attributed to the causal effect of time pressure on positive affect. Thus, the causal impact of one variable on another could be overestimated (or in

Within- or Between-Subject Design for Mediation

some cases underestimated) if confounders are not included in the model. When X is randomly assigned, as it is in the between-subject design, both the X - Y and X - M relationships are unconfounded. However, the M - Y relationship remains unaddressed.

The third assumption is that there are no measured or unmeasured confounders of the M - Y relationship that are affected by X . For example, if the type of purchase influenced feelings of wealth, which then caused both time pressure and positive affect, this would invalidate the causal claim that time pressure mediates the relationship between type of purchase and positive affect, because it could be that feelings of wealth is the true mediator, and time pressure is an outcome of feelings of wealth. Random assignment on X does not assist in confirming or disconfirming this assumption. Violations of this assumption can potentially be addressed by considering multiple mediator models (e.g., where feeling of wealth is measured and included as an additional mediator; Daniel et al., 2015).

Ultimately, a researcher interested in applying mediation analysis to estimate causal effects must consider these assumptions in advance of collecting data to provide as much support as possible for these assumptions through design characteristics, measurement of potential confounders, and measuring the variables in their hypothesized order. MacKinnon and Pirlott (2015) provide additional ways to improve the validity of causal claims in mediation analysis in between-subject designs. Researchers should also consider sensitivity analyses to evaluate the extent to which violations of these assumptions could affect their results (Imai, Keele, Yamamoto, 2010; VanderWeele, Chiba, 2014; Tingley, Yamamoto, Hirose, Keele, & Imai, 2014).

Within-Subject Designs

Within- or Between-Subject Design for Mediation

The potential outcomes framework and assumptions for causal inference in mediation analysis have been widely discussed for between-subject designs. However, within-subject designs are discussed less, even though they have the unique benefit that each participant is observed on the mediator and the outcome in both conditions. This means that each person's treatment effect on the mediator and outcome is identified (Joseph et al., 2015). Except that each participant receives these treatments in a specific order, and the potential outcomes must be considered for the different orders of administration, or we must assume that order does not affect the treatment effects.

Just as in a between-subject design, the within-subject design requires temporal precedence. Similarly, the no omitted confounder assumptions for the X - Y and X - M relationships are unneeded. But for the M - Y relationship, the no omitted confounders assumption is still needed. A unique aspect of the within-subject design is that the researcher can consider both upper- and lower-level confounders separately (i.e., person-level confounders and time-level confounders). For example, sleep quality might be a lower-level confounder if it differs across the two conditions within a person. Alternatively, an upper-level confounder would be something constant for the individual across the two repeated-measurements, for example occupation. Joseph et al. (2015) show that if unmeasured upper-level confounders of the M - Y relationship have an additive effect on both M and Y and there is no unmeasured heterogeneity in the effect of X on M , M on Y , and X on Y , then upper-level confounders of the M - Y relationship can be ignored for within-subject designs. This is an advantage of the within-subject design, because when these assumptions are met unmeasured person-level confounders (particularly of the M - Y relationship) do not bias the estimates of the indirect effects. Whereas, if a researcher is using a between-subject design they must assume that there are no unmeasured person-level

Within- or Between-Subject Design for Mediation

confounders. Similar to between-subject designs, within-subject mediation analysis requires the assumption that there are no measured or unmeasured confounders of the M - Y relationship that are affected by X (Joseph, et al., 2015).

The assumption that there are no carry-over effects is also needed for within-subject designs (Joseph, et al., 2015). This means that previous levels of the measured variables cannot affect future measures of these variables. Joseph and colleagues (2015) describe a method of sensitivity analysis for evaluating the assumptions in a within-subject mediation analysis.

Comparing Assumptions

Many assumptions are shared for between- and within-subject designs for identifying indirect effects. Temporal precedence of X - Y and X - M are established in both between- and within-subject designs. Similarly, in both designs the X - M and X - Y relationships are unconfounded. This results in a reduced list of assumptions required by both designs:

- 1) Temporal precedence (M comes before Y).
- 2) No omitted confounders of the M - Y relationship
- 3) No measured or unmeasured confounders of the M - Y relationship that are affected by X .

However, Assumption 2 can be weakened to “no omitted lower-level confounders of the relationship between the mediator and the outcome” in the within-subject case under certain conditions described above. This potentially suggests an advantage of the within-subject design for making causal inference because the assumptions are weaker, except that an additional assumption is required for within-subject designs: no carry-over effects. Because of this additional assumption, there is no clear dominant approach to causal inference in mediation. If the researchers feel confident they can meet the no carry-over assumption, within-subject

Within- or Between-Subject Design for Mediation

mediation models may be more likely to meet these assumptions, because upper-level confounders of the M - Y relationship are ignorable under certain conditions.

Statistical Power

Researchers should begin their evaluation of which design to use by considering threats to validity and causality. If each design is considered plausible for both validity and causality, they should also consider statistical power. The ability for within-subject designs to detect effects with greater precision than between-subject designs is well-documented for tests comparing means (Maxwell, Delaney, & Kelley, 2018; Senn, 1993; Venter & Maxwell, 1999). Previous work in multivariate analysis (of which both mediation analysis and within-subject analysis are special cases) would suggest that using within-subject factors in a mediation analysis should increase power (e.g., Muller & Peterson, 1984; Muller, LaVange, Ramey & Ramey, 1992). While theory suggests there ought to be a power advantage for the within-subject design, the extent of that advantage is unclear and may vary across different research scenarios. In this section, I use a Monte Carlo simulation to explore the relative power of each design for mediation analysis under a variety of conditions. This simulation answers two important sets of questions for researchers choosing a design for their mediation analysis:

- 1) Do within-subject designs have higher empirical power for detecting indirect effects than between-subject designs in all cases? How large is this power advantage? Are there specific cases when they do not have an advantage?
- 2) What are the population parameters that affect empirical power in a within-subject mediation analysis, and how do they do so?

One previous study has examined power for within-subject mediation analyses (Yzerbyt et al., 2018); however, this study did not directly compare power across designs. Additionally,

Within- or Between-Subject Design for Mediation

this previous research did not fully explore the parameter space relevant to researchers considering using within-subject designs. The present study expands on Yzerbyt, et al. (2018) by exploring additional inferential methods and a broader range of sample sizes, varying correlation among repeated-measurements and the moderation parameter, and comparing models with and without the moderation parameter. Methods of inference are varied for both between- and within-subject designs; however, for brevity in the main manuscript, only the results for percentile bootstrap confidence intervals are presented. The percentile bootstrap was selected because previous research suggests it has good balance of type I error and power and performs well under non-normality and with incomplete data (Biesanz, Falk, and Savalei, 2010; Hayes and Scharkow, 2013). Results for other inferential methods are included in the Supplementary Material. Correlation among repeated measurements was varied because power for detecting mean differences in within-subject designs depends on this correlation, a factor that was ignored in Yzerbyt et al. (2018). No previous research in either the between-subject or within-subject literature has examined the impact of the d parameter on the power to detect the indirect effect, which defines the differences between Equations 2a vs. 2b and 9a vs. 9b. It is possible that by estimating this parameter when it is not needed, researchers could lose statistical power, but they also must consider the impact of omitting this parameter if it is non-zero. The moderation parameter, d , was always set to 0 in Yzerbyt et al. (2018); in this simulation, d is systematically varied and its impact on power is investigated. Finally, and most importantly, by simulating data in a way that is equivalent for between- and within-subject designs, this simulation allows for direct comparison of empirical power in each of the two designs and allows researchers to decide which design to use for their own study, after first considering validity and causality.

Simulation Methods

Within- or Between-Subject Design for Mediation

For this simulation study I manipulated seven parameters, each of which might influence the rejection rates of the tests of inference on the indirect effect. Sample size was set at 20, 50, 100, and 200. The coefficients a -path, b -path, c' -path, and d -path were set at 0, 0.14, 0.39, and 0.59 (completely crossed). These values were selected based on previous simulations for between-subject mediation (Biesanz, Falk, & Savalei, 2010; Hayes & Scharkow, 2013). The values for these paths correspond to zero, small (2%), medium (13%), and large (26%) effects with respect to explained variance as defined by Cohen (1988). Correlation among mediators (ρ_M) and correlation among outcome variables (ρ_Y) were set at 0, .3, .6, and .9 (completely crossed). I selected these values to span the range from uncorrelated repeated measures to very highly correlated repeated measures².

There were 16,384 conditions; however, 4,736 conditions resulted in non-positive definite covariance matrices, so only 11,648 conditions were successfully run. The conditions that did not run aligned systematically with certain values (e.g. $\rho_M = .9$; See Appendix for more information). Two thousand samples were generated in each condition (1,000 for within-subject design and 1,000 for between subject designs). The Monte Carlo error will depend on population power, which can vary from near zero to near one, but is maximized at a population power of 0.50, at which point the Monte Carlo error is 0.016.

For each combination of possible conditions, I calculated an implied population covariance matrix and mean vector and generated data to follow these covariance and mean patterns (See Appendix for details). I fit two models in each sample: the mediation model excluding the moderation parameter (d , Equations 2a and 9a for between- and within-subject

² A small literature review (26 papers, 109 correlations) of published examples of within-subject mediation published in 2020 examined the absolute value of correlations of repeated measurements. We found that 20% were between 0 and .3, 32% between .3 and .6, 35% between .6 and .9, and 13% between .9 and 1. Most of the correlations were positive, and 15% were negative. More complete results are in the Supplementary Materials.

Within- or Between-Subject Design for Mediation

designs respectively), and the model including the moderation parameter (Equations 2b and 9b for between- and within-subject designs respectively). In each sample for each of the two models, a 95% percentile bootstrap confidence interval was calculated using 1,000 resamples, and I recorded whether the confidence interval excluded zero (rejecting the null hypothesis), providing an overall rejection rate for each model in each condition.

All conditions where the indirect effect is zero result in estimated type I error rates, which are discussed in the Supplementary Materials. By examining all the conditions where neither a nor b is zero, I estimated the statistical power of each test. I calculated empirical power by computing the proportion of samples that were deemed significant in each condition where the population indirect effect does not equal zero. I aggregated the results across manipulated conditions but also broke apart useful and informative subsets of the data to investigate cases in which manipulated factors may interact to influence empirical power.

Simulation Results

Complete simulation data and analysis code are available at osf.io/3njw6/.

Between vs. Within-Subject Designs. Comparing across all conditions with non-zero indirect effects in the simulation, within-subject inferential methods have clear dominance over the between-subject inferential methods with respect to empirical power. In 86% (no moderation estimated) and 89% (moderation estimated) of all conditions, the within-subject method has higher empirical power than the between-subject counterpart. Three questions emerge from this result: (1) Which cases have higher empirical power for the between-subject design, (2) How much more empirical power is gained by using a within-subject design across different sample sizes and effect sizes, and (3) What unique characteristics of within-subject designs impact their empirical power?

Within- or Between-Subject Design for Mediation

To understand this first question, I examined the 652 (out of 6,192) conditions for which both within-subject models (estimation and no-estimation of moderation) had lower empirical power than their between-subject counterparts. See Table S3 for comparisons involving all inferential methods. These cases came from all sample sizes, *a*-paths, *b*-paths, and direct effects; however, there were distinct patterns with the correlation and moderation parameters. Conditions with high correlations among the mediators, low correlations among the outcomes, and smaller moderation parameters more commonly showed higher empirical power for between-subject designs than within-subject designs. Empirical power is lowest when these 3 conditions align, so much so that there is no benefit of within-subject designs with regard to power. These patterns are explored more in-depth in the next section on Within-Subject Test Performance.

[Figure 3 Here]

Figure 3 shows empirical power for percentile bootstrap confidence intervals in both between- and within-subject designs across sample size, indirect effects, and estimating moderation (See Table S2 and Figure S1 for all inferential methods). For simplicity this I focus on the model with moderation, though all patterns are consistent across moderation and non-moderation models. Across sample sizes and indirect effects, within-subject designs have higher empirical power on average than between-subject designs for mediation. In general, there seems to be a larger benefit of within-subject designs for larger indirect effects and larger sample sizes, except that when empirical power gets very close to 1, the designs get closer in empirical power. The benefit of within-subject designs can be quite large, up to about a .3 higher than between-subject designs. Notably, in Figure 3 it is common for the between-subject line at a specific *N* to run parallel or near parallel to the within-subject line at *N*/2. For example, tests seem to have similar empirical power with a between-subject design at *N* = 200 and within-subject design at *N*

Within- or Between-Subject Design for Mediation

= 100, or between-subject design at $N = 100$ and within-subject design at $N = 50$. In general, there is a clear advantage of the within-subject design in average empirical power, and this benefit is approximately equal to a doubling of sample size within the range of sample sizes and effect sizes explored. Though, as previously noted, in specific cases, this benefit can be quite small, nonexistent, or reversed, so researchers should evaluate the appropriateness of within-subject versus between-subject designs for their specific case.

The next section dives deeper into empirical power specifically for within-subject designs, exploring factors that are unique to these designs and how those factors affect empirical power. These results help to elucidate the situations in which within-subject designs have lower empirical power than between-subject designs and help to emphasize the importance of considering power for one's specific research problem, rather than relying on one design over another in all cases.

Within-Subject Test Performance. I describe the results with respect to the empirical power separately for the different factors examined in the study. Results for type I error, other inferential methods, and an in-depth comparison of these results with Yzerbyt et al. (2018) are reported in the Supplementary Materials.

I calculated empirical power by averaging the rejection rates across all conditions where the population indirect effect did not equal zero, grouped by sample size (592 conditions per cell). Unsurprisingly, empirical power increases as sample size increases and as the indirect effect increases (See Figure 3).

Correlation among repeated measurements. Figure 4 (top panel) provides a visualization of the pattern of empirical power across correlations among the mediators for each sample size.

Within- or Between-Subject Design for Mediation

Empirical power increases up to $\rho_M = .6$, then the empirical power is lower for $\rho_M = .9$ ³. There seems to be a point at which increased correlation among the mediators is detrimental to power. Supplemental simulations suggest the peak of empirical power is around $\rho_M = .75$.

Overall, as correlation among outcome variables increases, so too does empirical power (See Figure 4, bottom panel). The increase is nonlinear but monotonic; though, it would be reasonable to expect this trend to asymptote at 1 as the correlation approaches 1.

[Figure 4 Here]

Moderation Effect and Whether Moderation is Estimated. To examine the role of the moderation parameter, I averaged empirical power across all conditions with the same level of moderation parameter, sample size, and b -path. There were no systematic differences with the size of the a -path. In general, empirical power does not merely increase or decrease with the moderation parameter (Figure 5, see Figure S2 for all inferential methods). Instead, as the moderation parameter increases, empirical power decreases if the moderation is not estimated. But power stays relatively stable when moderation is estimated and even increases in some cases. These differences between estimating and not estimating the moderation parameter seem to decrease as the b -path gets larger; however, it is difficult to evaluate this completely due to missing data from sets of impossible parameters (when both b and d are 0.59, the population correlation matrix is not positive definite). Overall, there tends to be a power advantage when estimating the moderation parameter.

