Published in final edited form as:
J Exp Soc Psychol. 2016 September ; 66: 29–38. doi:10.1016/j.jesp.2015.09.012.

Design approaches to experimental mediation☆

Angela G. Pirlott[a,*] and **David P. MacKinnon**[b]

[a]Department of Psychology, Saint Xavier University, United States

[b]Department of Psychology, Arizona State University, United States

Abstract

Identifying causal mechanisms has become a cornerstone of experimental social psychology, and editors in top social psychology journals champion the use of mediation methods, particularly innovative ones when possible (e.g. Halberstadt, 2010, Smith, 2012). Commonly, studies in experimental social psychology randomly assign participants to levels of the independent variable and measure the mediating and dependent variables, and the mediator is assumed to causally affect the dependent variable. However, participants are not randomly assigned to levels of the mediating variable(s), i.e., the relationship between the mediating and dependent variables is correlational. Although researchers likely know that correlational studies pose a risk of confounding, this problem seems forgotten when thinking about experimental designs randomly assigning participants to levels of the independent variable and measuring the mediator (i.e., "measurement-of-mediation" designs). Experimentally manipulating the mediator provides an approach to solving these problems, yet these methods contain their own set of challenges (e.g., Bullock, Green, & Ha, 2010). We describe types of experimental manipulations targeting the mediator (manipulations demonstrating a causal effect of the mediator on the dependent variable and manipulations targeting the strength of the causal effect of the mediator) and types of experimental designs (double randomization, concurrent double randomization, and parallel), provide published examples of the designs, and discuss the strengths and challenges of each design. Therefore, the goals of this paper include providing a practical guide to manipulation-of-mediator designs in light of their challenges and encouraging researchers to use more rigorous approaches to mediation because manipulation-of-mediator designs strengthen the ability to infer causality of the mediating variable on the dependent variable.

Keywords

Mediation; Experimental mediation; Causal inference

Journal editors almost require mediation for publication, as noted by past Associate Editor of the *Journal of Personality and Social Psychology* (*JPSP*), Robert Cialdini (2009): "who could argue the importance of understanding what mediates the effects of interest to

☆We presented portions of this paper at the Society for Personality and Social Psychology Conference in January 2010 and the Society for Experimental Social Psychology in September 2015. This research was supported in part by National Institute on Drug Abuse DA09757.

*Corresponding author at: Department of Psychology, Saint Xavier University, 3700 W 103rd St, Chicago, IL 60655, United States. pirlott@sxu.edu. .

psychologists? Mediation is about what research psychologists care about—locating causality—and sophisticated psychometric techniques now allow mediational accounts of our major findings…" (p. 5). Associate Editor of *Journal of Experimental Social Psychology*, Jamin Halberstadt (2010) stated in a *Society for Personality and Social Psychology* electronic mailing list email that he outright rejects manuscripts that fail to examine mediation: "I will desk reject all papers that are unlikely to survive the review process, or do not on their face satisfy the standards or goals of the Journal. This includes, in my opinion, […] studies with no insight into psychological mechanism." Furthermore, in his editorial as incoming editor of *JPSP*, Eliot Smith (2012) identified mediation as a critical component of social psychology: An "explanation of observed effects in terms of underlying processes is almost a signature of articles that *JPSP* has historically published. Only rare articles demonstrate an effect without making at least some progress toward identifying the contributing processes. The most common approach to identifying those processes is mediation analysis. Thus recent developments in both the theory and the methods of mediation analysis are particularly significant for this journal" (p. 1–2).

Indeed, identifying causal mechanisms lies as a cornerstone of experimental social psychology. Many articles in top social psychology journals include mediation in at least one study: 59% of articles in *JPSP* and 65% in *Personality and Social Psychology Bulletin* (*PSPB*) from 2005 to 2009 (Rucker, Preacher, Tormala, & Petty, 2011); 41% of studies in *PSPB* within a six month period in 2007 (Kashy, Donnellan, Ackerman, & Russell, 2009); and 16% of studies in *Psychological Science* from 2011 to 2012 (Hayes & Scharkow, 2013).

In typical experimental designs examining mediation, researchers randomly assign participants to levels of the independent variable (X) and measure the mediating (M) and dependent (Y) variables—known as *measurement-of-mediation* designs (Spencer, Zanna, & Fong, 2005). (For clarity, we use M to reflect the *measured* mediator and M* to reflect the *manipulated* mediator). Researchers then perform statistical analyses to provide estimates for the models summarized in Fig. 1. Three regression equations comprise the single mediator model (shown in Fig. 1):

$$\mathrm{M}=i_1+a\mathrm{X}+e_1 \quad (1)$$

$$\mathrm{Y}=i_2+c\mathrm{X}+e_2 \quad (2)$$

$$\mathrm{Y}=i_3+c^{'}\mathrm{X}+b\mathrm{M}+hXM+e_3 \quad (3)$$

where X is the independent variable, M is the measured mediator, and Y is the dependent variable; i_1, i_2, and i_3 are the intercepts in each equation; and e_1, e_2, and e_3 are the residual errors reflecting misprediction and unobserved omitted influences. Eq. (1) regresses M on X and *a* represents the statistical effect of X on M. Eq. (2) regresses Y on X and *c* represents the statistical effect of X on Y, or the total effect. Eq. (3) regresses Y on X and M simultaneously; c' represents the statistical effect of X on Y, adjusting for M, or the direct effect of X (i.e., the effect of X not mediated by M); b represents the statistical effect of M on Y adjusting for X; and h represents the statistical interaction effect of X and M but this is often assumed to be zero. The quantity, ab, is the estimator of the mediated effect (also called the indirect effect), if a series of assumptions are met.

Inferring that M mediates the relationship between X and Y rests upon the following assumptions: (1) no confounding of the X to Y relation; (2) no confounding of the X to M relation; (3) no confounding of the M to Y relation; and (4) no effects of X that confound the M to Y relation (VanderWeele & Vansteelandt, 2009). Furthermore, the temporal order of the variables X to M to Y is assumed correct and X, M, and Y are assumed reliable and valid measures of the intended constructs (MacKinnon, 2008).

Baron and Kenny's (1986) causal steps approach to mediation article is the most widely cited article in *JPSP* (Quinones-Vidal, Lopez-Garcia, Penaranda-Ortega, & Tortosa-Gil, 2004) at 20,326 times, according to Web of Science in June 2013. This suggests a reliance upon designs in which researchers measure the mediator and perform statistical mediation analyses (called *measurement-of-mediation designs* by Spencer et al., 2005)—particularly the causal steps approach—to provide evidence of a mediation relationship, over designs that randomly assign participants to levels of the proposed mediator, which we term *manipulation-of-mediator* designs (e.g., Smith, 2012). Measurement-of-mediation designs contain serious limitations (see for instance, Jacoby & Sassenberg, 2011; Spencer et al., 2005; and Stone-Romero & Rosopa, 2008); we focus our discussion on their limitations to causal inference of the M to Y relationship.

The mediation model is a theoretical model implying causality: The independent variable causes a change in the mediator that causes a change in the dependent variable. According to Shadish, Cook, and Campbell (2002), three requirements exist to infer that one variable causes another. First, temporal precedence, such that the causal variable precedes the dependent variable in time. Second, covariation between the hypothesized causal and dependent variables, such that the independent and dependent variables vary together. Third, no plausible alternative explanations account for the relation between the hypothesized causal and outcome variables. Other approaches to causal inference exist that share some requirements of those outlined by Shadish et al. (2002), e.g., considerations for causal relations, Hill (1971), the potential outcomes model (Holland, 1988; Rubin, 1974, 1977), and related causal inference models (Pearl, 2000; Robins & Greenland, 1992) but the work by Shadish et al. (2002) remains widely used in social psychology. Therefore, we rely upon the classic Campbellian approach to causal inference (West & Thoemmes, 2010), which focuses upon a pattern of research results logically consistent with a research hypothesis.

Only well-designed experiments satisfy all three criteria of temporal precedence, covariation, and lack of alternative explanations: Researchers randomly assign participants to levels of the independent variable and significant differences between conditions on the dependent variable suggest the independent variable caused change in the dependent variable. The manipulation of the independent variable occurs before measurement of the dependent variable, which satisfies the temporal precedence criterion. An effect of the independent variable on the dependent variable satisfies the covariation criterion. Finally, random assignment to conditions ensures that no pre-existing individual differences between conditions account for the differences between conditions. Assuming no confounds exist between conditions, no alternative explanations should account for the pattern of findings. Thus, in a study examining the effects of X on Y, an experiment satisfies the causal inference criteria, therefore enabling the causal inference of X on Y.

A mediation model, however, is a more complicated model containing three causal paths—the effects of X on M, X on Y, and M on Y. In common social psychology experiments including mediation, researchers randomly assign participants to levels of the independent variable, measure the mediator and dependent variables, and perform statistical mediation analyses to demonstrate the ability of the mediator to *statistically* account for relationship between X and Y as shown in Fig. 1. However, providing statistical evidence of a mediation relationship fails to provide *causal evidence* of the mediation relationship. Random assignment of participants to levels of the independent variable enables causal interpretation of the X to M and X to Y relationships as it satisfies all three criteria for causal inference—temporal precedence, covariation, and the elimination of alternative explanations. Although measuring M and Y satisfies the criterion for covariation between M and Y, it does not demonstrate temporal precedence of M to Y or the elimination of alternative explanations for the relationship between M and Y. This design cannot differentiate whether M causes Y, Y causes M, or some unmeasured confounding variable causes M and Y. Due to lack of random assignment to levels of the mediator, claims regarding the causal relation of M to Y are unjustified. Participants self-select to levels of the mediator; their values of the mediator are not randomly assigned. Other variables could confound the relationship between M and Y if not included in the statistical analysis. These omitted variables provide alternative explanations for the relation between M and Y instead of the hypothesized mediation process. The assumption that no omitted variables influence the mediation relation is called the no omitted variable assumption or no confounding assumption.

Some other variable could correlate and vary systematically with the mediator—a confounding variable. Using measurement-of-mediation designs, statistical analyses cannot differentiate whether the measured mediator or an unmeasured confound mediates the relationship between independent and dependent variables (Sigall & Mills, 1998). If the unmeasured confounding variable mediates the relationship between the independent and dependent variables, statistical mediation analyses could lead to the false conclusion that the measured mediator mediates the relationship. Accordingly, one cost to omitting confounding variables is that it is unclear whether the measured mediator is actually the true mediator. A second cost is that omitting confounding variables could bias mediation analyses by suppressing the relationship between M and Y and causing a Type II error, or bias analyses by spuriously increasing a mediated effect and leading to a Type I error. Additionally, as more variables are omitted (i.e., the more potentially confounded or omitted alternative mediators), the greater the likelihood of bias (Bullock et al., 2010).

Given the importance of detecting alternative explanations and confounds to experimental social psychology, it is surprising that the level of scrutiny typically applied to experimental manipulations has not yet been applied to mediation. Bullock et al. (2010) noted this warning was raised over 30 years ago by Judd and Kenny (1981) and the problem of confounding variables in mediation has been repeatedly made in the psychological literature (Holland, 1988; James, 1980; MacKinnon, Lockwood, Hoffman, West, & Sheets, 2002, p. 99–100; MacKinnon, 2008, Chapters 3 and 13; MacKinnon & Pirlott, 2015; McDonald, 1997), but these warnings have been largely ignored.

Given these shortcomings, researchers and journal editors championed the use of manipulation-of-mediator designs when appropriate (e.g. Smith, 1982, 2012; Spencer et al., 2005). For example in his editorial as incoming editor of *JPSP*, Eliot Smith noted, "Where appropriate, authors should consider adopting the experimental [mediation] approach, which is still underrepresented in the literature compared with measurement-of-mediation designs" (2012, p. 2). However, manipulation-of-mediator designs also contain challenges. Bullock et al. (2010) contended that significant challenges arise in manipulation-of-mediator designs that undermine its efficacy as a design.

Given the attention surrounding manipulation-of-mediator designs, the goals of this paper include describing new advances in manipulation-of-mediator designs while nonetheless reminding readers of the difficulties of these designs raised by Bullock et al. (2010) and others (MacKinnon, 2008; Chapter 14). In this paper, we describe types of experimental manipulations targeting the mediator (e.g., manipulations demonstrating a causal effect of the mediator on the dependent variable and manipulations targeting the strength of the mediator) and types of experimental designs (double randomization, concurrent double randomization, and parallel) while discussing their strengths and weaknesses and providing examples of research studies implementing those designs.

1. Types of manipulation-of-mediator manipulations

In this section, we detail two types of manipulations—one that experimentally manipulates the mediator to demonstrate an effect on Y and one that experimentally manipulates the *effect* of the mediator to either strengthen or weaken the effect of X on Y through the mediator.

1.1. Experimental manipulations demonstrating a causal effect of the mediator

Experimental designs that seek to demonstrate a causal effect of the mediator randomly assign participants to one condition of M*—either in a present versus absent or high versus low operationalization of the mediator—and measure the effects of M* on Y to demonstrate a causal effect of the mediator on the dependent variable. Mean differences in Y should correspond with the conditions of M*. Manipulating the value of the mediator in a dichotomized fashion (e.g., high versus low) creates systematic variance in Y corresponding with the levels of M*.

Note that Imai, Tingley, and Yamamoto (2013) argue that mediators typically cannot be manipulated to levels reflecting a true presence or absence of the mediator, and so differentiate between manipulations of mediators in an *absolute* sense (i.e., present versus absent) versus *encouragement* and *discouragement* manipulations (see also Holland, 1988).

In the *encouragement* condition, the value of the mediator is *encouraged* or increased, and in the *discouragement* condition, the value of the mediator is *discouraged* or decreased.

Consider including a manipulation check that measures the effect of the M* manipulation on a synonymous measured construct, i.e., measure the mediator. Mean differences in the measured mediator should reflect differences occurring as a function of the manipulated mediator M*. This demonstrates that the manipulation of M* affected the mediator in the theorized manner. This also provides evidence of the construct validity of the manipulation. In addition, measuring other mediators which the manipulation M* potentially could have also affected (i.e., confounding mediators) provides evidence of whether other alternative mediators actually drove the effects of M* on Y, rather than the hypothesized construct. These effects could be examined in a regression analysis which regresses Y on M* and the other measured potential mediators to determine which variables account for the effects on Y, whether M* or one of the other measured M variables, or a combination to determine whether M* affects Y.

To use an example of an experimental manipulation of the mediator from the literature, Li et al. (2012) conducted a manipulation-of-mediator design in which their manipulation sought to demonstrate the causal effects of M* on Y. Their theoretical model suggested that belief in a soul is a mediator of the relationship between religious affiliation (X: Catholic vs. Protestant) and internal attributions (Y), such that Protestants' greater belief in a soul drives their greater propensity to make internal attributions relative to Catholics. They manipulated Protestants' belief in a soul (M*) by randomly assigning them to either write an essay suggesting that souls do exist (encouragement condition) or do not exist (discouragement condition). The encouragement condition reflects an operationalization of the mediator at a high level—high belief in a soul, and the discouragement condition reflects an operationalization of the mediator at a low level—low belief in a soul. Li and colleagues argued that the manipulation strengthened or weakened the mediator, rather than fully changing the mediator. Thus, Protestants were not actually randomly assigned to *stop* or *start* believing in a soul, but instead *encouraged* to either *increase* or *decrease* their belief in a soul by writing an essay generating reasons why souls do or do not exist. As predicted by their theoretical model, mean differences in Y corresponded with conditions of M*: Protestants in the high belief in a soul condition made significantly greater internal attributions relative to those in the low or no belief in a soul condition. In other words, the means in Y significantly differed as a function of condition of M*. Although Li et al. did not report measuring the mediator (belief in a soul) as a manipulation check to demonstrate that the manipulation of belief in a soul (M*) actually affected belief in a soul (M), we expect that the value of the measured mediator M would be higher in the encouragement condition relative to the discouragement condition, which brought about the increase in internal attributions Y.