If the moderation parameter is truly zero, should it be estimated? Summarizing across all cases when the moderation parameter was zero and comparing empirical power in each

³ This pattern was true of all inferential methods except the two normal theory tests at $N = 20$ which show monotonic increases in empirical power. However, these methods exhibit the same non-monotonic pattern as the other methods at higher sample sizes. More information provided in the Supplementary Materials.

Within- or Between-Subject Design for Mediation

condition, 69% of cases saw a power advantage when omitting the moderation coefficient when it was zero at the population level. Overall, the average empirical power gain is very small, 0.004. Therefore, though one may have a power gain by omitting the moderation parameter, this gain is relatively small and relies on the assumption that the moderation parameter is zero. Based on these results I recommend that researchers estimate the moderation parameter, unless there is strong a priori information to suggest it is zero.

Size of Indirect Effect. Though the empirical power in some of the previous tables and figures may seem quite low, it is important to note that many of these estimates are averaged across the size of the indirect effect (ranging from 0.0196 to .3481). Table 1 shows that across indirect effects and sample sizes the empirical power of most tests approaches 1.0 at large samples and/or large indirect effects. Because the data generated were standard normal, the indirect effects generated can be interpreted as standardized. These power estimates are very idealistic considering the data is perfectly normal and generated with equal variances for the outcome variables and mediators. With these precautions in mind, it seems that between 100 and 200 participants in a within-subject design would be sufficient to achieve the traditionally recommended 0.8 power for an effect of 0.15 using a percentile bootstrap confidence interval. As can be seen in Table 1, much larger sample sizes are needed to detect smaller effects.

[Table 1 Here]

Based on these simulation results, it seems that the correlation among repeated-measurements and the moderation parameter can play an important role in the power to detect an indirect effect in a within-subject design.

Empirical Example

Within- or Between-Subject Design for Mediation

As described previously, Whillans and colleagues (2017) used a within-subject mediation analysis to examine whether a time-saving purchase as compared to a material purchase influences positive affect through the mechanism of reduced subjective time pressure. A study with a within-subject design was selected to extract the required statistics for conducting power for a within-subject mediation analysis, otherwise the researcher may need to make some guesses about certain parameters (e.g., ρ_M and ρ_Y) to conduct the power analysis. In this section, the notation using a subscript D indicates standard deviation of the difference score (σ_D) and an effect standardized by the standard deviation of the difference score (d_D), whereas the subscript P indicates pooled standard deviation (σ_P) and an effect standardized by the pooled standard deviation (d_P). Typically, pooled standard deviations are used to calculate t-statistics and effect sizes for between-subject designs but standard deviations of the difference score are used for within-subject designs.

Unstandardized effect sizes for a -, b - and c -paths were extracted from Whillans et al. (2017) and converted to effect sizes standardized by the pooled standard deviation (d_P) using the reported individual standard deviations of the mediators and outcomes in each condition. This ensures the effect sizes will be comparable between the two potential designs. This gives new standardized effect sizes, using the pooled standard deviation, for both the a -path ($d_P = \frac{-0.74}{1.68} = -0.44$) and for the c -path ($d_P = \frac{0.29}{0.73} = 0.39$). The pooled standard deviations were used to standardize the b -path as $-0.15 \times 1.68/0.73 = -0.34$. Information about the XM interaction could be included in the power analysis; however, because no information was reported in Whillans et al. (2017) about this parameter estimate, it was assumed to be zero.

Given this information, power analysis for the between-subject design can be conducted, but the correlations among repeated measures are needed for calculating the power for a within-

Within- or Between-Subject Design for Mediation

subject design. While reporting the correlation among repeated measures has been recommended to assist with meta-analytic research (Looney, Feltz, & VanVleet, 1994), they were not reported in the article. However, using the same methods described in the Supplementary Materials the correlations among the mediators is $\rho_M = 1 - \frac{\left(\frac{2.09}{1.68}\right)^2}{2} = 0.22$ and the correlation among the outcomes is $\rho_Y = 0.32$.

Estimating Sample Size or Power

An attempt to replicate Whillans et al., (2017) would require the research team to decide whether to use a within-subject or between-subject design, and they might want to consider either of the two following questions:

- (1) Given a specific sample size (e.g., $N = 60$) what power do I have to detect the indirect effect using either a between-subject design or a within-subject design?
- (2) What sample size would I need to achieve a pre-specified level of power (e.g., 90%) using either a between-subject design or a within-subject design?

To address this first question, there are numerous tools that can calculate power for between-subject designs (e.g., MCpowrMed by Shoemann, Boulton, & Short, 2017; WebPower by Zhang & Yuan, 2018; MedPower by Kenny, 2017; pwr2ppl by Aberson, 2019). MedPower was used because it accepts standardized coefficients. Kenny (2017) uses the joint significance test as it has similar power to the other recommended methods and is less computationally intensive. Given a sample size of 60, a between-subject design, and an α -level of 0.05, the power is 0.32. I used the R script in the Supplemental Materials to calculate the power for a within-subject design using the joint significance test, but there are also options for bootstrapping and Monte Carlo confidence intervals. Power is estimated to be 0.64. The estimated powers for the two designs

Within- or Between-Subject Design for Mediation

very closely align with the observation from the simulation study that power for within-subject designs tends to be about twice the power for between-subject designs at the same sample size.

To address the second question, I used MedPower (Kenny, 2017) to compute the required sample size to detect the desired indirect effect with a power of 0.90. To achieve a power of 0.90 in a between-subject design, the research would need a sample size of 211 (See Supplemental Materials for details). I ran the script in Supplemental Materials iteratively to find a sample size for which a within-subject design would achieve a power of 0.90. The result was a sample size of 99. Overall, this section demonstrated how to conduct power analysis for within- and between-subject mediation analysis, and illustrates how the Monte Carlo simulation results are mimicked in an applied example.

Discussion

This paper discusses three primary factors that should be used for selecting a research design for mediation analysis: validity, causality, and statistical power. Within-subject designs tend to introduce more threats to validity than between-subject designs, but these designs can, in some cases, be more ecologically valid. With respect to causality, between-subject designs require stronger assumptions about omitted confounders, whereas within-subject designs require the additional assumption of no carry-over effects. Overall, statistical power is greater for within-subject designs, with notable exceptions. Together this information suggests that there is no singularly dominant design for testing mediation, and that researchers selecting a design must weigh validity, causality, and statistical power for their specific case.

Recommendation for Applied Researchers

Researchers choosing a design to conduct mediation analysis should first weigh the possible designs on validity and causality. Carry-over effects are of primary concern for both

Within- or Between-Subject Design for Mediation

validity and causality, and researchers who do not feel they can implement a research study without the threat of carry-over effects should not use within-subject designs for estimating indirect effects. If carry-over effects are not of concern, then within-subject designs require fewer causal assumptions and tend to have higher power. However, if carry over is of direct concern then a between-subject design may be most appropriate.

After considering validity and causality, researchers can consider the balance of statistical power and cost. While the simulation study suggests that within-subject designs can require approximately half the sample size of between-subject designs to achieve the same level of power, this is an overly general rule of thumb and can only be used in limited circumstances. Available tools, such as the R script in the Supplementary Materials as well as a multitude of available tools for between-subject mediation analysis, make it easy to compute estimated sample size or power for both within-subject and between-subject designs given reasonable guesses about the parameter sizes. However, power analyses are only as good as the estimates they are based upon, and so trying a variety of effect sizes to see how the power estimates vary is an important part of conducting power analysis for mediation. Researchers should not rely on rules of thumb for choosing between designs or selecting sample sizes, but rather use power analyses informed by the expectations of the research team.

Based on the results of the simulation study, there are specific cases when researchers should be particularly wary about using a within-subject design, as there may be little or no power advantage in these cases. When effect sizes and sample size are small or when effect size and sample size are large, power is very low and very high respectively, and so the choice of design has only a very minimal effect on power. Additionally, there were cases when the within-subject design had lower power than between-subject designs: high correlations among the

Within- or Between-Subject Design for Mediation

mediators, low correlations among the outcomes, and small moderation parameters. If researchers believe a priori that this combination of parameters is plausible, they should prefer using a between-subject design, particularly if the within-subject design may cost more.

Why does power drop at high correlation among mediators?

Power drops when the mediators have a very high correlation with each other (See Figure 4). Typically, higher correlation among outcomes often results in a monotonic increase in power, but because mediators serve both as outcomes and as predictors, there is a trade-off with the increase in correlation. As the correlation among repeated-measures among the mediators increases, the variance in the error term in Equation 8 will decrease, improving the precision of the estimation of the a -path, as there is a reduction in the error in estimation. This reduction in error in estimation means greater power to detect the a -path. However, the variance of the mediator difference score will also decrease as the correlation among the mediators increases. The difference between the mediators is used as a predictor of the difference between the outcomes. All else being equal, as the variance of a predictor variable decreases, the standard error of the regression coefficient will increase. In this case, increasing the correlation among the mediators decreases the variance of the difference between the mediators, and the standard error for the b -path increases. So, as the correlation among the mediators increases, the precision of the estimate of the b -path decreases, and thus the power to detect the b -path decreases. Altogether, increasing the correlation between the mediators increases the power to detect the a -path but decreases power to detect the b -path. Ultimately this results in a non-monotonic function with respect to power of the indirect effect, ab . In general, higher correlation among the mediators is beneficial, but only to a point. Researchers should be aware that if the correlation among the repeated measurements of the mediators is too high, statistical power may be low.

Within- or Between-Subject Design for Mediation

Should I Estimate the Moderation Parameter (d)?

When estimating indirect effects with either between- or within-subject designs, researchers should include the d -path to improve their statistical power, unless they have strong reason to believe this coefficient is zero in the population. The simulation study suggests that there is perhaps a very small cost in power if this parameter is zero, but the gains in power if it is non-zero are much larger. Estimating this coefficient is currently common practice in within-subject mediation, but not in between-subject mediation. Previous work in between-subject mediation has emphasized the importance of estimating this parameter (Kraemer, Wilson, Fairburn, Agras, 2002; MacKinnon et al., 2020), but it has not been adopted by many applied researchers. Additionally, this parameter is required for estimation of causal indirect estimates such as the *pure natural* and *natural indirect effects*.

Can We Use Smaller Sample Sizes?

Many applied researchers may be excited to learn that they can use smaller sample sizes to achieve the same amount of power for mediation analysis by using a within-subject design. However, between-subject mediation studies are historically underpowered: 75% of mediation studies have less than 80% power to detect an indirect effect involving one small effect (a or b ; Fritz & MacKinnon, 2007). Average effects sizes are not well known for mediation, but what information is available suggests that they are often small (Götz et al., 2021). The median sample size for between-subject designs observed in Fritz & MacKinnon (2007) was 187. Based on the results of the current study, a similar sample size (200) would have 80% power to detect indirect effects of approximately 0.13 or higher using a within-subject design. This is just slightly lower than the average standardized indirect effect from Götz and colleagues (2021). Ultimately, this means that the typical researcher who currently uses between-subject mediation analysis and is

Within- or Between-Subject Design for Mediation

considering switching to within-subject designs should consider maintaining similar sample sizes to what they've previously used, if a power of 80% is what is desired. Otherwise, if smaller sample sizes are used, power will continue to be low.

Do I Need to Worry about Type I Errors?

While the focus of this study was on power, type I error is also an important characteristic of a statistical method to consider. The Supplementary Material shows that tests of the indirect effect were generally conservative in the simulation results, but this is not a reason to ignore concerns about type I errors in the mediation literature. The tests are not as conservative when one of the paths involved in the indirect effect is non-zero and the other is zero. This seems to be a likely case for when mediation analysis might be implemented. In practice, mediators are often introduced into models when there is strong previous evidence of either the association between the manipulation and the mediator, or the mediator and the outcome. Additionally, the data generated in this simulation come from perfect normal distributions, do not contain notable outliers, and were not subject to any researcher degrees of freedom— all factors that can increase type I error rates in real data application settings (Finch, West, & MacKinnon, 1997; Zu & Yuan, 2010; Simmons, Nelson, & Simonsohn, 2011). Recent meta-analytic studies of mediation analysis show there is an over-abundance of confidence intervals that are closer to zero than would be expected from studies of high power or studies generated by a null distribution (Götz, et al. 2021; Charlton et al., 2021). This suggests that multiple testing, publication bias, or other questionable research practices may be overpopulating published mediation results with false-positives. Researchers conducting mediation analyses should be wary of these practices, preregister their planned mediation analyses, and practice transparency in their reporting of the procedures they undertook in conducting mediation analyses.

Within- or Between-Subject Design for Mediation

Methodological Contribution

This research expands on the methodological research on power in mediation analysis in several ways. Perhaps most importantly, it provides the first comparison of empirical power in mediation across designs. While theory can tell us that we should expect some power advantage of within-subject designs in mediation analysis, the size of this advantage and the contingencies of it are not immediately apparent without a simulation study. This study demonstrates which population characteristics can reduce statistical power in within-subject mediation: high correlations among mediators and low correlations among outcomes. These findings should be explored for more complex repeated-measures designs (e.g., longitudinal and multilevel designs).

The current simulation study also fills a much-needed gap for the development of within-subject mediation analysis. Montoya & Hayes (2017) described multiple inferential methods that could be used for within-subject mediation analysis but did not provide recommendations for which method performs best. Based on the simulation results presented in the Supplementary Material, the recommended inferential approach is a percentile bootstrap confidence interval or, if the original data is not available, the Monte Carlo confidence interval. The computational tool MEMORE for SPSS and SAS can be used to estimate these within-subject mediation models and provide inferential tests including percentile bootstraps and Monte Carlo confidence intervals (Montoya & Hayes, 2017), and the `JSmediation` R package can be used to conduct the joint significance test for within-subject mediation models (Yzerbyt et al., 2018).

Limitations & Future Directions

Though the current research provides new and useful information to researchers interested in using mediation analysis, there are a variety of limitations. Many of these

Within- or Between-Subject Design for Mediation

limitations suggest avenues for future research in the ever-growing literature on mediation analysis, particularly focusing on within-subject and repeated-measures designs.

An important limitation of this simulation study is one that should be applied to most other simulation studies in this area: there is very little evidence that the effect sizes examined (or the combinations of these effect sizes) are realistic for the field of behavioral science. There is a pervasive norm of selecting “small”, “medium”, and “large” effects according to Cohen (1988); however, it is not clear if these effect sizes are common in real data. Very little meta-research on indirect effects is available; however, what is available suggests that indirect effects in psychology tend to be very small. A recent meta-analysis on the indirect effect of contact on prejudice through threat found an average standardized indirect effect of $-.072$ (Aberson, 2019). The empirical example has an observed standardized indirect effect of approximately 0.07 (Whillans, et al., 2017). Götz and colleagues (2021) recently conducted a large-scale study examining indirect effects in psychology by looking at more than 2,500 tests of mediation published in 5 APA journals in 2018 and 2019. Very few of the reported effects were standardized (11%), but of the standardized effects there was an average indirect effect of 0.145 and a median of 0.08 . Notably, Götz and colleagues (2021) find strong evidence for questionable research practices in mediation analysis. This means that the estimates from this paper are likely overestimates of the population parameters. Overall this suggests that of the effect sizes investigated in this and other simulation studies, the smallest effects investigated may be the closest to realistic, and sometimes even smaller effect sizes would be more realistic. However, the range of observed effects is very large, and the simulation study can speak to this broad range. Ultimately, as researchers move forward into sample size planning, it would be more beneficial to use computational tools such as those in the Supplemental Materials, which allow

Within- or Between-Subject Design for Mediation

them to select their estimated effect sizes and conduct the power analyses using exact values specific to their research question at hand. The tables and figures in this article are not recommended for selecting sample size, but rather illustrate the patterns in empirical power as different population parameters are systematically manipulated.