Note that the design used by Li et al. manipulated M* within one level of X to show the effect of M* on Y within Protestants. An alternative to this design is to examine the effect of M* on Y within both levels of X, or to concurrently manipulate both X and M* (see concurrent double randomization designs below). Whether M* demonstrates an effect on Y only within a certain level of X or across both levels of X might depend upon the substantive research question being answered.

1.2. Experimental manipulations targeting the effect of the mediator

Manipulation-of-mediator experimental designs that manipulate the effect of the mediator (Mark, 1986, 1990; see also Jacoby & Sassenberg, 2011) seek to either increase or decrease the effect of the mediator, in conjunction with an experimental manipulation of X (see these concurrent manipulation designs discussed further below). By influencing the effect of the mediator, these manipulations also influence the effect of X on Y. This manipulation, which targets the effect of the mediator, occurs relative to a control condition that allows the mediator to vary freely. Therefore, experimental manipulations targeting the effect of the mediator manipulate the *systematic variance* in the mediator typically caused by the manipulation of X, by either increasing or decreasing the systematic variance in the mediator associated with X. This results in, minimally, a 2 by 2 design in which both X and M* are manipulated concurrently. The two levels of M* are either increasing or decreasing the effect of the mediator (i.e., increasing or decreasing the systematic variance in the mediator caused by X) versus allowing it to vary freely.

A *blockage manipulation* is one type of experimental design manipulating the effect of the mediator that seeks to block or "neutralize" the mediator to prevent the mediator from operating, or more probably, reduce the effect of the mediator. The blockage manipulation blocks the mediator by eliminating the systematic variance in the mediator created by X and shows a decreased effect of the mediator. To block or decrease the effect of the mediator by eliminating systematic variance in the mediator due to change in X, researchers create a setting in which the mediator is controlled or blocked across levels of X, i.e., unable to vary freely as a function of X. No mean differences should exist in Y as a function of X (due to the blockage of the mediator). Contrast this with the control conditions in which the mediator is allow to vary freely and the expected mean differences in Y exist as a function of X. In the blockage conditions, because the mediator is manipulated and not measured, this will result in a decreased effect of X on Y, i.e., smaller mean differences in Y as a function of X. A true blockage manipulation occurs when no mean differences in Y exist as a function of X and suggests that the entire effect of X on Y is through that particular mediator. However, a more likely scenario is that mean differences in Y still exist as a function of X, although to a lesser extent, because X still affects Y via additional mediators and/or X affects Y to an extent independent of the targeted mediator.

As an example of a blockage manipulation, the cognitive dissonance model suggests that counterattitudinal behaviors (X) increase arousal dissonance (M) which in turn causes attitude change (Y). Cooper, Zanna, and Taves (1978) targeted the variability of the mediator —arousal—by administering a stimulant, depressant, or placebo to participants. The placebo condition allowed arousal to vary freely whereas the stimulant and depressant conditions blocked arousal from varying freely. Researchers then either asked or told participants to write a counterattitudinal essay pardoning Richard Nixon (X: dissonance manipulation—choice), and measured attitudes toward pardoning Richard Nixon (Y). By administering a depressant meant to suppress arousal, Cooper et al. (1978) blocked the mediator of arousal from varying freely, and which should elicit a smaller effect (or no effect) of choice condition on attitudes in the blockage conditions relative to placebo conditions. Results were consistent with expectations of a blockage manipulation. Replicating previous findings,

Author Manuscript

Author Manuscript

Author Manuscript

participants in the placebo M* condition displayed more favorable attitudes toward the Nixon pardon in the high choice condition relative to the low choice condition (mean difference: 14.7–8.3 = 6.4, see Table 1). In the depressant M* condition, the effects of X on Y were blocked as arousal was blocked by the depressant M*: the attitudes toward the Nixon pardon did not differ in the high versus low choice conditions (mean difference: 8.6–8.0 = 0.6). Although Cooper et al. (1978) did not report effect sizes comparing the effects of X on Y in the placebo and depressant conditions, we expect those effects to differ, thus showing the manipulation of the mediator targeting the effect of the mediator by manipulating its systematic variance affected the effect of X on Y. Given the suppression of variability in arousal in the stimulant condition, we would except similar effects, such that no effects of X on Y occurred in the stimulant condition, although researchers nonetheless found effects of X on Y in the stimulant condition (mean difference: 20.2–13.9 = 6.3) in a similar magnitude to the placebo condition (mean difference = 6.4).

An *enhancement manipulation* manipulates the mediator and seeks to increase the effects of the mediator to show its enhanced effects when manipulated, i.e., larger mean differences in Y corresponding to levels of X relative to the control conditions in which the mediator varies freely. The enhancement manipulation increases the effect of the mediator by increasing the systematic variance in the mediator created by X. To increase the effect of the mediator by increasing systematic variance in the mediator due to change in X, researchers create a setting in which the mediator differs maximally across corresponding levels of X. This essentially suggests that researchers create two conditions of X which vary systematically with the two corresponding levels of the mediator, and contrast the effects of those two conditions with the effects of the two conditions of X that allow the mediator to vary freely, thus yielding a 2 (X: high versus low) by 2 (M*: co-occurring at high and low levels with X versus varying freely) design. This uses the aforementioned manipulation-of-mediator manipulation which targets the value of the mediator, but in this case the mediator varies and co-occurs with its corresponding level of X, to enhance systematic variance stemming concurrently from X and M*.

To use Cooper et al.'s (1978) cognitive dissonance by arousal design as example of an enhancement manipulation, researchers randomly assigned participants to dissonance condition (X: high or low choice) and arousal manipulation (M*: stimulant, depressant, placebo) and measured their attitudes toward the Nixon pardon. In this particular example, the enhancement manipulation would contrast the high versus low levels of X co-occurring with high versus low levels of the mediator (i.e., stimulant versus depressant) against the high versus low levels of X with the mediator varying freely (i.e., placebo condition). The effects of X on Y should be larger when associated with the corresponding levels of the mediator relative to the effects of X on Y when allowing the mediator to vary freely. Therefore, in this example, we would expect the mean differences in Y to be greater in the high choice/high arousal condition versus low choice/low arousal condition, relative to the high choice/vary freely mediator condition versus low choice/vary freely mediator condition. As shown in Table 1, the effect of X on Y when the mediator was allow to vary freely yielded a mean difference of 6.4 (14.7–8.3). The effect of X on Y when the mediator varied maximally with levels of X was larger—a mean difference of 12.2 (20.2–8.0).

Similar to a manipulation check as described above in experimental manipulations seeking to demonstrate an effect of M* on Y, again consider including a manipulation check that measures the mediator. This should demonstrate mean differences (or a lack of mean differences) in M corresponding to conditions of X as a function of whether the mediator was blocked, enhanced, or allowed to vary freely. This demonstrates that the manipulation of the mediator affected the mediator in the theorized manner and again, provides evidence of the validity of the manipulation. Also as described above, measuring other alternative mediators which could have been affected by the manipulation of the mediator provides evidence of whether other alternative mediators actually caused the effects, and including the other measured mediators in a regression analyses could demonstrate which measured mediator drove the effects to rule out alternative explanations.

Measuring the mediator as a manipulation check should reveal a pattern of means specific to the type of manipulation (blockage, enhancement, or vary freely). Allowing the mediator to vary freely should reveal a pattern of mean differences in M corresponding with the appropriate levels of X. An enhancement manipulation which manipulates X with the corresponding levels of M* should reveal a pattern of mean differences in M corresponding to the manipulation—high X/high M* versus low X/low M*. If comparing the enhancement condition to the vary freely condition, the pattern of mean differences in M in the enhancement condition should reveal greater effects in M relative to the vary freely condition. A blockage manipulation which blocks the mediator should reveal a pattern of no mean differences in M corresponding with levels of X, or if comparing the blockage manipulation to the vary freely manipulation, the blockage manipulation should show smaller mean differences in M as a function of X relative to the mean differences in M as a function of X in the vary freely condition.

2 Types of manipulation-of-mediator designs

The next section focuses on different types of manipulation-of-mediator designs: double randomization, concurrent double randomization, and parallel designs.

2.1. Double randomization designs

Double randomization designs (also called *process manipulations* by Mark, 1990 and *experimental-causal-chain* designs by Spencer et al., 2005; Stone-Romero & Rosopa, 2008) randomly assign participants to X and measure M and Y in Experiment 1 to allow clear interpretation of the X to M and X to Y paths as estimates of causal influence. Then in Experiment 2, randomly assign participants to levels of M* defined by how X changed M in Experiment 1. If M* significantly affects Y in Experiment 2 (i.e., if there are significant differences in Y as a function of M*—a significant main effect of M* on Y), evidence exists in support for mediation and the causal effect of M* on Y. As noted above, a manipulation check measuring M also demonstrates the causal effects of X and M* on M.

Using this approach, Word, Zanna, and Cooper (1974) tested mediation across two experimental studies. They examined self-fulfilling prophecy in interracial interviews, such that interracial interaction (X) predicted quality of interview (M), which predicted interview performance (Y). To experimentally test mediation in Study 1, they randomly assigned

White participants to interview either a Black or a White confederate interviewee. Race of interviewee (X) significantly affected interview quality (M), such that Black interviewees received poorer interviews—less immediacy, more speech errors, and shorter interviews; and race of interviewee affected interview performance (Y) such that Black interviewees performed worse relative to White interviewees. Study 2 then experimentally tested the M* to Y relationship, such that White participant interviewees were randomly assigned to one of two interview conditions—they either received a poor quality interview like Black applicants—less immediacy, more speech errors, and shorter interviews, or a high quality interview like White applicants. Participants receiving poor interviews (M*—like Black applicants) performed worse (Y).

As described above when discussing encouragement/discouragement mediator manipulations, work by Li et al. (2012) also provides an example of the double randomization design. Their theory predicted belief in a soul (M) mediated the effect of religious group differences (X: Protestant versus Catholic) in internal attributions (Y), such that Protestants' greater belief in a soul drove their elevated use of internal attributions relative to Catholics. To test this, in one study, they demonstrated belief in a soul (M) statistically mediates the relationship between religious group (X) and internal attributions (Y). In a second study they randomly assigned Protestants' levels of M* and measured internal attributions. To manipulate belief in a soul, they used an encouragement/discouragement manipulation-of-mediator design, in which they randomly assigned Protestants to write an essay either suggesting that souls do or do not exist; writing that souls exist increased internal attributions relative to the no soul condition.

The primary strength of double randomization designs is that they provide important experimental evidence about the mediation relation. Random assignment of participants to levels of the mediator satisfies the temporal precedence criterion, demonstrates covariation of M and Y, and beyond the abilities of measurement-of-mediation designs, reduces the plausibility of alternative interpretations of the M to Y relationship. Furthermore, demonstrating that M* causes Y in the second study provides a conceptual replication of the pattern of findings found in the first study. Replicating the pattern of findings using a conceptual replication (i.e., here using different operationalizations of the mediating variable) strengthens the ability to infer that the construct of interest (here, the mediator) is responsible for the given pattern of findings (Brewer & Crano, 2014; see also Fabrigar & Wegener, in press and Hüffmeier, Mazei, & Schultze, in press, for a detailed discussion on the benefits of conceptual replications).

2.2. Concurrent double randomization designs

Concurrent double randomization designs (also called testing-a-process-hypothesis-by-an-interaction strategy by Jacoby & Sassenberg, 2011 and *moderation-of-process designs* by Spencer et al., 2005; see also Sigall & Mills, 1998) experimentally manipulate both the mediator and independent variable simultaneously in a two-factor experimental design. This enables an examination of the causal effect of X on Y and M* on Y. Including a manipulation check measuring M would also provide evidence of the causal effects of X on M and M* on M.

As described above as an example of a blockage and enhancement manipulation-of-mediator design, we use Cooper et al. (1978)'s cognitive dissonance study to demonstrate a concurrent double randomization design. Cooper et al. (1978) manipulated the hypothesized mediator—arousal (M*)—by randomly assigning participants to receive a placebo, stimulant, or depressant meant to affect arousal thereby either allowing it to vary naturally, enhancing it, or blocking it while simultaneously manipulating X—participants were either asked or told to write a counterattitudinal essay (X: dissonance manipulation—high or low choice) and then reported attitudes toward Richard Nixon (Y). Given that the manipulation of arousal included a manipulation to either set the mediator to a specified value (e.g., high arousal via stimulant or low arousal via depressant) or to allow it to vary freely, for the sake of clarity in this example we focus interpretations on the high versus low arousal manipulation of M*.

Looking at the interaction between M* (high versus low arousal; omitting the placebo condition) by X (high versus low choice) in Table 1 reveals the causal effects of M* and X on Y. There appears to be a main effect of X on Y: averaging the effects of X on Y over conditions of M* yields a mean difference of 3.45 [X_{high}: (8.6 + 20.2)/2; X_{low}: (8.0 + 13.9)/2]. There also appears to be a main effect of M* on Y, across levels of X, with a mean difference of 16.75 [M^*_{high}: (8.6 + 8.0)/2; M^*_{low}: (20.2 + 13.9)/2]. This reveals a causal effect of both M* and X on Y. There also appears to be an interaction between X and M* suggesting the strongest effects occur in the conditions in which levels of X occur with the corresponding levels of M*. The pattern of findings might also demonstrate that the strongest effects (i.e., mean differences) occur in the conditions in which levels of X correspond with the appropriate levels of M*, for example, X_{high}/M^*_{high} versus X_{low}/M^*_{low}.

A limitation to concurrent double randomization is the inability to demonstrate that X affects M, although measuring the mediator would demonstrate if both X and M* affect M. Nonetheless, concurrent double randomization designs provide important experimental evidence about the mediation relation. Through random assignment of observations to levels of M*, these designs establish temporal precedence of M* to Y and reduce the plausibility of alternative explanations of the M* to Y relationship. They thus also go beyond measurement-of-mediation designs in allowing for the interpretation of a potential covariation between M* and Y as causal.

2.3. Parallel designs

Parallel designs (Imai et al., 2013) essentially combine concurrent double randomization manipulation-of-mediator designs with measurement-of-mediation designs. Researchers randomly assign participants into one of two studies assessing the same mediation model—either a measurement-of-mediation design in which participants are randomly assigned to levels of X and M and Y are measured, or a concurrent double randomization manipulation-of-mediator design in which participants are randomly assigned to a level of X and a level of M*, and Y is measured. This is conceptually the equivalent of randomly assigning participants to a level of X and a level of M* in which the mediator is allowed to vary freely (i.e., no manipulation of the mediator), or fixed to a high versus low value of the mediator, and identical to the Cooper et al. design described above.

As described above as an example of a blockage and enhancement manipulation-of-mediator design, we use the Cooper et al. (1978)'s cognitive dissonance study to demonstrate a concurrent double randomization design as a modification of the parallel design. Cooper et al. (1978) manipulated the hypothesized mediator—arousal (M*)—by randomly assigning participants to receive a placebo, stimulant, or depressant meant to affect arousal thereby either allowing it to vary naturally, enhancing it, or blocking it while simultaneously manipulating X—participants were either asked or told to write a counterattitudinal essay (X: dissonance manipulation—high or low choice) and then reported attitudes toward Richard Nixon (Y). Given the complexity of this design, the results can be examined in several different ways:

(1) Looking at the effects of X on Y within the M* vary freely condition (here, M* = placebo condition), consistent with previous cognitive dissonance studies, participants in the placebo condition demonstrated the predicted pattern of findings: Participants in the high choice condition produced more supportive attitudes relative to the low choice condition.