The simulations in this paper assume that between- and within-subject effects are equal. In multilevel modeling, when between- and within-effects differ, these effects are called contextual effects (Enders & Tofighi, 2007; Raudenbush & Bryk, 2002; Kreft et al., 1995) or ergodicity (Molenaar, 2004). To reasonably compare between- and within-subject designs, the indirect effects are assumed to be equal; however, this does not necessarily mean that these effects will be equal in a given research context, in particular the relationship between M and Y . In choosing a design, researchers should reflect on whether they expect these effects to differ across levels (within and between) and which of these effects is of primary interest. Recent developments in multilevel mediation analysis allow for the estimation of both types of effects, with appropriate designs (Rockwood, 2018; Rockwood & Hayes, 2022; Preacher, Zyphur, Zhang, 2010).

Throughout this paper, only experimental designs have been discussed where individuals are randomly assigned to conditions or all individuals experience both conditions. Mediation analysis is used for other designs; however, as previously noted, more assumptions are needed for causal inference when using quasi-experimental or completely cross-sectional design. The statistical methods discussed in this paper do not rely on using an experimental design, but the validity of the causal claims is more tenuous when random assignment or crossover designs are not used. With respect to the claims about power, balanced conditions were always assumed, allowing there to be alignment across between- and within-subject designs. In within-subject

Within- or Between-Subject Design for Mediation

studies, pre-post designs are frequently analyzed using the described mediation models. In the case of pre-post designs where all participants go through the same intervention but are measured before and after the intervention, the researchers must assume that all change over time is due to the intervention and there is no natural change over time (Montoya & Hayes, 2017).

This paper does not consider design with both within- and between-subject factors. Valente and MacKinnon (2018) discuss models available for estimating mediation with an experimental (between-subject factor) and pre-post design at the same time. Another example of integrating between- and within-subject designs is counterbalancing order and treating this as a between-subject factor. Testing order as a moderator of the indirect effect would help to evaluate if carry-over effects or other order effects invalidated a study (Montoya, 2018). More complex designs could be handled using moderated mediation in a multilevel (Rockwood, 2018), longitudinal (Bind, Vanderweele, Coull, & Schwartz, 2016), or multilevel SEM framework (Zyphur, Zhang, Preacher, & Bird, 2018). By combining between- and within- subject designs, researchers can often benefit from additional statistical power of repeated-measurements, while still benefiting from the causal interpretation that comes from random assignment. Additional development in this area, especially a deeper understanding of the causal assumptions for indirect effects in these models, is needed to bridge the gap between methodological innovation and current practice among applied researchers.

There are a variety of conditions that went uninvestigated in the current simulation study. Each of these provides a potential avenue for expansion of understanding the performance of the proposed methods under different circumstances. Previous simulation research in between-subject mediation analysis has examined the effects of non-normality (Biesanz, Falk, & Savalei, 2010; Finch, West, MacKinnon, 1997), missing data (Biesanz, Falk, & Savalei, 2010), and

Within- or Between-Subject Design for Mediation

heteroskedasticity (Yuan & MacKinnon, 2014). Each of these issues may present particularly interesting quandaries in a within-subject design as well. An important future direction is to consider how lack of equal variances (across M s or Y s) influences inference in the within-subject case. There is some evidence from Yzertbyt et al. (2018), who generated their data in such a way that the variances were not equal, and their results do not seem to notably differ from the results in this paper. However, a much more purposeful investigation of this issue is needed. Similarly, power in within-subject designs may be impacted by whether variances are equal across conditions and whether repeated-measures are positively or negatively correlated. All these conditions could be investigated in future simulations. Because the data generated in this simulation are generated under ideal circumstances, future research should examine required sample sizes for adequate power under less than ideal circumstances, such as non-normality, heteroskedasticity, missing data, and unequal variances for the outcomes and mediators (particularly for small sample sizes). Additionally, future computational tools should allow for consideration of these factors when conducting power analysis.

Summary

This study describes three major factors that researchers should use when selecting a research design for mediation: validity, causality, and statistical power. Validity and causality must be considered before weighing statistical power, and only if both proposed designs seem to be valid and meet assumptions for causality, should a researcher proceed to weighing statistical power. I used Monte Carlo simulation to demonstrate power advantages of using a within-subject versus between-subject design when conducting mediation analysis. In most cases, within-subject designs have higher empirical power than between-subject designs. Generally, for between-subject designs with a sample size of N , within-subject designs with a sample size of

Within- or Between-Subject Design for Mediation

approximately $N/2$ have comparable empirical power. However, within-subject designs have lower empirical power when correlation among outcomes is low and correlation among mediators is high. The applied example of power analysis for within-subject and between-subject mediation analysis demonstrates how to conduct power analysis for each design. The R script provided in the Supplemental Materials provides researchers with an easy to use tool for conducting power analysis for within-subject mediation analysis. Based on the results of this study, researchers should first consider validity and causality when assessing whether a within-subject design would be appropriate for their study investigating mechanisms. Then they should consider if using a within-subject design can achieve greater statistical power, and if so they can use the methods developed by Montoya & Hayes (2017) as well as the macro MEMORE for SPSS and SAS to conduct their analyses.

Within- or Between-Subject Design for Mediation

Acknowledgements: Special thanks to Dr. Andrew Hayes, Dr. Craig Enders, Dr. Michael Edwards, Jessica Fossum, Tristan Tibbe, Valerie Polad, and the Quantitative Research Collaboratory for providing feedback on iterations of this manuscript. Thanks to the panel of expert reviewers who provided their insight and recommendations for improvement of the manuscript. Any remaining errors or problems with the manuscript are the fault of the author. Parts of this manuscript were presented at the annual meetings for the Southeastern Psychological Association (2016) and the Society for Personality and Social Psychology (2019). This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant DGE-1343012.

Declaration of Interest Statement: Dr. Montoya reports grants from National Science Foundation, during the conduct of the study.

Within- or Between-Subject Design for Mediation

References

- Aberson, C. L. (2019). Pwr2ppl: Power analysis for common designs. R package version 0.1. Retrieved from <https://cran.r-project.org/web/packages/pwr2ppl/index.html>. DOI 10.4324/9781315171500.
- Aberson, C. L. (2019) Indirect effects of threat on the contact-prejudice relationship. *Social Psychology*, 50(2), 105-126. DOI: 10.1027/1864-9335/a000364.
- Allen, M. (2017). *The sage encyclopedia of communication research methods* (Vols. 1-4). Thousand Oaks, CA: SAGE Publications, Inc DOI: 10.4135/9781483381411.
- Baron, R. M., & Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6), 1173. DOI: 10.1037/0022-3514.51.6.1173.
- Biesanz, J. C., Falk, C. F., & Savalei, V. (2010). Assessing mediational models: Testing and interval estimation for indirect effects, *Multivariate Behavioral Research*, 45(5), 661-701. DOI: 10.1080/00273171.2010.498292.
- Bind, M.-A. C., Vanderweele, T. J., Coull, B. A., Schwartz, J. D. (2016). Causal mediation analysis for longitudinal data with exogenous exposure. *Biostatistics*, 17(1), 122 – 134. DOI: 10.1093/biostatistics/kxv029.
- Breitborde, N. J. K., Wastler, H., Pine, J. G., & Moe, A. M. (2020). Suicidality and social problem-solving skills among individuals with first-episode psychosis participating in Coordinated Specialty Care. *Early Intervention in Psychiatry*, 15, 497 – 504. DOI: 10.1111/eip.12967.

Within- or Between-Subject Design for Mediation

Charlton, A., Montoya, A. K., Price, J., & Hilgard, J. (2021). Noise in the process: an assessment of the evidential value of mediation effects in marketing journals.

<https://doi.org/10.31234/osf.io/ck2r5>

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Hillsdale, NJ: Routledge. DOI: 10.4324/9780203771587.

Craig, C. C. (1936). On the frequency function of xy . *The Annals of Mathematical Statistics*, 7(1), 1 – 15.

Daniel, R. M., De Stavola, B. L., Cousens, S. N., Vansteelandt, S. (2015). Causal mediation with multiple mediators. *Biometrics*, 71(1), 1 – 15. DOI: 10.1111/biom.12248.

Efron, B., & Tibshirani, R. J. (1994). *An Introduction to the Bootstrap*, Chapman & Hall/CRC. DOI: 10.1201/9780429246593.

Enders, C. K., & Tofighi, D. (2007). Centering predictor variables in cross-sectional multilevel models: A new look at an old issue. *Psychological Methods*, 12(2), 121 – 138, DOI: 10.1037/1082-989X.12.2.121.

Finch, J. F., West, S. G., & MacKinnon, D. P. (1997). Effects of sample size and nonnormality on the estimation of mediated effects in latent variable models. *Structural Equation Modeling*, 4, 87-107. DOI: 10.1080/10705519709540063.

Fisher, A. C., & Wallenstein, S. (1981). Crossover designs in medical research. In C. R. Buncher & J. Y. Tsay (Eds.), *Statistics in the pharmaceutical industry* (pp. 139-159). New York: Dekker.

Fritz, M. S., & MacKinnon, D. P. (2007). Required sample size to detect the mediated effect. *Psychological Science*, 18(3), 233-239. DOI: 10.1111/j.1467-9280.2007.01882.x.

Within- or Between-Subject Design for Mediation

- Goodboy, A. K., Dillow, M. R., Knoster, K. C., & Howard, H. A. (2021). Relational turbulence from the COVID-19 pandemic: Within-subjects mediation by romantic partner interdependence. *Journal of Social and Personal Relationships*, 38(6), 1800 – 1818. DOI: 10.1177/02654075211000135
- 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). DOI:10.1037/bul0000315.
- Greenwald, A. G. (1976). Within-subjects designs: To use or not to use? *Psychological Bulletin*, 83(2), 314 – 320. DOI: 10.1037/0033-2909.83.2.314.
- Grice, G. R. (1966). Dependence of empirical laws upon the source of experimental variation. *Psychological Bulletin*, 66(6), 488-498. DOI: 10.1037/h0023914.
- 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. DOI: 10.1177/0956797613480187.
- Hayes, A. F., Montoya, A. K., & Rockwood, N. J. (2017). The analysis of mechanisms and their contingencies: PROCESS versus structural equation modeling. *Australasian Marketing Journal*, 25(1), 76 – 81. DOI: 10.1016/j.ausmj.2017.02.001.
- Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, 15(4), 309-334. DOI: 10.1037/a0020761.
- Imai, K., Keele, L., Yamamoto, T. (2010). Identification, inference and sensitivity analysis for causal mediation effects. *Statistical Science*, 25(1), 51 – 71. DOI: 10.1214/10-STS321

Within- or Between-Subject Design for Mediation

- Josephy, H., Vansteelandt, S., Vanderhasselt, M.-A., & Loeys, T. (2015). Within-subject mediation analysis in AB/BA crossover designs. *International Journal of Biostatistics*, 11(1), 1 – 22. DOI: 10.1515/ijb-2014-0057.
- Judd, C. M., & Kenny, D. A. (1981). Process analysis: Estimating mediation in treatment evaluations. *Evaluation Review*, 5(5), 602-619. DOI: 10.1177/0193841X8100500502.
- Judd, C. M., Kenny, D. A., & McClelland, G. H. (2001) Estimating and testing mediation and moderation in within-subject designs. *Psychological Methods*, 6(2), 115-134. DOI: 10.1037//1082-989X.6.2.115.
- Judd, C. M., & McClelland, G. H. (1989). *Data analysis: A model comparison approach*. New York: Harcourt Brace Jovanovich.
- Kenny, D. A. (2017, February). MedPower: An interactive tool for the estimation of power in tests of mediation [Computer software]. Available from <https://davidakenny.shinyapps.io/MedPower/>.
- Kraemer, H. C., Wilson, T., Fairburn, C. G., Agras, S. (2002). Mediators and moderators of treatment effects in randomized clinical trials. *Archives of General Psychiatry*, 59, 877 – 883. DOI: 10.1001/archpsyc.59.10.877.
- Kreft, I. G., De Leeuw, J., & Aiken, L. S. (1995). The effect of different forms of centering in hierarchical linear models. *Multivariate Behavioral Research*, 30 (1), 1- 21. DOI: 10.1207/s15327906mbr3001_1.
- Ledgerwood, A., & Shrout, P.E. (2011). The trade-off between accuracy and precision in latent variable models of mediation processes. *Journal of Personality and Social Psychology*, 101, 1174–1188. DOI: 10.1037/a0024776.

Within- or Between-Subject Design for Mediation

- Looney, M. A., Feltz, C. J., & VanVleet, C. N. (1994). The reporting and analysis of research findings for within-subject designs: Methodological issues for meta-analysis. *Research Quarterly for Exercise and Sport*, 65(4), 363 – 366. DOI: 10.1080/02701367.1994.10607641.
- MacKinnon, D. P., Valente, M. J., & Gonzalez, O. (2020) The correspondence between causal and traditional mediation analysis: The link is the mediator by treatment interaction. *Prevention Science*, 21, 147-157. DOI: 10.1007/s11121-019-01076-4.
- MacKinnon, D. P., & Pirlott, A. G. (2015). Statistical approaches for enhancing causal interpretation of the M to Y relation in mediation analysis. *Personality and Social Psychology Review*, 19(1), 30-43. DOI: 10.1177/1088868314542878.
- Magson, N. R., Freeman, J. Y. A., Rapee, R. M., Richardson, C. E., Oar, E. L., & Fardouly, J. (2021) Risk and protective factors for prospective changes in adolescent mental health during the COVID-19 pandemic. *Journal of Youth and Adolescence*, 50, 44–57. DOI: 10.1007/s10964-020-01332-9
- Maxwell, S. E., Delaney, H. D. & Kelley, K. (2018). *Designing experiments and analyzing data: A model comparison perspective* (3 ed.), Mahwah, NJ: Lawrence Erlbaum Associates.
- Molenaar, P. C. M., (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement*, 2(4), 201 – 218. DOI: 10.1207/s15366359mea0204_1.
- Montoya, A. K. & Hayes, A. F. (2017). Two-condition within-participant statistical mediation analysis: A path-analytic framework. *Psychological Methods*, 22(1), 6 – 27. DOI: 1082-989X/17.

Within- or Between-Subject Design for Mediation

Montoya, A. K. (2018). Conditional process analysis in two-instance repeated-measures designs.

(Unpublished dissertation). The Ohio State University, Columbus, OH. Retrieved from

http://rave.ohiolink.edu/etdc/view?acc_num=osu1530904232127584.

Muller, K. E., LaVange, L. M., Ramey, S. L., & Ramey, C. T. (1992). Power calculations for general linear multivariate models including repeated measures applications. *Journal of the American Statistical Association*, 87(420), 1209 – 1266.