(2) As an example of a blockage design: Looking at the effects of X on Y within the M* = depressant condition relative to the effects of X on Y within the M* = placebo condition demonstrates that when constraining the variance of the mediator (depressant condition), there are no effects of X on Y relative to the effects of X on Y in the placebo condition. The same effects should also occur when examining the effects of X on Y within the M* = stimulant condition relative to the effects of X on Y in the M* = vary freely condition, although this did not occur in this example.

(3) As an example of an enhancement design: Looking at the effects of X on Y when X co-occurs at the corresponding levels of M* (i.e., high choice/high arousal and low choice low arousal) relative to M varying naturally reveals a larger effect of X on Y in the enhancement condition relative to the M varying naturally condition.

(4) Looking at the interaction between M* (high versus low arousal; omitting the placebo condition) by X (high versus low choice) reveals the causal effects of M* and X on Y: there appears to be a main effect of X on Y and a main effect of M* on Y.

A strength of this design is that by randomly assigning participants to the measurement-of-mediation and concurrent double randomization manipulation-of-mediator designs, any differences between the two studies cannot be attributed to individual differences. In other words, when measurement-of-mediation and concurrent double randomization manipulation-of-mediator designs are used to provide evidence of the mediated effect and the causal relationship between M and Y, the differences between the two studies could be attributed to individual differences because the samples could systematically differ. Random assignment to design eliminates this confound. Thus, if the results of both experiments provide convergent evidence that the mediation relation occurred, there is decreased likelihood that the results were due to an unmeasured confounding variable, and increased evidence of arousal as a mediator (in this example).

Furthermore, replicating the mediation model in two different types of designs—measurement-of-mediation and manipulation-of-mediation—provides evidence of a conceptual replication, which further strengthens the ability to infer that the mediator was the variable responsible for the process by which X affected Y (Brewer & Crano, 2014; Fabrigar & Wegener, in press; Hüffmeier et al., in press).

A drawback to this design is that it requires a large sample size—at least enough participants to be randomly assigned to six different conditions, and this potentially undermines its effectiveness in studies with limited funding or participants. However, the primary strength of the parallel design lies in that, if the findings of the measurement-of-mediation and the concurrent double randomization designs manipulation-of-mediator converge, evidence exists that the mediated relation found in the measurement-of-mediation design is not due to unmeasured variables confounded with the mediator. This provides strong causal evidence for the mediation relation.

3. Challenges underlying manipulation-of-mediator designs

Manipulation-of-mediator designs impose their own challenges and limitations. Below we discuss challenges underlying manipulation-of-mediator designs, including distinguishing between mediators and moderators, alternative explanations for the relationship between M and Y, the precision and ease of manipulating mediators, construct validity of the manipulated mediator, demonstrating that X causes M, assessing the indirect effect, knowing the causal effect at the individual-versus group-level, and heterogeneity of causality.

3.1. Manipulated mediators inherently become moderators

Manipulating mediators to demonstrate how the effects of X on Y differ as a function of M* inherently causes a mediator to also become a moderator. Therefore, in manipulation-of-mediator designs, it can be difficult to distinguish between mediators and moderators (Bullock et al., 2010; Imai et al., 2013; Mark, 1990; Spencer et al., 2005).

Mediators and moderators differ importantly. To define a moderator, from Baron and Kenny (1986, p. 1174): "a moderator is a qualitative (e.g., sex, race, class) or quantitative (e.g., level of reward) variable that affects the direction and/or strength of the relation between an independent or predictor variable and a dependent or criterion variable. Specifically within a correlational analysis framework, a moderator is a third variable that affects the zero-order correlation between two other variables. […] In the more familiar analysis of variance (ANOVA) terms, a basic moderator effect can be represented as an interaction between a focal independent variable and a factor that specifies the appropriate conditions for its operation." To define a mediator, from Baron and Kenny (1986, p. 1176) as the generative mechanism by which an independent variable influences the dependent variable and in this way "it accounts for the relation between the predictor and criterion. […] Whereas moderator variables specify when certain effects will hold, mediators speak to how or why such effects occur."

Experimental manipulations-of-the-mediator designs that target the effect of the mediator (i.e., blockage and enhancement designs) manipulate the effect of the mediator which in turn effects the strength of the relationship between X and Y as a function of the effect on the mediator. From the outside, this appears to be a moderation of the X to Y effect. However, the conceptual difference is that it targets the effect of the mediator that in turn affects the effect of X on Y, whereas a moderator directly affects the effect of X on Y. Also measuring the mediator can also demonstrate statistically that M mediates the relationship between X, M*, and Y.

Experimental manipulations-of-the-mediator designs that seek to show an effect of the mediator on Y manipulate M*. When combined with a manipulation of X, i.e., a concurrent double randomization design, this manipulation of the mediator is equivalent to a moderation design in which the manipulated mediator M* moderates the effect of X on Y to show the effects of X on Y differ as a function of M*. Again, measuring the mediator can also yield evidence that M mediates the relationship between X, M*, and Y.

Furthermore, double randomization designs clarify the problem of whether a variable is or is not a mediator. Demonstrating that X causally influences M, and then in a second experiment, that X and M* causally influence Y provides empirical evidence that M is a mediator, although it does not prove that M* is *not* a moderator. Concurrent double randomization designs in which levels of M* reflect artificially created high and low levels of M* relative to varying freely levels of M* while also measuring M allow comparisons of manipulations-of-mediation effects with measurement-of-mediation effects to demonstrate that the mediating variable is, in fact, a mediator.

A program of research using multiple methods (including measurement-of-mediation and manipulation-of-mediator designs) provides evidence that a variable is a mediator and not a moderator. This limitation challenges theories to better specify models of moderators and mediators. Including a concurrent double randomization design within a series of studies or program of research in which evidence demonstrates that the proposed mediator is affected by the independent variable provides evidence that it is, in fact, a mediator, and not only a moderator.

3.2. Alternative explanations

One challenge underlying manipulation-of-mediator designs arises in that experimental manipulation of the mediator cannot *automatically* rule out alternative explanations of the relationship between M* and Y. Even when the mediator is manipulated in the manipulation-of-mediator designs, it is still possible that confounds may lead to incorrect conclusions. As in all experimental research (whether manipulating X or M*), researchers must demonstrate and/or argue persuasively for why the particular experimental manipulation did not include a covarying confounding variable or simultaneously activate multiple variables (in this case, mediators). A way to address this issue is by systematically manipulating multiple mediators, perform pilot testing of the M* manipulation to ensure it targets one mediator and not others, and measuring all possible relevant confounding mediating variables.

3.3. Precision needed to manipulate mediators

Another challenge is whether the researcher can manipulate mediators with the precision needed to demonstrate mediation. Given that proposed mediators are often continuous, the manipulation of the mediator at a precise level necessary for change could be difficult. In double randomization designs, manipulating X to estimate the causal effect of X on M in Experiment 1 provides evidence for the necessary levels to manipulate the mediator in Experiment 2. Further, experimental manipulations of the mediator could affect other confounded mediators and then the effect of X on Y would have to be considered to be mediated by a combination/package of mediators, and exclusive attribution of the indirect effect to the particular mediator in focus would be difficult. MacKinnon (2008) and Bullock et al. (2010) recommend that if researchers manipulate one mediator, investigators ought to provide justification as to why other variables are not also affected, and to derive a list of alternative mediators and demonstrate that those mediators are unaffected and not driving the results. Similar to challenges of operationalizing an independent variable, the operationalization of a mediator can include confounds. Precise operationalization of the experimental manipulations of the mediator can rectify this situation and experimental researchers are already familiar with this construct validity challenge (Cook & Campbell, 1979). Furthermore, these criticisms are not specific to manipulation-of-mediator designs; they are challenges embedded within all experimental designs.

In past research, these alternative explanations are sometimes addressed in the discussion section but the use of methods to assess confounder bias should help bring possible confounding variables into the forefront in mediation designs (MacKinnon & Pirlott, 2015). To counter claims about mediators being concurrently manipulated, one could manipulate multiple mediators in one experiment (e.g., Sheets & Braver, 1999), or vary one mediator by manipulation while blocking alternative mediators by experimentally keeping them constant (e.g., Neuberg & Fiske, 1987; O'Carroll, Drysdale, Cahill, Shajahan, & Ebmeier, 1999).

3.4. Ease of manipulating mediators

An additional challenge to mediation analysis in general is the extent to which mediating mechanisms easily lend themselves to manipulation and measurement. Cognitive and affective processes are common mediating mechanisms that might be challenging to measure in a way to detect effects of an experimental manipulation (Bullock et al., 2010; MacKinnon, 2008). Similarly, ethical issues may also limit the type of statistical or experimental approach to mediation analysis (Jo, 2013; Mark, 1990). Therefore, manipulation-of-mediator designs are best applied when the mediator can easily be manipulated and measured (Spencer et al., 2005).

For practical purposes, manipulation-of-mediator manipulations that seek to manipulate the mediator are probably easier than manipulation-of-mediator manipulations that target the *effect* of the mediator. Manipulations seeking to vary the mediator minimally require an operationalization of the mediator in a high versus low two-condition dichotomy to demonstrate that the mediator casually affects the dependent variable. Manipulations targeting the *effect* of the mediator require a manipulation of the mediator (blocked or enhanced relative to vary freely) in conjunction with a manipulation of X to demonstrate effects on Y, thus creating a more complicated design, although concurrent double randomization designs which manipulate X and M* require the same minimal number of conditions (minimally a 2 by 2 design).

3.5. Construct validity of manipulated mediators

A similar challenge is how the mediator—as a construct—differs according to whether it is measured or manipulated. Although this construct validity challenge is present in all research, the nature of mediators—that they are typically physiological or psychological processes and therefore more abstract—might make this more difficult than some types of independent variable manipulations, which can be more concrete. Nonetheless, when developed in a larger body of research, replications across alternative operationalizations of the mediator—including measured and manipulated operationalizations—provide converging evidence of the process (Brewer & Crano, 2014; Fabrigar & Wegener, in press; Hüffmeier et al., in press).

3.6. Demonstrating that X causes M

Experimentally manipulating the mediator enables causal inference that M causes Y, but the mediation relationship posits that X causes M that causes Y. Therefore, manipulation-of-mediator designs demonstrate that M* causes Y, but not that X causes M. Double randomization designs show that X causes M (and Y) in Experiment 1 and that M* causes Y in Experiment 2, which enable an understanding of the causal relationship between X and M, X and Y, and M* and Y, but note that the understanding of these relationships comes from two separate experiments. Measuring M as a manipulation check allows testing whether X (and M*) indeed affect M.

3.7. Assessing the indirect effect

These experimental designs are also limited because they cannot assess the entire indirect effect. However, designs which contrast the effects of X on Y when the mediator varies freely versus specifying the mediator to specific levels enables an examination of the indirect effect in the $X \rightarrow M \rightarrow Y$ conditions. Furthermore, designs which manipulate both X and M* and measure M and Y enable an assessment of the effects of X and M* on M, X and M* on Y, and M on Y, thus providing additional evidence of the mediated effect, although not obtained identically to traditional measurement-of-mediation designs.

3.8. Causal effects at the participant- versus group-level

In typical between-subject experimental designs manipulating X and measuring Y, analyses reveal whether a significant difference in Y between levels of X occurred *on average*, not necessarily for individual participants. For accurate causal inference, however, according to the counterfactuals philosophy, one must obtain the causal effect for each participant, i.e., for the participant to simultaneously serve in all conditions of the independent variable and measure their dependent variable. This is, of course, impossible (Imai et al., 2013). Assessing the causal effect of the mediator on the dependent variable further compounds this challenge. For example, it is impossible to observe values of the mediator in the treatment group for participants in the control group and impossible to observe mediator values in the control group for participants in the treatment group. The experimental designs described in this article do not solve this problem ubiquitous to all between-subjects experimental designs, but provide further insight into mediation from a perspective focusing on consistency of the pattern of effects, i.e., X to M, X to Y, and M* to Y.

3.9. Heterogeneity of causality

A final challenge underlying all mediation designs—regardless of measurement-of-mediation or manipulation-of-mediation designs, is whether the effects of X on M and M* on Y occur for different subsets of people, i.e., heterogeneity of causality (Bullock et al., 2010). In measurement-of-mediation designs, the relationship of X on M, X on Y, and M on Y are *average* effects in the sample, rather than effects to specific to each participant. Likewise, in double randomization, concurrent double randomization, and parallel designs the effects of X on M, X on Y, and M* on Y are average effects. In other words, it is possible that the effects of X on M and M on Y (or M* on Y) occur for differ subsets of participants (Bullock et al., 2010; Cerin & MacKinnon, 2009; Imai et al., 2013; Glynn & Quinn, 2011). For example, X affects M only for males (no effect for females) whereas M (or M*) affects Y only for females (no effect for males). Analyses ignoring participant sex would estimate a nonzero indirect effect, *ab*, but the actual indirect effect would be zero for all participants. Importantly, however, this critique is not unique to manipulation-of-mediator designs; this same critique falls upon measurement-of-mediation designs and all research designs (Bullock et al., 2010; Cerin & MacKinnon, 2009).

The easiest way to overcome the problem of heterogeneity of causality is to use within-subjects or repeated measures designs in which participants are randomly assigned to the order of levels of X. Within-subjects designs, however, incur carryover effects that limit their broad usage. Within-subjects designs satisfy the potential outcomes or counterfactual model of causal inference, which suggests the only way to infer causality is to know the outcome for a participant in *both conditions* of the independent variable. Assuming that the use of within-subjects designs is possible, a simple within-subjects measurement-of-mediation design in which participants are randomly assigned to the order of X condition and M and Y measured, enables an examination of whether the difference between conditions on M is in the same direction as Y for all participants. Although examining the *mean differences* of X on M and X on Y could show significant differences between groups when the mediated effect is zero, correlational analyses would detect whether there was a significant relationship between M and Y using the difference scores in each condition for M and Y.

Although using a within-subjects design in which only X is randomly assigned overcomes the heterogeneity of causality problem, it fails to provide sufficient causal inference of the M to Y relationship. The concurrent double randomization design provides further evidence of the causal effect of M* on Y while also overcoming the heterogeneity of causality problem. The heterogeneity of causality limitation suggests two main effects: the effect of X on M and the effect of M* on Y, which could occur for two different subsets of people. Using a within-subjects concurrent double randomization manipulation-of-mediation design in which X and M* are concurrently manipulated and M and Y are both measured enables researchers to examine whether the effects of X on M occur in the same direction for the same participants as the effects of M* on Y by examining a correlation between the difference scores of M as a function of X and Y as a function of M*.

The within-subjects study design, however, is limited to studies that can accommodate carry over effects. Imai et al. (2013) described a modified version of the within-subjects design that they call *crossover designs*: participants participate in both the treatment and control conditions, in a randomized order. This design assumes there are no carry-over effects and appears to require a significant amount of time to pass between participation in the first wave of the study and the second wave of the study, so that the order effects wear off and so that the indirect effect from wave one can be calculated and implemented into wave two. In the first wave, participants are randomly assigned to levels of X and M and Y are measured. In the second wave, participants are exposed to the other level of X that they did not participate in during the first wave and M and Y are measured. This enables a calculation of the difference scores from both conditions of X on M and Y and correlational analyses to determine whether the mean differences in M and Y parallel the correlation. Extending this to manipulation-of-mediator designs, participants would be randomly assigned to levels of X and M*, in a randomized order, and measure M and Y. Calculate difference scores in M as a function of X and Y as a function of M* and examine the correlation between the difference scores.