DOI:10.1080/01621459.1992.10476281.

Muller K. E., & Peterson B. L. (1984). Practical methods for computing power in testing the multivariate general linear hypothesis. *Computational Statistics and Data Analysis*, 2, 143–158. DOI: 10.1016/0167-9473(84)90002-1.

Nguyen, T., Carnevale, J. J., Scholer, A. A., Miele, D. B., & Fujita, K. (2019, May 23).

Metamotivational Knowledge of the Role of High-Level and Low-Level Construal in Goal-Relevant Task Performance. *Journal of Personality and Social Psychology*.

Advance online publication. DOI: 10.1037/pspa0000166.

Pazda, A. D., & Thorstenson, C. A. (2019). Color intensity increases perceived extraversion and openness for zero-acquaintance judgements. *Personality and Individual Differences*, 147, 118 – 127. DOI: 10.1016/j.paid.2019.04.022.

Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*. San Francisco, CA: Morgan Kaufman. DOI: 10.5555/2074022.2074073.

Poulton, E. (1973). Unwanted range effects from using within-subject experimental designs. *Psychological Bulletin*, 80(2), 113. DOI: 10.1037/h0034731.

Within- or Between-Subject Design for Mediation

- Preacher, K. J. (2015). Advances in mediation analysis: A survey and synthesis of new developments. *Annual Review of Psychology*, 66, 825 – 852. DOI: 10.1146/annurev-psych-010814-015258.
- Preacher, K. J., & Hayes, A. F. (2004) SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, & Computers*, 36(4), 717-731. DOI: 10.3758/BF03206553.
- Preacher, K. J., Zyphur, M. J., & Zhang, Z. (2010). A general multilevel SEM framework for assessing multilevel mediation. *Psychological Methods*, 15(3), 209 – 233. DOI: 10.1037/a0020141.
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 1). Sage.
- Robins, J. M., & Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(2), 143 – 155. DOI: 10.1097/00001648-199203000-00013.
- Rockwood, N. J. & Hayes, A. F. (2022). Multilevel mediation analysis. To appear in A. A. O’Connell, D. B. McCoach, & B. Bell (Eds.), *Multilevel Modeling Methods with Introductory and Advanced Applications*. Charlotte, NC: Information Age Publishing.
- Rockwood, N. J. (2018). Advancing the formulation and testing of multilevel mediation and moderated mediation models. (Unpublished master’s thesis). The Ohio State University, Columbus, OH. Retrieved from http://rave.ohiolink.edu/etdc/view?acc_num=osu1489578419777238.
- Rothstein, L. D. (1974). Reply to Poulton. *Psychological Bulletin*, 81(3), 199 – 200. DOI: 10.1037/h0035841.

Within- or Between-Subject Design for Mediation

Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions.

Journal of the American Statistical Association, 100(469), 322 – 331. DOI:

10.1198/016214504000001880

Senn, S. (1992). Is the ‘simple carry-over’ model useful? *Statistics in Medicine*, 11(6), 715–726.

DOI: 10.1002/sim.4780110603.

Senn, S. (1993). *Cross-over trials in clinical research*. Ed 2. John Wiley & Sons, LTD. DOI:

10.1002/0470854596.

Schoemann, A. M., Boulton, A. J., & Short, S. D. (2017). Determining power and sample size for simple and complex mediation models. *Social Psychological and Personality Science*, 8(4), 379-386. DOI: 10.1177/1948550617715068.

Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Houghton Mifflin Company: Boston, MA.

Shrout, P. E., & Bolger, N. (2002). Mediation in experimental and nonexperimental studies: new procedures and recommendations. *Psychological Methods*, 7, 422-445.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359 – 1366. DOI: 10.1177/0956797611417632.

Stone, C. A., & Sobel, M. E. (1990). The robustness of estimates of total indirect effects in covariance structure models estimated by maximum likelihood. *Psychometrika*, 55(2), 337-352. DOI: 10.1007/BF02295291.

Thorstenson, C. A., Pazda, A. D., & Lichtenfeld, S. (2020). Facial blushing influences perceived embarrassment and related social functional evaluations. *Cognition and Emotion*, 34(3), 413-426. DOI: 10.1080/02699931.2019.1634004.

Within- or Between-Subject Design for Mediation

- Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). Mediation: R package for causal mediation analysis. *Journal of Statistical Software*, 59(5). DOI: 10.18637/jss.v059.i05.
- Valente, M. J., & MacKinnon, D. P. (2017). Comparing models of change to estimate the mediated effect in the pretest-posttest control group design. *Structural Equation Modeling: A Multidisciplinary Journal*, 24(3), 428 – 450. DOI: 10.1080/10705511.2016.1274657.
- VanderWeele, T. J. (2010). Bias formulas for sensitivity analysis for direct and indirect effects. *Epidemiology*, 21(4), 540 – 551. DOI: 10.1097/EDE.0b013e3181df191c
- VanderWeele, T. J. (2015). *Explanation in causal inference: Methods for mediation and interaction*. Oxford University Press. DOI: 10.3102/1076998617698112.
- VanderWeele, T. J., & Chiba, Y. (2014). Sensitivity analysis for direct and indirect effects in the presence of exposure-induced mediator-outcome confounders. *Epidemiology, biostatistics and public health*, 11(2), e9027. DOI: 10.2427/9027.
- Venter, A. & Maxwell, S. E. (1999). Maximizing power in randomized designs in N is small. In R. H. Hoyle (Ed.), *Statistical Strategies for Small Sample Research*. Thousand Oaks: Sage.
- Whillans, A. V., Dunn, E. W., Smeets, P., Bekkers, R., & Norton, M. I. (2017). Buying time promotes happiness. *Proceedings of the National Academies of Science*, 144(32), 8523 – 8527. DOI: 10.1073/pnas.1706541114.
- Yuan, Y., & MacKinnon, D. P. (2014). Robust mediation analysis based on median regression. *Psychological Methods*, 19, 1-20. DOI: 10.1037/a0033820.

Within- or Between-Subject Design for Mediation

Yzerbyt, V., Muller, D., Batailler, C., & Judd, C. M. (2018). New recommendations for testing indirect effects in mediational models: The need to report and test component paths.

Journal of Personality and Social Psychology: Attitudes and Social Cognition, 115(6), 929 – 943. DOI: 10.1037/pspa0000132.

Zhang, Z., & Yuan, K. H. (2018). *Practical Statistical Power Analysis Using Webpower and R* (Eds). Granger, IN: ISDSA Press.

Zu, J., & Yuan, K.-H. (2010). Local influence and robust procedures for mediation analysis.

Multivariate Behavioral Research, 45(1), 1 – 44. DOI: 10.1080/00273170903504695.

Zyphur, M. J., Zhang, Z., Preacher, K. J., & Bird, L. J. (2018). Moderated mediation in multilevel structural equation models: Decomposing effects of race on math achievement within versus between high schools in the United States. In S. E. Humphrey & J. M. LeBreton (Eds.), *The handbook of multilevel theory, measurement, and analysis* (pp. 473-494). Washington, D. C.: American Psychological Association.

Within- or Between-Subject Design for Mediation

Table 1

Empirical power at $\alpha = .05$, in within-subject designs across sample size, size of indirect effect, and whether or not moderation was allowed in the model for the outcome.

<i>ab</i>	N = 20		N = 50		N = 100		N = 200	
	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod
0.02	.02	.02	.05	.04	.12	.09	.27	.21
0.05	.06	.05	.17	.14	.34	.29	.55	.49
0.08	.10	.09	.24	.21	.41	.35	.59	.53
0.15	.19	.17	.50	.46	.77	.73	.92	.89
0.23	.29	.27	.64	.61	.85	.82	.94	.92
0.35	.42	.42	.78	.77	.92	.90	.97	.97

Within- or Between-Subject Design for Mediation

Figure 1.

Path diagram of simple mediation model two-condition between-subject design. Model for Y using M and X can be estimated using either Equation 2a (coefficients not in parentheses and excluding dashed line in estimation) or 2b (coefficients in parentheses and including dashed line in estimation).

Within- or Between-Subject Design for Mediation

Figure 2. Path diagram of simple mediation model for two-condition within-subject design.

Model for $Y_1 - Y_2$ using $M_2 - M_1$ and intercept can be estimated using either Equation 9a (coefficients not in parentheses and excluding dashed line in estimation) or 9b (coefficients in parentheses and including dashed line in estimation).

Note: A 1 contained in a triangle indicates an intercept. The ^c indicates grand-mean centered.

Within- or Between-Subject Design for Mediation

Figure 3. *Empirical power at $\alpha = .05$, across design, sample size, and indirect effect when moderation is estimated*

Within- or Between-Subject Design for Mediation

Figure 4. Empirical power at $\alpha = .05$, across levels of correlation among mediators (top panel) and correlation among outcomes (bottom panel).

Within- or Between-Subject Design for Mediation

Figure 5. Empirical power at $\alpha = .05$, across moderation parameter, b -path, and estimation of moderation.

Within- or Between-Subject Design for Mediation

Appendix: Covariance Matrix Derivation for Simulation

To generate data for the Monte Carlo simulation, the covariance matrix for M_1 , M_2 , Y_1 , and Y_2 was derived. All variables had a variance of 1, to allow coefficients to have a standardized interpretation. Equations 6 and 7 can be rewritten as:

$$Y_{1i} = g_{10} + (b + d)M_{1i} + 0M_{2i} + e_{Y_{1i}}$$

$$Y_{2i} = g_{20} + 0M_{1i} + (b - d)M_{2i} + e_{Y_{2i}}$$

where g_{11} and g_{21} are replaced with combinations of b , which is the average relationship between M and Y in the same conditions ($\frac{g_{11}+g_{21}}{2}$), and d , which is half the difference in the relationships between M and Y in the same conditions ($\frac{g_{11}-g_{21}}{2}$). Additionally, the new versions of these equations make very explicit that the mediators only predict the outcome in the same condition and not the other condition (i.e., no carry-over effects).

A basic property from multivariate regression can be used: $\beta = \Sigma_{YZ}\Sigma_{ZZ}^{-1}$ where β is a $m \times r$ matrix of regression coefficients for predicting m outcome variables stored in the $N \times m$ matrix Y using r predictor variables stored in the $N \times r$ matrix Z . Σ_{YZ} is an $m \times r$ covariance matrix between the outcome and the predictors, whereas Σ_{ZZ} is the $r \times r$ covariance matrix for the predictors. From this Σ_{YZ} can be solved for:

$$\Sigma_{YZ} = \beta\Sigma_{ZZ} = \begin{bmatrix} b+d & 0 \\ 0 & b-d \end{bmatrix} \begin{bmatrix} 1 & \rho_M \\ \rho_M & 1 \end{bmatrix} = \begin{bmatrix} b+d & \rho_M(b+d) \\ \rho_M(b-d) & b-d \end{bmatrix}.$$

Thus, the full covariance matrix for M_1 , M_2 , Y_1 , and Y_2 is:

$$\begin{array}{cccc} M_1 & M_2 & Y_1 & Y_2 \\ \begin{bmatrix} 1 & \rho_M & b+d & \rho_M(b-d) \\ \rho_M & 1 & \rho_M(b+d) & b-d \\ b+d & \rho_M(b+d) & 1 & \rho_Y \\ \rho_M(b-d) & b-d & \rho_Y & 1 \end{bmatrix} \end{array}$$

Within- or Between-Subject Design for Mediation

From each population covariance matrix and pattern of means, samples were generated using a Cholesky decomposition with standard normal deviates. After the variables were generated, M_2 and Y_2 were linearly transformed to set a and c to their assigned values. For example, M_2 would be calculated by taking the M_2 which results from the Cholesky decomposition and subtracting a from all values of M_2 . These transformations do not change the covariance structure of the four variables, and they also allow for sampling variance in a and c , as the sample differences will not be exact each time (due to sampling variance in the mean differences between M_1 and M_2 as well as Y_1 and Y_2).

Notably, this derivation also gives key insights into which conditions in the simulation did not result in positive definite matrices. When the sum of b and d is high, the covariance matrix will not be positive definite, so these conditions were not run. Additionally, when the correlation among the mediators is high, this can result in a non-positive definite matrix.

**Supplemental Materials: Selecting a Within- or Between-Subject Design for Mediation:
Validity, Causality, and Statistical Power**

Comparing Inferential Methods.....	1
Type I Error Simulations Results.....	16
Comparison to Yzerbyt et al. (2018).....	25
Literature Review of Correlations Among Repeated Measurements	31
Using MedPower for Between-Subject Power Analysis in Empirical Example	38
R script for Power Analysis in Empirical Example	40
Supplementary References.....	43

Within- or Between-Subject Design for Mediation

Comparing Inferential Methods

In the simulation study, multiple methods of inferences were tested, but only percentile bootstrap confidence intervals are discussed in the main manuscript. In this section I overview the results using all inferential methods and comparing among the inferential methods.

Methods of Inference for Indirect Effects

There are many methods of inference for between-subject mediation and the growing body of research examining the performance of these different methods mostly focuses on between-subject designs. A variety of simulation studies have compared the performance of different tests of inference for the indirect effect in between-subject designs (Biesanz, Falk, Savalei, 2010; Fritz & MacKinnon, 2007; Hayes & Scharkow, 2013; MacKinnon et al., 2002; Stone & Sobel, 1990; Yzerbyt, Muller, Batailler, & Judd, 2018), but only Yzerbyt et al. (2018) has systematically studied the performance of inferential methods in within-subject mediation analyses. All the methods for testing indirect effects described in this section can be implemented for within-subject designs with MEMORE, a freely-available macro for SPSS and SAS (Montoya & Hayes, 2017).

Causal Steps Method

The causal steps method was introduced by Baron & Kenny (1986; See also Judd & Kenny, 1981). According to this method, the researcher has evidence for mediation if the *c*-path, *a*-path, and *b*-path are all significantly different from zero. For two-condition, within-subject designs, Judd, Kenny, and McClelland (2001) provided instructions on how to conduct these steps using Equation 8 – 10: just as in Baron and Kenny (1986), the *c*-, *a*-, and *b*-paths must all be statistically significant to support a claim of mediation.

Within- or Between-Subject Design for Mediation

This method has been criticized for a variety of reasons: one concern is the requirement that the total effect (c -path) be significant to make claims of mediation (Collins et al., 1998; Kenny, Kashy, Bolger, 1998; MacKinnon, 2008; MacKinnon et al., 2000; Shrout & Bolger, 2002). It is possible for the direct and indirect effects to be opposing in sign (e.g., one is positive and the other is negative), producing a non-significant total effect. This pattern does not indicate a lack of mediation, but rather that opposing forces are operating through X , making the total effect small or nonexistent (MacKinnon, 2008; O'Rourke & MacKinnon, 2018; Shrout & Bolger, 2002). Additionally, there can be greater statistical power to detect an indirect effect compared to a total effect even when the magnitude of the two effects is the same (i.e., $ab = c$, Judd & Kenny, 2014). The causal steps method does not provide estimates of the indirect effect or a standard error of the indirect effect, nor can a researcher create a confidence interval for the indirect effect using this method (Hayes, 2018; MacKinnon, 2008; MacKinnon, et. al., 2002). Most simulation studies exclude the causal steps method because of these limitations. One study which did examine the causal steps method found it had conservative empirical type I error and was underpowered (MacKinnon et al., 2002). The causal steps method is included in the current simulation, as it was the only existing inferential approach for within-subject designs prior to 2017.