The extended time between studies in Imai et al.'s cross over design makes the study more feasible for some between-subjects studies that cannot be within-subjects studies occurring at the same time point while also managing the heterogeneity of causality problem. This design, however, is not a perfect solution for all traditional between-subjects studies in which carryover effects limit its practicality. Therefore, the problem of heterogeneity of causality is one that warrants further consideration.

4 Discussion

Manipulation-of-mediator and measurement-of-mediation designs to determine causal mediation provide overlapping assessments of the mediating mechanisms (e.g., Brewer, 2000; Cook & Groom, 2004; Mark, 1986, 1990; Sigall & Mills, 1998; Smith, 2000; Spencer et al., 2005). What benefits do manipulation-of-mediator designs provide researchers relative to measurement-of-mediation designs, what are the costs, and do the benefits ultimately outweigh the costs?

Measurement-of-mediator designs provide evidence of the causal relationship of X on M and X on Y and provide an estimate of the indirect effect, given all paths can be determined from one experimental study. However, the largest shortcoming of measurement-of-mediation designs is that they cannot actually demonstrate the causal effect of M on Y, and so the ability to infer that M caused Y is limited and subject to alternative explanations.

Manipulation-of-mediation designs can solve the question of whether M* causes Y by randomly assigning participants to levels of M*, measuring Y, and examining the effects of M* on Y as demonstrated through significant mean differences on Y as a function of conditions of M*. Yet manipulation-of-mediation designs are not without challenges, for example including distinguishing between mediators and moderators, alternative explanations for the relationship between M and Y, the precision and ease of manipulating mediators, construct validity of the manipulated mediator, demonstrating that X causes M, assessing the indirect effect, knowing the causal effect at the individual- versus group-level, and heterogeneity of causality.

However, the greatest benefit of manipulation-of-mediator designs is to increase the accuracy of conclusions about causal inference of the M* to Y relation. To infer that one variable causes another, the hypothesized variable must precede the outcome variable in time, the two variables must covary, and alternative explanations cannot account for the relationship between the two variables (Shadish et al., 2002). Considering the M to Y relationship, both measurement-of-mediation and experimental approaches to mediation satisfy some of the criteria for causality. Measurement-of-mediation designs allow researchers to demonstrate covariation (via correlation between M and Y) as does an experiment (via an effect of M* on Y). Nonetheless, manipulation-of-mediator designs provide better evidence for causality than measurement-of-mediation approaches by providing additional information than measurement-of-mediation designs. Measurement-of-mediation designs do not always provide clear demonstration of temporal precedence of M to Y. If experimenters manipulate X and simply measure M and Y, temporal precedence that M precedes Y is not clear, although longitudinal designs provide information on whether change in M preceded change in Y. Experimental manipulations of M* provide evidence that M* preceded Y temporally and caused Y. Finally, only unconfounded manipulations of the mediating variable reduce the plausibility of alternative explanations for the relationship between M* and Y, such as Y causing M or M and Y being unrelated except through a confounding third variable. Manipulation-of-mediator designs provide stronger and more rigorous evidence of the causal relationship between M* and Y, and therefore should be included in programmatic research dedicated to demonstrating causal processes.

References

Baron RM, Kenny DA. The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. Journal of Personality & Social Psychology. 1986; 51(6):1173–1182. http://dx.doi.org/10.1037/0022-3514.51.6.1173. [PubMed: 3806354]

Brewer, MB. Research design and issues of validity. In: Reis, HT.; Judd, CM., editors. Handbook of research methods in social and personality psychology. Cambridge University Press; New York: 2000. p. 3-16.

Brewer, MB.; Crano, WD. Research design and issues of validity. In: Reis, HT.; Judd, CM., editors. Handbook of research methods in social and personality psychology. 2nd ed. Cambridge University Press; New York, NY: 2014. p. 11-26.

Bullock JG, Green DP, Ha SE. Yes, but what's the mechanism? (don't expect an easy answer). Journal of Personality and Social Psychology. 2010; 98:550–558. http://dx.doi.org/10.1037/a0018933. [PubMed: 20307128]

Cerin E, MacKinnon DP. A commentary on current practice in mediating variable analyses in behavioural nutrition and physical activity. Public Health Nutrition. 2009; 12(8):1182–1188. http://dx.doi.org/10.1017/S1368980008003649. [PubMed: 18778534]

Cialdini RB. We have to break up. Perspectives on Psychological Science. 2009; 4(1):5–6. http://dx.doi.org/10.1111/j.1745-6924.2009.01091.x. [PubMed: 26158821]

Cook, TD.; Campbell, DT. Quasi-experimentation: Design and analysis issues for field settings. Rand-McNally; Chicago: 1979.

Cook, TD.; Groom, C. The methodological assumptions of social psychology: The mutual dependence of substantive theory and method choice. In: Sansone, C.; Morf, CC.; Panter, AT., editors. The Sage handbook of methods in social psychology. Sage Publications; Thousand Oaks: 2004. p. 19-44.

Cooper J, Zanna MP, Taves PA. Arousal as a necessary condition for attitude change following induced compliance. Journal of Personality and Social Psychology. 1978; 36:1101–1106. http://dx.doi.org/10.1037/0022-3514.36.10.1101. [PubMed: 722473]

Fabrigar LR, Wegener DT. Conceptualizing and evaluation the replication of research results. Journal of Experimental Social Psychology: Special Issue on Rigorous and Replicable Methods in Social Psychology. 2015 (in press).

Glynn AN, Quinn KM. Why process matters for causal inference. Political Analysis. 2011; 19(3):273–286.

Halberstadt J. My plan for more efficient reviewing [Society for Personality and Social Psychology electronic mailing list email]. Jun 8.2010 Retrieved from https://groups.google.com/forum/#!topic/spsp-discuss/4Q-TZlsLrRg.

Hayes AF, Scharkow M. The relative trustworthiness of inferential tests of the indirect effect in statistical mediation analysis: Does method really matter? Psychological Science. 2013; 24:1918–1927. http://dx.doi.org/10.1177/0956797613480187. [PubMed: 23955356]

Hill, AB. Principles of medical statistics. 9th ed. Oxford; New York: 1971.

Holland PW. Causal inference, path analysis, and recursive structural equation models. Sociological Methodology. 1988; 18:449–484. http://dx.doi.org/10.2307/271055.

Hüffmeier J, Mazei J, Schultze T. Reconceptualizing replication as a sequence of different studies: A replication typology. Journal of Experimental Social Psychology: Special Issue on Rigorous and Replicable Methods in Social Psychology. 2015 (in press).

Imai K, Tingley D, Yamamoto T. Experimental designs for identifying causal mechanisms. Journal of the Royal Statistical Society. 2013; 176:5–51.

Jacoby J, Sassenberg K. Interactions not only tell us *when*, but can also tell us *how*: Testing process hypotheses by interaction. European Journal of Social Psychology. 2011; 41:180–190. http://dx.doi.org/10.1002/ejsp.762.

James LR. The unmeasured variables problem in path analysis. Journal of Applied Psychology. 1980; 65:415–421.

Jo B. Discussion on the paper by Imai, Tingley, and Yamamoto. Journal of the Royal Statistical Society. 2013; 176:40–41.

Judd CM, Kenny DA. Process analysis: Estimating mediation in treatment evaluations. Evaluation Review. 1981; 5:602–619. http://dx.doi.org/10.1177/0193841X8100500502.

Kashy DA, Donnellan MB, Ackerman RA, Russell DW. Reporting and interpreting research in PSPB: Practices, principles, and pragmatics. Personality and Social Psychology Bulletin. 2009; 35:1131–1142. http://dx.doi.org/10.1177/0146167208331253. [PubMed: 19458094]

Li YJ, Johnson KA, Cohen AB, Williams MJ, Knowles ED, Chen Z. Fundamental(ist) attribution error: Protestants are dispositionally focused. Journal of Personality and Social Psychology. 2012; 102(2):281–290. http://dx.doi.org/10.1037/a0026294. [PubMed: 22082060]

MacKinnon, DP. Introduction to statistical mediation analysis. Erlbaum; New York: 2008.

MacKinnon DP, Pirlott AG. Statistical approaches for enhancing causal interpretation of the M to Y relation in mediation analysis. Personality and Social Psychology Review. 2015; 19(1):30–43. http://dx.doi.org/10.1177/1088868314542878. [PubMed: 25063043]

MacKinnon DP, Lockwood CM, Hoffman JM, West SG, Sheets V. Comparison of methods to test mediation and other intervening variable effects. Psychological Methods. 2002; 7:83–104. http://dx.doi.org/10.1037//1082-989X.7.1.83. [PubMed: 11928892]

Mark, MM. Validity typologies and the logic and practice of quasi-experimentation. In: Trochim, WMK., editor. Advances in quasi-experimental design and analysis. Jossey-Bass; San Francisco: 1986. p. 47-66.

Mark MM. From program theory to tests of program theory. New Directions for Program Evaluation. 1990; 47:37–51. http://dx.doi.org/10.1002/ev.1553.

McDonald RP. Haldane's lungs: A case study in path analysis. Multivariate Behavioral Research. 1997; 32:1–38. http://dx.doi.org/10.1207/s15327906mbr3201_1. [PubMed: 26751104]

Neuberg SL, Fiske ST. Motivational influences on impression formation: Outcome dependency, accuracy-driven attention, and individuating processes. Journal of Personality and Social Psychology. 1987; 53:431–444. http://dx.doi.org/10.1037/0022-3514.53.3.431. [PubMed: 3656080]

O'Carroll RE, Drysdale E, Cahill L, Shajahan P, Ebmeier KP. Stimulation of the noradrenergic system enhances and blockade reduces memory for emotional material in man. Psychological Medicine. 1999; 29:1083–1088. http://dx.doi.org/10.1017/S0033291799008703. [PubMed: 10576300]

Pearl, J. Causality: models, reasoning, and inference. Cambridge University Press; 2000.

Quinones-Vidal E, Lopez-Garcia JJ, Penaranda-Ortega M, Tortosa-Gil F. The nature of social and personality psychology as reflected in JPSP, 1965–2000. Journal of Personality and Social Psychology. 2004; 86(3):435–452. http://dx.doi.org/10.1037/0022-3514.86.3.435. [PubMed: 15008647]

Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. Epidemiology. 1992; 3:143–155. Stable URL: http://www.jstor.org/stable/702894. [PubMed: 1576220]

Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of Educational Psychology. 1974; 66:688–701. http://dx.doi.org/10.1037/h0037350.

Rubin DB. Assignment to treatment group on the basis of a covariate. Journal of Educational Statistics. 1977; 2:1–26. http://dx.doi.org/10.2307/1164933.

Rucker DD, Preacher KJ, Tormala ZL, Petty RE. Mediation analysis in social psychology: Current practices and new recommendations. Social and Personality Psychology Compass. 2011; 5(6): 359–371. http://dx.doi.org/10.1111/j.1751-9004.2011.00355.x.

Shadish, WR.; Cook, TD.; Campbell, DT. Experimental and quasi-experimental designs for generalized causal inference. Houghton-Mifflin; Boston: 2002.

Sheets VL, Braver SL. Organization status and perceived sexual harassment: Detecting the mediators of a null effect. Personality and Social Psychology Bulletin. 1999; 25:1159–1171. http://dx.doi.org/10.1177/01461672992512009.

Sigall H, Mills J. Measures of independent variables and mediators are useful in social psychology experiments: But are they necessary? Personality and Social Psychology Review. 1998; 2:218–226. [PubMed: 15647156]

Smith ER. Beliefs, attributions, and evaluations: Nonhierarchical models of mediation in social cognition. Journal of Personality and Social Psychology. 1982; 43(2):248–259. http://dx.doi.org/10.1037/0022-3514.43.2.248.

Smith, ER. Research design. In: Reis, HT.; Judd, CM., editors. Handbook of research methods in social and personality psychology. Cambridge University Press; New York: 2000. p. 17-39.

Smith ER. Editorial. Journal of Personality and Social Psychology. 2012; 102(1):1–3. http://dx.doi.org/10.1037/a0026676. [PubMed: 22514799]

Spencer SJ, Zanna MP, Fong GT. Establishing a causal chain: Why experiments are often more effective than mediational analyses in examining psychological processes. Journal of Personality and Social Psychology. 2005; 89:845–851. http://dx.doi.org/10.1037/0022-3514.89.6.845. [PubMed: 16393019]

Stone-Romero EF, Rosopa PJ. The relative validity of inferences about mediation as a function of research design characteristics. Organizational Research Methods. 2008; 11(2):326–352. http://dx.doi.org/10.1177/1094428107300342.

VanderWeele TJ, Vansteelandt S. Conceptual issues concerning mediation, interventions, and composition. Statistics and Its Interface (Special Issue on Mental Health and Social Behavioral Science). 2009; 2:457–468.

West SG, Thoemmes F. Campbell's and Rubin's perspectives on causal inference. Psychological Methods. 2010; 15(1):18–37. http://dx.doi.org/10.1037/a0015917. [PubMed: 20230100]

Word CO, Zanna MP, Cooper J. The nonverbal mediation of self-fulfilling prophecies in interracial interaction. Journal of Experimental Social Psychology. 1974; 10:109–120. http://dx.doi.org/10.1016/0022-1031(74)90059-6.

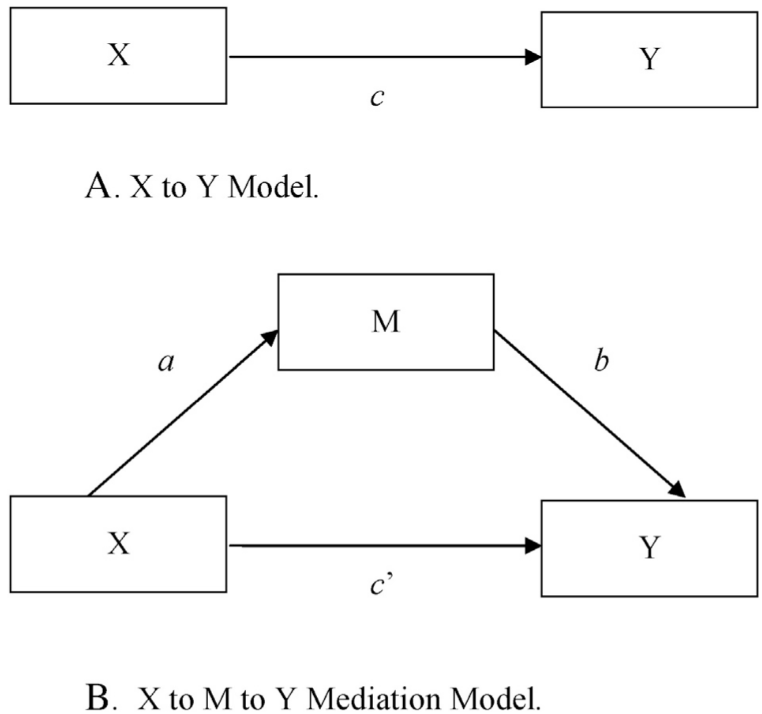

Fig. 1.

Single mediator model in which X is randomized and M and Y are measured. The *a* coefficient reflects the effect of X on M; the *b* coefficient reflects the statistical effect of M on Y, controlling for X; the *c* coefficient reflects the total effect of X on Y, not controlling for M; and the *c*' represents the direct effect of X on Y, controlling for M.

Table 1

Example of data from a concurrent double randomization design. Adapted from Cooper et al. (1978).

X Condition	M* condition		
	Sedative	**Placebo**	**Stimulant**
High choice	8.6_a	14.7_b	20.2_c
Low choice	8.0_a	8.3_a	13.9_b

Note, *n*=10 subjects per cell. Higher means on the 31-point scale indicate greater agreement with the attitude-discrepant essay. Cell means with different subscripts are different from each other at the .05 level by the Newman–Keuls procedure. The mean in the survey control condition is 7.9_a.