Test of Joint Significance

The test of joint significance is very closely related to the causal-steps method (Kenny et al., 1998; MacKinnon et al., 2002). This test does not impose the requirement that the total effect be significantly different from zero, rather only the a - and b -paths are required to be significantly different from zero. In between-subject designs, significance tests on a and b would be generated from the estimation of Equations 1 and 2. A similar process can be done for within-

Within- or Between-Subject Design for Mediation

subject designs but using estimates from Equations 8 and 9 instead. If a and b are both significantly different from zero, there is evidence for mediation.

This method eliminates some but not all criticisms of the causal steps approach: there is still no estimate of the indirect effect or its standard error, no single p -value, or a confidence interval (but see Yzerbyt, et al., 2018). The joint significance test has been shown to have conservative empirical type I error but good empirical power in a variety of conditions for both between- and within-subject designs (Biesanz, Falk, & Savalei, 2010; MacKinnon et al., 2002; Yzerbyt, et al., 2018).

Normal Theory Tests

The normal theory tests (commonly known as Sobel tests or delta method) are inferential methods for the indirect effect (ab) that assume that the sampling distribution of ab is normal. In a between-subject design, normal theory tests use an asymptotic standard error derived from the multivariate delta method to compute a Z -score (ab/se_{ab}) then a p -value, which is compared to a previously determined significance level (e.g., $\alpha = .05$). There are multiple derived standard errors for this test, most popularly the Sobel (Equation S1; Sobel, 1982, 1986) and Aroian standard errors (Equation S2; Aroian, 1947) based on the first- and second-degree Taylor expansions, respectively.

$$se_{ab1} = \sqrt{a^2 se_b^2 + b^2 se_a^2} \quad (S1)$$

$$se_{ab2} = \sqrt{a^2 se_b^2 + b^2 se_a^2 + se_b^2 se_a^2} \quad (S2)$$

In a between-subject design, the a - and b -path estimates along with their standard errors (se_a and se_b respectively) would come from estimating Equations 1 and 2. Within-subject

Within- or Between-Subject Design for Mediation

designs can apply Equation S1 or S2 to compute the estimated standard error of the indirect effect using information from estimating Equations 8 and 9. From that point, calculation of a Z -score and p -value follow just as in the between-subject design.

One of the main criticisms for the normal theory tests is the assumption of normality of the indirect effect, which has been shown to be violated in finite samples (Stone & Sobel, 1990); however, MacKinnon and Dwyer (1993) showed that both the Sobel and Aroian estimates of the standard error were very close to simulated population standard errors in samples of size 50 for one set of population parameters a , b , and c' (see also Stone & Sobel, 1990). Though the standard errors may be accurate, the shape of the distribution is still non-normal; the distribution of the indirect effect has skew and kurtosis that can be calculated as a function of the population parameters for the a - and b -paths (Craig, 1936). In between-subject designs, normal theory tests have lower empirical power than alternative methods (Biesanz, Falk, & Savalei, 2010; Fritz & MacKinnon, 2007; MacKinnon, Lockwood, and Williams, 2004). Additionally, normal theory tests are slightly biased, particularly in small samples and with small non-zero indirect effects (MacKinnon, Lockwood, & Williams, 2004; Stone & Sobel, 1990).

Bootstrap Confidence Intervals

See main manuscript for description of percentile bootstrap confidence intervals.

Monte Carlo Confidence Intervals

The final method of inference discussed in this paper is the Monte Carlo confidence interval (Preacher & Selig, 2012). This method uses the estimates of the a - and b -paths and their standard errors to simulate sampling distributions for both a and b . These simulated estimates are then used to create an estimated sampling distribution for the indirect effect. To create a Monte Carlo confidence interval, a sample of size K , where K is recommended to be large (e.g. 5,000),

Within- or Between-Subject Design for Mediation

is generated from a random normal distribution with mean a , where a is the point estimate of the a -path from the observed sample, and standard deviation se_a . These estimates approximate the sampling distribution of the a -path under the assumptions of linear regression. Next, a sample of size K is generated from a random normal distribution with mean b , where b is the point estimate of the b -path from the observed sample, and standard deviation se_b . These estimates approximate the sampling distribution for the b -path under the assumptions of linear regression. The estimates of the a -path and b -path are then multiplied together listwise, such that the first estimate from the a -distribution is multiplied by the first estimate from the b -distribution, to generate a sample of K estimates of the indirect effect. This distribution is used as an approximation of the sampling distribution of the indirect effect. The estimates of the indirect effect are sorted from smallest to largest, selecting the upper and lower bounds for the confidence interval such that $100(\frac{\alpha}{2})\%$ of the estimates lie below the lower bound and $100(1 - \frac{\alpha}{2})\%$ of the estimates lie below the upper bound of the confidence interval.

Similar to the extension of the normal theory tests in within-subject designs, extending the Monte Carlo confidence interval to a within-subject design only requires the researcher to use estimates of the paths and the standard errors of the paths that are appropriate to the design. As such, a would be the mean difference between the two mediator variables, and b would be the partial regression coefficient b from Equation 9. The standard errors are the same as those to be used in calculating the asymptotic variance of the indirect effect for the normal theory tests (Equations 11 and 12). The rest of the Monte Carlo procedure can be carried out as it was in the between-subject case.

Monte Carlo confidence intervals have a few unique advantages and disadvantages. First, they do not require the original data to generate confidence intervals for the indirect effect: only

Within- or Between-Subject Design for Mediation

the path estimates and their standard errors. However, Monte Carlo confidence intervals rely on the assumption that the distributions of the a -path and b -path are normal, which is true in large enough samples without excessive violations of the assumptions of linear regression. Any violations of regression assumptions that would bias the estimates of the standard errors (e.g., outliers, heteroskedasticity) would influence the Monte Carlo confidence interval method. The Monte Carlo confidence interval allows for the unusual shape of the sampling distribution of the indirect effect and works well as an inferential test for the indirect effect in the between-subject mediation literature (Hayes & Scharkow, 2013; MacKinnon, Lockwood, & Williams, 2004; Preacher & Selig, 2012). The Monte Carlo confidence interval and percentile bootstrap confidence interval have been shown to have lower than expected empirical type I error and lower empirical power than some other methods, suggesting that the confidence intervals are too wide, resulting in slightly more conservative tests and lower power (Hayes & Scharkow, 2013; MacKinnon, Lockwood, Williams, 2004).

Many previous studies have compared the relative performance of methods of inference for mediation in between-subject designs, but only one previous study has done so for within-subject designs (Yzerbyt et al., 2018). Previous research has not fully explored the parameter space relevant to researchers considering using within-subject designs. The present study expands on Yzerbyt, et al. (2018) primarily by comparing within- and between-subject designs, but also by exploring additional inferential methods and a broader range of sample sizes, varying correlation among repeated-measurements and the moderation parameter, and comparing models with and without the moderation parameter. Methods of inference are varied for both between- and within-subject designs as there is no single inferential method that is universally used for detecting indirect effects. The present study adds the causal steps approach as proposed by Judd,

Within- or Between-Subject Design for Mediation

Kenny, and McClelland (2001) and normal theory tests using both first and second-order Taylor expansions. These methods are important to investigate because of their common use, particularly the causal steps method, as it was the only proposed method of inference for mediation in within-subject designs up until 2017. Correlation among repeated measurements was varied because power for detecting mean differences in within-subject designs depends on this correlation, a factor that was ignored in Yzerbyt et al. (2018). No previous research in either the between-subject or within-subject literature has examined the impact of the d parameter on the power to detect the indirect effect, which defines the differences between Equations 2a vs. 2b and 9a vs. 9b. It is possible that by estimating this parameter when it is not needed, researchers could lose statistical power, but they also must consider the impact of omitting this parameter if it is non-zero. The moderation parameter, d , was always set to 0 in Yzerbyt et al. (2018); in this simulation, d is systematically varied and its impact on empirical type I error and power is investigated. Finally, and most importantly, by simulating data in a way that is equivalent for between- and within-subject designs, this simulation allows for direct comparison of empirical power in each of the two designs and allows researchers to decide which design to use for their own study.

Methods

For each model in each sample, I applied the six tests of inference for the indirect effect as described in the Method of Inference for Indirect Effects section. I used six inferential tests of the indirect effect, each of which can be conducted in either between- or within-subject designs (as described in the previous section). Bootstrap and Monte Carlo confidence intervals were generated at the 95% level. For the causal steps method and joint significance tests, the individual component tests were conducted at the 5% level. Normal theory tests were conducted

Within- or Between-Subject Design for Mediation

at the 5% level. I used two versions of the normal theory tests, each with different standard errors: the delta method first-order Taylor series approximation (Equation S1, Sobel, 1982) and the delta method second-order Taylor series exact solution (Equation S2, Aroian, 1947). I used both percentile bootstrap confidence intervals and Monte Carlo confidence intervals with 1,000 samples.

Results

Between vs. Within-Subject Designs. Table S1 provides description of the proportion of conditions where the within-subject or between-subject design had the power advantage, broken down by inferential method and whether moderation was estimated. When examining the conditions under which within-subject designs have greater power than between-subject designs, there are 315 (out of 6,192) conditions for which all within-subject methods had lower empirical power than their between-subject counterparts. The patterns of parameters which produced these conditions are the same as those presented in the primary manuscript: large a -paths, high correlations among the mediators, low correlations among the outcomes, and small moderation parameters.

Figure S1 shows empirical power for each inferential method in both between and within-subject designs across sample sizes, indirect effects, and with and without estimating moderation. The patterns for all inferential methods are similar to the percentile bootstrap confidence interval presented in the main manuscript.

Within-Subject Test Performance. In this section I discuss the results of the simulation study specifically looking at the within-subject test performance. Generally, I follow the organization of the main manuscript to aid comparison between the results with the percentile

Within- or Between-Subject Design for Mediation

bootstrap confidence interval and the other methods included here. First, we begin by comparing power across the methods, which is not described in the main manuscript.

Comparing Inferential Methods. Table S2 provides power estimates for each inferential method at different indirect effect levels, sample size, and whether moderation is estimated. Comparing methods of inference in this case is difficult because many methods had overly conservative empirical type I error, and very few methods were reliably close to the nominal 0.05 level. As such, comparisons among methods should account for the fact that the methods did not all have equal empirical type I error rates. Overall, the test of joint significance, percentile bootstrap confidence interval, and Monte Carlo confidence intervals have similarly high empirical power; note that these methods also had similar empirical type I error. The causal steps method and normal theory tests are under-powered (typically 4 – 7% lower), but also these were the methods with lowest empirical type I error rates.

Moderation Effect and Whether Moderation is Estimated. Figure S2 summarizes power across the range of the moderation parameter, *b*-path size, inferential method, and whether moderation was estimated or not. Summarizing across all cases when the moderation parameter was zero and comparing empirical power in each condition, between 56% (Causal Steps) and 69% (Bootstrap) of cases saw a power advantage when omitting the moderation coefficient when it was zero at the population level. Overall, the average empirical power gain is very small for each method, ranging from 0.0009 (Causal Steps) to 0.004 (Bootstrapping). These results do not notably differ from those presented in the main manuscript, leading to the recommendation that the moderation effect should be estimated unless researchers have a strong a priori hypothesis that this effect is zero.

Within- or Between-Subject Design for Mediation

Size of Indirect Effect. Table S2 and Figure S1 show that the power to detect the indirect effect increases as the indirect effect increases, and this is true across inferential methods.

Discussion

All tests performed in a relatively similar manner to previous simulation research with between-subject designs (for reference see Hayes & Scharkow, 2013 and MacKinnon et al., 2002). These results also replicate results from Yzerbyt et al., (2018), showing that the joint significance, Monte Carlo confidence interval, and percentile bootstrap confidence interval tend to have the most accurate empirical type I error rate (closest to .05) and highest empirical power. For a more in-depth comparison of Yzerbyt et al. (2018) and the current manuscript see page 25 of Supplementary Materials. Overall, the general pattern of findings from Yzerbyt et al. (2018) are replicated in this study and expanded upon by exploring additional factors such as correlation among repeated-measurements and the moderation parameter. Additionally, based on these simulation results, I recommend that researchers use a percentile bootstrap confidence interval, Monte Carlo confidence interval, or joint significance test for inferential decisions in within-subject mediation. Notably, the joint significance test does not provide a single p -value or a confidence interval, and so for that reason either of the interval estimates may be preferred. For reasons previously explained, the Monte Carlo confidence intervals and the joint significance test may be sensitive to common violations of multiple regression, such as heteroskedasticity; however, evaluation of these sensitivities is beyond the scope of this paper.

An important aspect of this study is the opportunity to compare the effect of selection of inferential method to the effect of changing design. Much research has focused on understanding which inferential methods have the best balance of empirical type I error and power (Biesanz, Falk, & Savalei, 2010; Fritz & MacKinnon, 2007; MacKinnon et al., 2002; Yzerbyt, et al., 2018),

Within- or Between-Subject Design for Mediation

but the difference between inferential methods is miniscule in comparison to the difference between designs. Researchers looking to increase statistical power in their mediation analyses should consider design first and consider the selection of inferential method as secondary. Especially when choosing among the joint significance test, percentile bootstrap confidence intervals, and Monte Carlo confidence intervals, the differences are very small.

Within- or Between-Subject Design for Mediation

Table S1.

Proportion of cases for which empirical power at $\alpha = .05$ is greater in each design.

Design	Method											
	CS		JS		Sobel		Aroian		Bootstrap		MC	
	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod
Within	.84	.81	.88	.86	.90	.86	.90	.87	.89	.86	.89	.85
Tie	.03	.03	.01	.01	.01	.01	.01	.02	.00	.01	.00	.00
Between	.13	.16	.12	.15	.08	.12	.09	.12	.11	.14	.11	.15

Note. CS = causal steps method. JS = joint significance test. MC = Monte Carlo confidence interval. Mod = Moderation estimated. No Mod = No moderation estimated.

Within- or Between-Subject Design for Mediation

Table S2

Empirical power at $\alpha = .05$, in within-subject design across sample size, size of indirect effect, method, and whether or not moderation was allowed in the model for the outcome.

ab	Method											
	CS		JS		Sobel		Aroian		Bootstrap		MC	
	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod
<i>N</i> = 20												
0.02	.01	<.01	<i>.01</i>	.01	<.01	<.01	<.01	<.01	.02	.02	.02	.01
0.05	.02	.02	<i>.05</i>	.04	.02	.02	.02	.01	.06	.05	.06	.05
0.08	.03	.03	<i>.08</i>	.07	.05	.04	.04	.03	.10	.09	.09	.08
0.15	.07	.06	<i>.16</i>	.14	.11	.09	.10	.08	.19	.17	.19	.16
0.23	.10	.09	<i>.26</i>	.23	.20	.17	.18	.16	.29	.27	.29	.27
0.35	.14	.14	<i>.39</i>	.38	.33	.31	.31	.30	.42	.42	.43	.42
<i>N</i> = 50												
0.02	.03	.02	.05	.04	.02	.01	.01	.01	.05	.04	<i>.05</i>	.04
0.05	.09	.07	<i>.16</i>	.13	.09	.07	.08	.06	.17	.14	.16	.13
0.08	.12	.11	<i>.23</i>	.19	.16	.13	.15	.12	.24	.21	.23	.20
0.15	.25	.22	<i>.49</i>	.44	.39	.33	.37	.31	.50	.46	.50	.45
0.23	.30	.29	<i>.63</i>	.59	.57	.52	.55	.50	.64	.61	.64	.60
0.35	.36	.35	<i>.77</i>	.75	.75	.73	.74	.72	.78	.77	.78	.77
<i>N</i> = 100												
0.02	.07	.06	.12	.09	.05	.03	.05	.03	.12	.09	<i>.11</i>	.09
0.05	.20	.17	<i>.33</i>	.28	.25	.19	.23	.18	.34	.29	.33	.28
0.08	.24	.21	<i>.40</i>	.34	.34	.29	.33	.28	.41	.35	.40	.35
0.15	.44	.42	<i>.77</i>	.73	.73	.67	.71	.66	.77	.73	.77	.73
0.23	.47	.46	<i>.84</i>	.81	.83	.80	.83	.79	.85	.82	.84	.82
0.35	.51	.50	<i>.91</i>	.90	.91	.90	.91	.90	.92	.90	.92	.90
<i>N</i> = 200												
0.02	.18	.15	.28	.22	.17	.12	.15	.10	.27	.21	<i>.27</i>	.21
0.05	.36	.32	<i>.55</i>	.49	.50	.42	.48	.41	.55	.49	.55	.49
0.08	.38	.34	<i>.58</i>	.52	.56	.50	.56	.50	.59	.53	.59	.53
0.15	.58	.56	<i>.92</i>	.89	.91	.88	.91	.88	.92	.89	.92	.89
0.23	.59	.58	<i>.94</i>	.92	.94	.92	.94	.92	.94	.92	.94	.92
0.35	.60	.60	<i>.97</i>	.97	.97	.97	.97	.97	.97	.97	.97	.97

Note. In each row the test with highest empirical power is highlighted in **bold**, the test with the second highest empirical power is highlighted in **bold italics**, and the test with the third highest empirical power is in *italics*. CS = causal steps method. JS = joint significance test. MC = Monte Carlo confidence interval. Mod = Moderation estimated. No Mod = No moderation estimated.

Within- or Between-Subject Design for Mediation

Figure S1. Empirical power at $\alpha = .05$, across design, method, sample size, and indirect effect

Within- or Between-Subject Design for Mediation

Figure S2. Empirical power at $\alpha = .05$, across moderation parameter, b -path, method, and estimation of moderation.

Within- or Between-Subject Design for Mediation

Type I Error Simulations Results

While the focus of the primary manuscript is on statistical power, it is also important to understand the type I error rates of inferential methods. This was explored in the simulation study, and the results are presented here. The type I error results for all inferential methods are presented in this section, and the power results for the inferential methods omitted from the main manuscript are presented in the next section.

The population indirect effect was zero in all conditions where the a -path or the b -path was set to zero at the population level. I calculated empirical type I error by computing the proportion of samples that were deemed significant by a specific inferential method in each condition where the population indirect effect equaled zero. Each method used an α -level of .05, so each method was expected to reject 5% of samples under conditions where there is no indirect effect. I used two common criteria to evaluate empirical type I error rates when $\alpha = .05$: Bradley's (1987) liberal criterion (.025 to .075) and Serlin's (2000) robustness criterion (.035 to .065). Empirical type I error rates within these ranges would be deemed acceptable under each criterion.

Empirical type I error. For each sample size there are 812 conditions where the population indirect effect is zero, with 1,000 samples per condition. The conditions are grouped by the type of true null hypotheses: either a and b are both zero, or a or b (but not both) is zero and the other path is non-zero (e.g., $a = .59$, $b = 0$). Empirical type I errors across these two null hypotheses are reported in Tables S3 and S4. Empirical type I errors outside of Bradley's criteria (worst performers) are marked in italics, and those inside Serlin's criteria are marked in bold (best performers). Note that errors within Serlin's criteria are also within Bradley's criteria, but the converse is not necessarily the case. In Tables S3 and S4, the inferential test with empirical

Within- or Between-Subject Design for Mediation

type I error rate closest to .05 is underlined on each row to examine trends of most accurate performance. In general, empirical type I error increases as sample size increases. No method exceeds Serlin's robustness criteria, indicating that no method has overly high empirical type I error, with the highest empirical type I error rate as 0.0593 (Percentile bootstrap, $N = 100$, $b = 0$, $a = .59$, no moderation estimated). Based on these results, it seems that all the tests are either accurate or conservative. Additionally, percentile bootstrap, Monte Carlo confidence intervals, and the joint significance test clearly perform most accurately with respect to having empirical type I error rates that are closest to the desired level of the test. Percentile bootstrap is closest to the desired level in 57% of cases (16/28), joint significance in 36% of cases (10/28), and Monte Carlo confidence interval in 7% of cases (2/28).

For simplicity this section focuses on the model with moderation, though all patterns are consistent across moderation and non-moderation models. In the simulation, when a and b are both zero, all tests are very conservative, with all empirical type I errors at or below .0039 (See Tables S3 and S4). Empirical type I error increases slightly as sample size increases for the causal-steps method and the test of joint significance (.0007 to .0017 and .0024 to .0025 respectively). Alternatively, the empirical type I error rate decreases for the normal theory tests (.0004 to .0001 [Sobel] and .0003 to .0001 [Aroian]), bootstrap confidence intervals (.0035 to .0016), and Monte Carlo confidence interval (.0028 to .0015) with increasing sample size (20 to 200). In comparison, the performance across tests varies more when a or b but not both are zero. Empirical type I error when both a and b are zero can be compared to when one path is zero and the other is large (e.g., $a = 0$, $b = .59$). For all tests, this difference is very large, often multiple orders of magnitude, and increases with increasing sample size. For example, the Monte Carlo confidence interval with sample size of 20 has an empirical type I error of .0028 when both a and

Within- or Between-Subject Design for Mediation

b are zero, but an empirical type I error of .0353 when $a = 0$ and $b = .59$ and .0444 when $a = .59$ and $b = 0$, a difference of .0325 ($a = 0$) and .0416 ($b = 0$). This difference is .0493 ($a = 0$) and .052 ($b = 0$) when sample size is increased to 200. When either a or b but not both are zero, the causal steps method and the two normal theory tests are more conservative than the test of joint significance, the percentile bootstrap confidence interval, and the Monte Carlo confidence interval, which all seem to approach the desired rejection rate of .05 as sample size increases.

Correlation among repeated measurements. Figure S3 provides a visualization of empirical type I error across the range of the mediator correlation for different sample sizes, methods, and coefficient paths for models with moderation (the patterns are all similar for models with no moderation). An interesting pattern emerges such that when $a = 0$ and $b \neq 0$, empirical type I error decreases as the correlation among the mediators increases, and this effect is stronger for larger b coefficients. Alternatively, when $b = 0$ and $a \neq 0$, the empirical type I error increases as the correlation among the mediators increases, and this effect is stronger for larger a -coefficients. These findings suggest that as the correlation among the mediators increases, detection of a true a -path increases but detection of a true b -path decreases. This effect is discussed more in-depth in the Discussion section.

Figure S4 visualizes the empirical type I error across the range of the outcome correlation for different sample sizes, methods, and coefficient paths for models with moderation (the patterns are all similar for models with no moderation). Overall, what is clear is that when $b = 0$ the correlation among the outcomes does not affect empirical type I error; however, when $a = 0$ and $b \neq 0$, as the correlation among the outcomes increases, so too does empirical type I error. For most methods, as the correlation increases the empirical type I error approaches .05, but as seen previously there are no cases where empirical type I error is much higher than .05.

Within- or Between-Subject Design for Mediation

Whether Moderation is Estimated. Comparing models that estimate moderation to those that do not estimate moderation with respect to empirical type I error, leads to an interesting pattern, which can be seen through column-wise comparisons in Tables S3 and S4. When $b = 0$ there is not a clear pattern of which model (moderation vs. no moderation) has higher empirical type I error. In a condition by condition comparison, three methods more frequently have higher empirical type I error in the moderation model than the no-moderation model (proportion of conditions with higher empirical type I error: 68.75%, 56.25%, and 62.5% for Sobel, Aroian, and Monte Carlo respectively). Three methods have higher empirical type I error more frequently for the no-moderation model (75%, 62.5%, and 75% for causal steps, joint significance, and bootstrap respectively). However, when $a = 0$, all methods overwhelmingly show that the moderation model has higher empirical type I error. All exceptions were cases when $b = 0$, except when $N = 20$, $b = .39$, for bootstrapping. Though this difference is consistent, it is typically very small, with a maximum of .0053 when $N = 200$, $b = 0.14$, for bootstrapping.

Moderation parameter. When examining empirical type I error across the range of the moderator parameter, two patterns emerge. First, when $b = 0$ the empirical type I error is unaffected by the moderation parameter, for both the models that estimate and don't estimate moderation. However, when $b \neq 0$ there is a clear difference between the models that estimate the moderation parameter and those that do not: When no moderation is estimated, as the moderation parameter increases the empirical type I error decreases slightly. However, when moderation is estimated, empirical type I error stays stable across the values of the moderation parameter, and sometimes increases slightly (See Figure S5). This suggests that if there is moderation, then estimating it will result in more accurate empirical type I error rates than not estimating it.

Within- or Between-Subject Design for Mediation

Table S3

Empirical type I error at $\alpha = .05$, when $a = 0$, across sample size, values of b , inferential method, and whether or not moderation was allowed in the model for the outcome.

<i>b</i>	Method											
	CS		JS		Sobel		Aroian		Bootstrap		MC	
	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod
<i>N</i> = 20												
0	.001	.001	.002	.002	<.001	<.001	<.001	.0002	.004	<u>.004</u>	.003	.003
0.14	.002	.002	.006	.005	.002	.001	.001	.0008	<u>.009</u>	.008	.008	.006
0.39	.007	.006	.019	.019	.009	.007	.007	.0055	<u>.025</u>	.023	.024	.021
0.59	.012	.011	.028	.027	.016	.015	.014	.013	.037	.037	.035	.034
<i>N</i> = 50												
0	.001	.001	<u>.002</u>	.002	<.001	<.001	<.001	<.001	.002	.002	.002	.002
0.14	.007	.005	.013	.010	.003	.002	.003	.002	<u>.013</u>	.010	.012	.009
0.39	.017	.016	.034	.031	.016	.013	.014	.011	.038	.034	.036	.033
0.59	.021	.020	.041	.040	.026	.025	.024	.022	.047	.046	.045	.044
<i>N</i> = 100												
0	.001	.001	.002	<u>.002</u>	<.001	<.001	<.001	<.001	.002	.002	.002	.002
0.14	.012	.010	<u>.020</u>	.016	.006	.004	.005	.003	.019	.015	.018	.015
0.39	.024	.023	.042	.039	.024	.021	.022	.019	.044	.041	.043	.040
0.59	.027	.027	.044	.044	.032	.031	.030	.029	.048	.047	.047	.046
<i>N</i> = 200												
0	.002	.002	.003	<u>.003</u>	<.001	<.001	<.001	<.001	.002	.002	.002	.002
0.14	.019	.016	<u>.029</u>	.024	.010	.007	.009	.006	.027	.022	.027	.022
0.39	.031	.030	.047	.046	.033	.030	.031	.028	.049	.047	.049	.047
0.59	.032	.032	.049	.048	.039	.038	.037	.036	.051	.051	.051	.050

Note. Empirical type I errors in *italics* are outside Bradley's liberal criteria (.025 – .075), and type I errors in **bold** are inside Serlin's robustness criteria (.035 – .065). In each row the test with the type I error rate closest to .05 is underlined. CS = causal steps method. JS = joint significance test. MC = Monte Carlo confidence interval. Mod = Moderation estimated. No Mod = No moderation estimated.

Within- or Between-Subject Design for Mediation

Table S4

Empirical type I error at $\alpha = .05$, when $b = 0$, across sample size, values of a , inferential method, and whether or not moderation was allowed in the model for the outcome.

<i>a</i>	Method											
	CS		JS		Sobel		Aroian		Bootstrap		MC	
	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod	Mod	No Mod
<i>N</i> = 20												
0	.001	.001	.002	.002	<.001	<.001	<.001	<.002	.004	<u>.004</u>	.003	.003
0.14	.002	.002	.006	.006	.002	.002	.001	.001	.009	<u>.010</u>	.008	.007
0.39	.007	.007	.023	.024	.011	.011	.010	.010	.033	.038	.029	.029
0.59	.010	.010	.035	.035	.022	.022	.020	.020	.049	.055	.044	.044
<i>N</i> = 50												
0	.001	.001	<u>.002</u>	.002	<.001	<.001	<.001	<.001	.002	.002	.002	.002
0.14	.006	.006	.013	.017	.002	.002	.002	.002	<u>.014</u>	<u>.014</u>	.011	.011
0.39	.017	.017	.037	.037	.018	.018	.016	.016	.045	.046	.039	.039
0.59	.022	.022	.047	.048	.030	.030	.027	.027	.059	.060	.052	.052
<i>N</i> = 100												
0	.001	.001	.002	<u>.002</u>	<.001	<.001	<.001	<.001	.002	.002	.002	.002
0.14	.011	.011	<u>.019</u>	.019	.004	.004	.003	.003	.019	.019	.017	.017
0.39	.026	.026	.045	.045	.025	.025	.023	.022	.051	.051	.047	.047
0.59	.028	.028	.050	.049	.038	.038	.035	.035	.059	.059	.054	.055
<i>N</i> = 200												
0	.002	.002	.003	<u>.003</u>	<.001	<.001	<.001	<.001	.002	.002	.002	.002
0.14	.018	.018	<u>.028</u>	.028	.010	.009	.008	.008	.027	.027	.026	.025
0.39	.031	.032	.050	.050	.034	.034	.032	.032	.055	.055	.052	.052
0.59	.032	.032	.050	.051	.044	.044	.042	.043	.056	.057	.054	.054

Note. Type I errors in *italics* are outside Bradley's liberal criteria (.025 – .075), and type I errors in **bold** are inside Serlin's robustness criteria (.035 – .065). In each row the test with the type I error rate closest to .05 is underlined. CS = causal steps method. JS = joint significance test. MC = Monte Carlo confidence interval. Mod = Moderation estimated. No Mod = No moderation estimated.

Within- or Between-Subject Design for Mediation

Figure S3. Empirical type I error at $\alpha = .05$, across the range of the correlation among the mediators, a -path, b -path, sample size, and inferential method in models with moderation.

Within- or Between-Subject Design for Mediation

Figure S4. Empirical *type I* error at $\alpha = .05$, across the range of the correlation among the outcomes, *a*-path, *b*-path, sample size, and inferential method in models with moderation.

Within- or Between-Subject Design for Mediation

Figure S5. Empirical type I error at $\alpha = .05$, when a -path is zero, across the range of the moderation parameter, b -path, sample size, and inferential method, and whether moderation is estimated or not.

Within- or Between-Subject Design for Mediation

Comparison to Yzerbyt et al. (2018)

Yzerbyt et al. (2018) is the closest published simulation study to the current study. In this appendix the information and results from this study and Yzerbyt et al. (2018) are compared. First the covariance structure for generating data in the within-subject design is described. Yzerbyt et al. (2018) generate data for the within-subject case using the following set of equations:

$$M_1 \sim N(0,1)$$

$$Y_1 \sim N(0,1)$$

$$M_2 = M_1 + a_{51} + e_{M_2} \quad e_{M_2} \sim N(0,1)$$

$$Y_2 = Y_1 + b_{62}(M_2 - M_1) + e_{Y_2} \quad e_{Y_2} \sim N(0,1)$$

Though not specified in Yzerbyt et al. (2018), I assumed that all standard normally distributed variables (i.e., M_1 , Y_2 , e_{M_2} , and e_{Y_2}) were independent. This leads to the conclusion that the covariance between M_1 and Y_1 is zero. Additionally, because M_1 and Y_1 are independent, and e_{M_2} and Y_1 are independent, M_2 and Y_1 are independent. Next to examine the variance of M_2 and its covariance with M_1 .

$$Var(M_2) = Var(M_1) + Var(e_{M_2}) = 1 + 1 = 2$$

$$Cov(M_2, M_1) = 1 \times Var(M_1) = 1 \times 1 = 1$$

Next, the relationship between all other variables and Y_2 is derived. Again, using the property $\Sigma_{YZ} = \beta \Sigma_{ZZ}$ where Y is Y_2 and Z is Y_1 , M_2 , and M_1 in that order.

$$\Sigma_{YZ} = \begin{bmatrix} 1 \\ b_{62} \\ -b_{62} \end{bmatrix}' \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 1 \\ 0 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ b_{62} \\ 0 \end{bmatrix}'$$

The variance of Y_2 is

Within- or Between-Subject Design for Mediation

$$\begin{aligned} Var(Y_2) &= Var(Y_1) + b_{62}^2(Var(M_1) + Var(M_2) - 2 Cov(M_1, M_2)) + Var(e_{Y_2}) \\ &= 1 + b_{62}^2(1 + 2 - 2) + 1 = 2 + b_{62}^2 \end{aligned}$$

This leads to the final solution for the variance-covariance matrix from Yzerbyt et al. (2018).

$$\begin{array}{cccc} M_1 & M_2 & Y_1 & Y_2 \\ \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 2 & 0 & b_{62} \\ 0 & 0 & 1 & 1 \\ 0 & b_{62} & 1 & 2 + b_{62}^2 \end{bmatrix} \end{array}$$

There are notable differences between the solution in the Appendix used for the current study, and that for Yzerbyt et al. (2018). First, Yzerbyt et al. (2018) did not generate variables with standard variances, and thus the coefficients used cannot be interpreted as standardized coefficients. In fact, because b_{62} appears in the variance for Y_2 this results in the variance of Y_2 being confounded with this coefficient. Given the variances and covariances of M_1 and M_2 this shows that the correlation between these two variables is fixed to about 0.7. Additionally, the correlation among the Y variables is confounded with the b_{62} parameter. So as the b_{62} parameter increases the correlation between the Y variables will decrease. The present study allows the effects of these two parameters on power to be estimated separately. Notably, the variances for M_1 and M_2 are not equal, which is different from what is assumed in Judd, Kenny, and McClelland (2001). This assumption is important for the orthogonality of the difference and sum scores used in Equation 9b. The fact that Yzerbyt et al. (2018) still found good performance of this method under unequal variances may suggest that perhaps this assumption is not needed, except that they only considered cases when the moderation parameter was zero. Future research should more systematically examine the impact of violations of this assumption and the correspondence between this assumption and the moderation parameter, as neither the current study nor Yzerbyt et al. (2018) examine this issue in depth.

Within- or Between-Subject Design for Mediation

Conditions from the current study were selected such that they would most closely align with the conditions from Yzerbyt et al. (2018) to allow for possible comparison. The closest approximation of Table S5 from Yzerbyt et al. (2018) is presented below in Table B1. From this study, only conditions where the moderation parameter is zero, the correlation among the mediators is 0.6, and the direct effect is zero are included, as these are the closest parameters to Yzerbyt et al. (2018). Additionally, because the correlation among the outcomes depends on the b parameter, the correlation condition that most closely aligns with that from Yzerbyt et al. (2018) was selected. For reference when b is 0, .14, .39, and .59 the correlation among the outcomes is 0.707, 0.703, 0.682, and 0.653 respectively. Because the closest condition observed in the present simulation study is $\rho_Y = 0.6$, this condition is used as an approximate. Only sample sizes of 20, 50, and 100 are used as these are closest to the sample sizes from Yzerbyt et al. (2018), who used 25, 50 and 100. Only inferential methods from Yzerbyt et al. (2018; joint significance, Monte Carlo confidence interval, and percentile bootstrap confidence interval) were used. Finally, only the models that estimate moderation were included, as this aligns with the model estimated in Yzerbyt et al. (2018).

Table S5.

Empirical type I error rates (Top Panel) and Power (Bottom Panel) at $\alpha = .05$, for Within-Subject Mediation Analysis by population coefficients, sample size, and inferential method, when correlation among mediators and outcomes are 0.7, the direct effect is 0, moderation is 0, and moderation is estimated.

		Type 1 Error						
		$a = 0$	$a = 0$	$a = .14$	$a = 0$	$a = 0.39$	$a = 0$	$a = .59$
		$b = 0$	$b = .14$	$b = 0$	$b = .39$	$b = 0$	$b = .59$	$b = 0$
$n = 20$								
	JS	.0018	.0045	.0043	.0177	.0255	.0410	.0415
	MC	.0020	.0050	.0053	.0205	.0328	.0522	.0558
	PB	.0023	.0048	.0075	.0230	.0345	.0525	.0630
$n = 50$								
	JS	.0030	.0065	.0105	.0368	.0475	.0560	.0528

Within- or Between-Subject Design for Mediation

MC	.0025	.0063	.0088	.0392	.0482	.0630	.0622		
PB	.0020	.0068	.0102	.0390	.0572	.0640	.0680		
<i>n</i> = 100									
JS	.0028	.0132	.0185	.0427	.0427	.0472	.0515		
MC	.0020	.0098	.0145	.0432	.0488	.0525	.0565		
PB	.0020	.0110	.0132	.0442	.0562	.0528	.0572		
Power									
	<i>a</i> = .14	<i>a</i> = .14	<i>a</i> = .14	<i>a</i> = .39	<i>a</i> = .39	<i>a</i> = .39	<i>a</i> = .59	<i>a</i> = .59	<i>a</i> = .59
	<i>b</i> = .14	<i>b</i> = .39	<i>b</i> = .59	<i>b</i> = .14	<i>b</i> = .39	<i>b</i> = .59	<i>b</i> = .14	<i>b</i> = .39	<i>b</i> = .59
<i>n</i> = 20									
JS	.0070	.0320	.0792	.0380	.1600	.3500	.0668	.3070	.6260
MC	.0085	.0380	.0965	.0452	.1780	.3870	.0838	.0350	.6740
PB	.0088	.0402	.0982	.0482	.1780	.3810	.0900	.3350	.6490
<i>n</i> = 50									
JS	.0262	.1600	.1860	.1320	.6790	.8500	.1560	.7970	.9930
MC	.0242	.1540	.2040	.1340	.6870	.8590	.1700	.8120	.9940
PB	.0290	.1580	.2080	.1400	.6890	.8630	.1740	.8090	.9940
<i>n</i> = 100									
JS	.0970	.3260	.3490	.2820	.9780	.9920	.2720	.9810	1.0000
MC	.0857	.3380	.3630	.2920	.9800	.9920	.2880	.9810	1.0000
PB	.0880	.3360	.3650	.3040	.9780	.9930	.2920	.9813	1.0000

Note. JS = Joint significance test. MC = Monte Carlo confidence interval. PB = percentile bootstrap.

Overall, the pattern of findings from Yzerbyt, et al. (2018) are replicated. The empirical type I error of the joint significance test is typically lower than .05 or very close; whereas the Monte Carlo confidence interval and the percentile bootstrap are sometimes higher than .05. The current study used Bradley's and Serlin's criteria to evaluate whether empirical type I error was "too high". According to those criteria, Monte Carlo confidence intervals and percentile bootstrap do not have empirical type I error that are "too high." There is only one case according to Serlin's criteria where empirical type I error is too high: percentile bootstrap where $a = 0.59$ and $b = 0$. By contrast, Yzerbyt et al. (2018) use significance tests to evaluate if empirical type I error rates are significantly different from .05. I view this to be an overly strict criterion, especially given its sensitivity to number of simulation replications run. With Yzerbyt and colleague's (2018) 10,000

Within- or Between-Subject Design for Mediation

simulation per condition, any empirical type I error rates higher than 0.0528 are considered too high. Using this criterion in the current simulation with 1,000 simulations per condition, any empirical type I error rate higher than 0.059 would be considered too high. Using this rule, the results of this study largely replicate the results of Yzerbyt et al. (2018) where in a few cases Monte Carlo confidence intervals and percentile bootstrapping had empirical type I error rates that were significantly higher than .05, and the joint significance test does not ever have this problem. However, while this does help account for Monte Carlo error in the simulation, it is specifically examining whether the true type I error rate is exactly .05 rather than whether the true type I error rate is near enough to .05 to be considered robust, which is what Serlin's and Bradley's criteria aim to do. Additionally, Yzerbyt et al. (2018) focus primarily on whether empirical type I error is significantly higher than .05 and do not consider conditions where empirical type I error is significantly lower than .05; whereas Serlin's and Bradley's criterion focus on type I error rates that are both too high and too low. Table S3 and S4 show that at lower sample sizes, bootstrap confidence intervals tend to have the empirical type I error rate closest to .05; whereas joint significance tests are closest when sample sizes are largest.

Regarding power, the results from the current simulation largely replicate those of Yzerbyt et al. (2018). The empirical power for the joint significance tests, percentile bootstrap, and Monte Carlo confidence interval are all very similar. However, while Yzerbyt et al. (2018) describe the joint significance test as “indistinguishable from the other two tests at the larger sample sizes”, the current study points to the pattern of highest empirical power across conditions as suggesting that while the empirical power may be similar, it is in fact distinguishable. Percentile bootstrap and Monte Carlo confidence intervals have higher empirical power than the joint significance test, except in a very interesting case of high sample size and

Within- or Between-Subject Design for Mediation

very small indirect effects. In this case it seems that the joint significance test may have higher empirical power. One notable difference between Yzerbyt et al. (2018) Table 3 and Table S5 in the current paper, is that the current paper tends to have higher empirical power. Some differences in empirical power may be attributed to slight differences in the parameter values (i.e., choosing $\rho_M = 0.6$ instead of 0.7 and selecting the exact ρ_Y for each condition instead of using 0.6 as a proxy). However, higher correlation among the mediators (below 0.75) and higher correlation between the outcomes tended to result in higher empirical power, suggesting that the results from Yzerbyt et al. (2018) should be slightly higher than the current study, rather than the opposite which is currently observed. This is likely due to the lack of standardized coefficients being used in the Yzerbyt et al. (2018) paper. As the b -path increases, the variance of Y_2 increases, thus as this parameter increases on its original scale, that does not mean it increases to the same extent on the standardized scale. So, the effect sizes examined in the current manuscript might be slightly larger than the effect sizes in Yzerbyt et al. (2018), leading to higher empirical power.

Within- or Between-Subject Design for Mediation

Literature Review of Correlations Among Repeated Measurements

The correlation among repeated measurements is not always reported in research papers, and thus it is difficult to know what reasonable values might be for these parameters in Monte Carlo simulations. This study explores what observed correlations among repeated measurements are common for within-subject mediation analysis. Recent publications were used to extract observed correlations, aiming for 100 correlations. These correlations were analyzed to evaluate what size of correlations are common, whether correlations tend to be negative or positive, and whether the correlations differ across mediators and outcomes.

Methods

Papers were identified to be included in the literature, then relevant statistics were extracted to calculate the correlations. From these statistics correlations were calculated. Multiple methods for calculating the correlations were used based on the reported statistics, and so at the end I needed to select among multiple calculated values if there were multiples. Each of these steps is described in the below sections.

Identifying Papers. Only papers published in 2020 that cited Montoya & Hayes (2017) were included. I used Google Scholar to identify the subset of papers that fit these criteria. Articles were sorted by relevance by Google Scholar, then papers were selected in batches of approximately 10 for coding by a trained undergraduate research assistant. I previewed papers and selected ones that seemed to have reported sufficient statistics to calculate the correlation, though some papers were excluded after coding because they did not report enough information. Fifty-six papers were identified, with each paper contributing one or more correlation.

Within- or Between-Subject Design for Mediation

Extracting Statistics. The following information was extracted from the papers to keep a record of the papers and calculate the statistics required for calculating the correlation. The undergraduate research assistant was instructed to extract statistics from studies that used within-subject mediation analysis. If analyses were divided by between-subject factors (e.g., experimental conditions, gender, etc.) separate entries were made for each between-subject group.

- Author names and year: used as a study identifier
- Study number (if there were multiple studies)
- Group description: If the statistics were reported by between-subject factors, an identifier to which group the statistics belong.
- Mediator/Outcome: Whether the information is related to the mediator or the outcome in the model.
- Mean 1: Mean for first reported within-subject condition.
- Mean 2: Mean for second reported within-subject condition.
- Mean Difference: Difference between the means for the first and second reported within-subject conditions.
- SD 1: Standard deviation for first reported within-subject conditions.
- SD 2: Standard deviation for second reported within-subject conditions.
- Pooled SD: Pooled standard deviation across the two within-subject conditions.
- SD of Difference: Standard deviation of the difference scores across within-subject conditions.
- SE of Mean Difference: Standard error of the mean of the differences scores across within-subject conditions.

Within- or Between-Subject Design for Mediation

- Cohen's d : standardized mean difference.¹
- T-statistic: For the test of the difference between the means of the within-subject conditions.
- DF: Degrees of freedom for the test of the difference between the means of the within-subject conditions.
- Correlation: If this correlation among repeated measures was directly reported in the paper.

Calculate correlations. Based on reported statistics there are a few ways that the correlation can be calculated. The first, and easiest case, is when the correlation is directly reported and no calculations are needed. The second case relies on the following equation:

$$\rho = 1 - \frac{1}{2} \left(\frac{\sigma_D}{\sigma_p} \right)^2$$

Where the correlation ρ is a function of the standard deviation of the difference and the pooled standard deviation. If the standard deviation of the difference and the pooled standard deviation are directly reported, then the above equation can be used to calculate the correlation.

Alternatively, if either or both statistics are not reported, they can typically be calculated from other statistics. To calculate the standard deviation of the difference, σ_D , we used the following methods:

¹ It is worth noting that Cohen's d can be calculated a few different ways in within-subject designs. We found there to be little regularity in the way this statistic was reported, and so did not rely on information from this statistic, if it could be avoided.

Within- or Between-Subject Design for Mediation

- 1) Calculated from Cohen's d and group means: $\sigma_D = \frac{M_1 - M_2}{d}$. This could be done using either the raw means or if the difference in means was directly reported, that could be used as the numerator.
- 2) Calculated from the t -statistic, degrees of freedom, and means: $\sigma_D = (M_1 - M_2) / \left(\frac{t}{\sqrt{df+1}} \right)$. Similar to Method 1, the raw means of the difference in means can be used for this calculation. Additionally, this could be done by using either reported t -statistic or if the difference in means and standard error of the difference is reported, their ratio could be used for the t -statistic.

To calculate the pooled standard deviation, σ_P , the two standard deviations were pooled using the following equation:

$$\sigma_P = \sqrt{\left(\frac{\sigma_1^2 + \sigma_2^2}{2} \right)}$$

The most common combination of approaches that yielded a valid correlation estimate was calculating the pooled standard deviation from the raw standard deviations and calculating the standard deviation of the difference from a reported t -statistic (Method #2). All but four correlations came from using this method. Two correlations came from using the reported standard deviation of the difference and calculating the pooled standard deviation. Two correlations came from using the reported pooled standard deviation and the standard deviation of the difference from the t -statistic.

Selecting Among Options. Some papers resulted in correlations from multiple methods, primarily those that reported Cohen's d and the t -statistic. When Cohen's d was reported, the

Within- or Between-Subject Design for Mediation

implied standard deviation of the difference was compared to the pooled standard deviation. In many cases, these were identical suggesting that the pooled standard deviation was used to calculate Cohen's d , rather than the standard deviation of the difference. When this was the case, the standard deviation of the difference from the t -statistic was used instead. Preference was given to the most direct information available. For example, if both standard deviations were reported and the pooled standard deviation was reported, the pooled standard deviation was used. However, we did examine if there were stark differences from the calculated and reported statistics. The only major difference was regarding Cohen's d . One correlation was selected for each case, and these were analyzed.

Results

Analyses were conducted to describe the pattern of correlations observed in the literature review. We looked at the proportion of correlations within the ranges described in the simulation study 0 – 0.3, 0.3 – 0.6, 0.6 – 0.9, 0.9 – 1. Additionally, some correlations were negative, though that was not explored in the simulation study, so we looked at the frequency of negative correlations, as well as the absolute value of correlations. These analyses were also broken up by correlations among mediators and correlations among outcomes.

Our final sample included 26 papers with 109 correlation estimates. Each paper had between 1 and 13 correlations. Figure S6 provides a histogram of all the correlations together.

Figure S6. Frequencies of correlations among repeated measurements

Within- or Between-Subject Design for Mediation

Overall 16 correlations were below zero (15%), and the rest were positive. Twenty correlations were between 0 and .3 (18%), 32 were between .3 and .6 (29%), 33 between .6 and .9 (30%), and 8 between .9 and 1 (7%). In general, it seems that higher correlations are more common, and positive correlations are more common than negative. Table S6 breaks down these ranges by mediators and outcomes. There were 51 correlations for mediators and 58 correlations for outcomes.

Table S6. Percentage of correlations among repeated measures broken down by numeric value and type of measure (mediator or outcome).

Correlation Range	Mediators (N = 51)	Outcomes (N = 58)
-1 to 0	14%	16%
0 to .3	14%	22%
.3 to .6	37%	22%
.6 to .9	31%	29%
.9 to 1	4%	10%

Within- or Between-Subject Design for Mediation

Discussion

Overall, both positive and negative correlations among repeated measures were observed, though they tend to be positive more frequently. Additionally, larger correlations (correlations further from zero) tend to be most common. And though not excessively common, correlations between .9 and 1 are present. These results suggest that the range of correlations selected for the current simulation study are appropriate for exploring the observed parameter space of repeated-measures correlations.

Within- or Between-Subject Design for Mediation

Using MedPower for Between-Subject Power Analysis in Empirical Example

MedPower was used to calculate power for $N = 60$ for effect sizes from Whillans et al. (2017)

(Kenny, 2017). MedPower requires standardized regression coefficients for all paths (β s).

Because the standard deviation of a dichotomous variable coded (0,1) is 0.50 the d_p effect sizes

can be converted to $\beta = d_p \times 0.50$. So, the β_a for the a -path is -.22 and the β_c for the c -path is

0.195. MedPower also requires the standardized direct effect, which can be calculated by

subtracting the standardized indirect effect from the standardized total effect: $\beta_{c'} = \beta_c -$

$\beta_a\beta_b = 0.195 - (-.22)(-.34) = 0.3152$. Below is a screenshot of the results of this analysis.

Figure S7: Output from MedPower for conducting power analysis from Whillans et al. (2017)

with $N = 60$.

Power and N Computations for Mediation

Compute Now!

Please [CLICK HERE](#) to make a small donation of \$2.50 to offset some of the costs of maintaining MedPower. I thank the more than 50 users who have already donated! Your turn to donate now.

Determine:

☒ Power given Sample Size

☐ Sample size given desired level of power

Sample Size

60

Effect Size Measure

☒ Beta ☐ partial r

Effect of X on M (path a)

-.22

Effect of M on Y (path b)

-.34

Effect of X on Y (path c')

.3152

Alpha

.05

Make sure review the assumptions of mediation analysis given in Details and Assumptions tab above.

Effect	Beta	Partial r	Power	N
c (total)	.390	.390	.887	60
a	-.220	-.220	.393	60
b	-.340	-.360	.817	60
c' (direct)	.315	.337	.757	60
ab (indirect)	.075		.322	60

Alpha for all power calculations set to .050. Effects (a, b, and c') are Betas.

Within- or Between-Subject Design for Mediation

MedPower was also used to calculate required sample size to reach a power of 0.90 given the same effect sizes as above (Kenny, 2017). Below is a screenshot of the results of this analysis.

Figure S8: Output from MedPower for conducting power analysis from Whillans et al. (2017) to reach power of 0.90.

Power and N Computations for Mediation

Compute Now!

Please [CLICK HERE](#) to make a small donation of \$2.50 to offset some of the costs of maintaining MedPower. I thank the more than 50 users who have already donated! Your turn to donate now.

Determine:

☐ Power given Sample Size

☒ Sample size given desired level of power

Desired Power

Effect Size Measure

☒ Beta ☐ partial r

Effect of X on M (path a)

Effect of M on Y (path b)

Effect of X on Y (path c')

Alpha

Make sure review the assumptions of mediation analysis given in Details and Assumptions tab above.

Effect	Beta	Partial r	Power	N
c (total)	.390	.390	.902	63
a	-.220	-.220	.902	211
b	-.340	-.360	.902	76
c' (direct)	.315	.337	.900	87
ab (indirect)	.075		.901	211

Alpha for all power calculations set to .050. Effects (a, b, and c') are Betas.

Within- or Between-Subject Design for Mediation

R script for Power Analysis in Empirical Example

```
#Function: WSmedpower
#Arguments
# nsims: number of simulations to run (default is 1000)
# boots: number of bootstrap samples (default is 0, which
doesn't run bootstrap analysis)
# MCsamples: number of monte carlo samples (default is 0, which
does run Monte Carlo analysis)
# alpha: level of test, acceptable type I error rate (default is
0.05)
# N: sample size to test
# corm: correlation among repeated measurements of mediator
# cory: correlation among repeated measurements of outcome
# aest: a-path estimates standardized by SD(M)
# best: b-path estimate standardized by SD(M) and SD(Y)
# cest: c-path estimate standardized by SD(Y)
# dest: d-path estimate standardized by SD(M) and SD(Y)

WSmedpower <- function(nsims=1000, boots=0, MCsamples=0,
alpha=0.05, N, corm, cory, aest, best, cest, dest){
  sigm1y1 <- best + dest
  sigm2y2 <- best - dest
  sigm1y2 <- corm*(best - dest)
  sigm2y1 <- corm*(best + dest)

  cormat <- matrix(c(1,          corm,      sigm1y1, sigm1y2,
                    corm,      1,          sigm2y1, sigm2y2,
                    sigm1y1, sigm2y1, 1,      cory,
                    sigm1y2, sigm2y2, cory,   1),
                  nrow = 4, ncol = 4)
  dets <- rep(NA, times = 3)
  for (i in 2:4){
    dets[i-1] <- det(cormat[1:i, 1:i])
  }

  if(all(dets > 0)){
    Results <- data.frame(matrix(data = NA, nrow = nsims, ncol =
3))
    names(Results) <- c("JointSig", "Bootstrap", "MonteCarlo")
    for(i in 1:nsims){
      M1M2Y1Y2 <- data.frame(matrix(rnorm(n = N*4), ncol =
4)%*%chol(cormat))
      names(M1M2Y1Y2) <- c("M1", "M2", "Y1", "Y2")
      M1M2Y1Y2$M2 <- M1M2Y1Y2$M2-aest
      M1M2Y1Y2$Y2 <- M1M2Y1Y2$Y2-cest
      contrast <- matrix(c( 1,  1,  0,
                           -1,  1,  0,
```

Within- or Between-Subject Design for Mediation

```
0, 0, 1,
0, 0, -1), nrow = 4, byrow = TRUE)
MDiffMSumYDiff <-
data.frame(as.matrix(M1M2Y1Y2)%*%contrast)
names(MDiffMSumYDiff) <- c("Mdiff", "Msum", "Ydiff")
MSumUncentered <- MDiffMSumYDiff$Msum
MDiffMSumYDiff$Msum <- MDiffMSumYDiff$Msum -
mean(MDiffMSumYDiff$Msum)

apath <- mean(MDiffMSumYDiff[,1])
seapath <- sd(MDiffMSumYDiff[,1])/sqrt(N)
apathsig <- 2*pt(abs(apath/seapath), df = N-1, lower.tail
= FALSE)
bmodel <- lm(Ydiff~Mdiff+Msum, data = MDiffMSumYDiff)
betas <- coef(bmodel)
bpath <- betas[2]
sebpath <- summary(bmodel)$coef[2,2]
bpathsig <- summary(bmodel)$coef[2,4]

#Joint Significance
JointSigReject <- ((apathsig < alpha)&(bpathsig < alpha))
Results[i,1] <- JointSigReject

#Bootstrap CI
if(boots > 0){
  confidence <- 1- alpha
  bootResults <- vector(length = boots)
  for (j in 1:boots){
    rows <- sample(1:N, size = N, replace = TRUE)
    bootdat <- MDiffMSumYDiff[rows, ]
    bootapath <- mean(bootdat$Mdiff)
    bootmodel <- lm(Ydiff~Mdiff+Msum, data = bootdat)
    bootbpath <- coef(bootmodel)[2]
    bootResults[j] <- bootapath*bootbpath
  }
  bootResults <- sort(bootResults)
  LowerCI = (1-confidence)/2
  UpperCI = 1 - LowerCI
  BootLCI = bootResults[ceiling(boots*LowerCI)]
  BootUCI = bootResults[ceiling(boots*UpperCI)]
  BootReject = ((BootLCI > 0)|(BootUCI < 0))
  Results[i,2] <- BootReject
}

#Monte Carlo CI
```

Within- or Between-Subject Design for Mediation

```
    if(MCsamples > 0){
      asamp <- rnorm(n = MCsamples, mean = apath, sd = seapath)
      bsamp <- rnorm(n = MCsamples, mean = bpath, sd = seopath)
      absamp <- asamp*bsamp
      absamp <- sort(absamp)
      MCLCI <- absamp[ceiling(MCsamples*LowerCI)]
      MCUCI <- absamp[ceiling(MCsamples*UpperCI)]
      MCReject <- ((MCLCI >0)|(MCUCI <0))
      Results[i,3] <- MCReject
    }

  }

  testpower <- colMeans(Results)
  return(testpower)

}

if(all(dets <= 0)){
  print("Impossible Parameter Combinations")
}

}

WSmedpower(nsims = 5000, boots = 0, MCsamples = 0, alpha = 0.05,
N = 100, corm = .22, cory = .32, aest = -.44, best = -.34, cest
= .39, dest = 0)
```


Within- or Between-Subject Design for Mediation

Supplementary References

- Aroian, L. A. (1947). The probability function of the product of two normally distributed variables. *Annals of Mathematical Statistics*, 18, 265–271. DOI: 10.1214/aoms/1177730442.
- Bradley, J. V. (1987). Robustness? *British Journal of Mathematical and Statistical Psychology*, 31, 144 – 152. DOI: 10.1111/j.2044-8317.1978.tb00581.x.
- Collins, L. M., Graham, J. J. & Flaherty, B. P. (1998) An alternative framework for defining mediation, *Multivariate Behavioral Research*, 33(2), 295 – 312, DOI: 10.1207/s15327906mbr3302_5.
- Hayes, A. F. (2018). *Introduction to mediation, moderation, and conditional process analysis*. 2nd ed. Guilford Press.
- Judd, C. M., & Kenny, D. A. (2014). Power anomalies in testing mediation. *Psychological Science*, 25(2), 334-339. DOI: 10.1177/0956797613502676.
- Kenny, D. A., Kashy, D. A., & Bolger, N. (1998). Data analysis in social psychology: Recent and recurring issues. In *Handbook of social psychology* (4 ed., Vol. 1). DOI: 10.1002/9780470561119.socpsy001004.
- MacKinnon, D. P. (2008). *Introduction to statistical mediation analysis*. Routledge.
- MacKinnon, D. P., & Dwyer, J. H. (1993). Estimating mediated effects in prevention studies. *Evaluation Review*, 17(2), 144-158. DOI: 10.1177/0193841X9301700202.
- MacKinnon, D. P., Krull, J. L., & Lockwood, C. M. (2000). Equivalence of the mediation, confounding and suppression effect. *Prevention Science*, 1, 173-181. DOI: 10.1023/A:1026595011371.

Within- or Between-Subject Design for Mediation

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(1), 99-128. DOI: 10.1207/s15327906mbr3901_4.

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, 83 – 104. DOI: 10.1037/1082-989x.7.1.83.

O'Rourke, H. P. & MacKinnon, D. P. (2018). Reasons for testing mediation in the absence of an intervention effect: A research imperative in prevention and intervention research. *Journal of Studies on Alcohol and Drugs*, 79(2), 171–181, DOI:10.15288/jsad.2018.79.171.

Preacher, K. J., & Selig, J. P. (2012) Advantages of Monte Carlo confidence intervals for indirect effects. *Communication Methods and Measures*, 6(2), 77-98. DOI: 10.1080/19312458.2012.679848.

Serlin, R. C. (2000). Testing for robustness in Monte Carlo studies. *Psychological Methods*, 5(2), 230 – 240. DOI: 10.1037/1082-989x.5.2.230.

Sobel, M. E. (1982). Asymptotic confidence intervals for indirect effects in structural equation models. In S. Leinhardt (Ed.), *Sociological Methodology*, (pp. 290 – 293). Washington, DC: American Sociological Association. DOI: 10.2307/270723.

Sobel, M. E. (1986). Some new results on indirect effects and their standard errors in covariance structure models. In N. Tuma (Ed.), *Sociological Methodology*, (pp. 159 – 86). Washington, DC: American Sociological Association.