Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Experimental Social Psychology 40 (2004) 424–431

Journal of Experimental Social Psychology

www.elsevier.com/locate/jesp

Adjusting researchers' approach to adjustment: On the use of covariates when testing interactions☆

Vincent Y. Yzerbyt,[a,*] Dominique Muller,[b,c] and Charles M. Judd[b]

[a] Department of Psychology, Catholic University of Louvain at Louvain-la-Neuve, Place du Cardinal Mercier 10, B-1348 Louvain-la-Neuve, Belgium
[b] University of Colorado at Boulder, USA
[c] University Pierre Mendes France at Grenoble, France

Received 6 June 2003; revised 29 September 2003

Abstract

A common design in social psychology involves the use of two independent variables, an experimental manipulation and a measured individual difference, and the interest is in the interaction between them. In such designs, there are often obvious covariate(s), correlated with the measured independent variable, which the researcher wishes to control. Typically this is done by including the covariate in the analytic model. We show that in most cases this is an inadequate model. In general, the interaction between the two independent variables will be estimated without bias only when the interaction between the covariate and the manipulated independent variable is included in the analysis. We present simulations showing the factors affecting the magnitude of the bias and provide a survey of recent social psychological literature illustrating the frequency of the problem. Finally, we discuss cases where both independent variables are manipulated and the covariate is a potential mediator.

Keywords: Covariate; Analysis of covariance; ANCOVA; Mediation; Adjustment; Power; Interaction; Stereotype threat

Introduction

It is well known in the methodological literature (e.g., Judd & McClelland, 1989) that there are two reasons why one might choose to include a covariate in an analytic model: power and adjustment. The first is served when the covariate is highly related to the dependent variable, but unrelated to the independent variables of interest. This condition is generally assured with a covariate that is measured prior to the delivery of independent variables to which participants have been randomly assigned. In this case, the inclusion of the covariate in the analytic model leads to a reduction in the variance of the residuals, thereby increasing the power of the test of the independent variables. Additionally, because there is no redundancy between the covariate and the independent variables, the inclusion of the covariate has no impact on the effect estimates associated with the independent variables. In other words, the adjusted means will not be different from the unadjusted means.

The adjustment function occurs when the covariate is related to the independent variables. The inclusion of the covariate in the analytic model accordingly results in "adjustments" in the effect estimates associated with the independent variables. The effects of the independent variables are then partial effects, controlling for or removing the impact of the covariate. There tend to be two classic cases where these adjusted effect estimates are sought out. The first is when one or more of the independent variables is confounded with the covariate and one wishes to remove the confound in assessing the impact of the independent variables. This is the classic adjustment function of analysis of covariance when used

☆ This article was written while the first author was on a sabbatical leave at the University of Colorado at Boulder, USA. The second author benefited from a Lavoisier scholarship of the Ministère des Affaires Étrangères, France. Additionally, this work was supported by a grant from the National Institute of Mental Health R01MH45049 to the third author. We express our gratitude to Markus Brauer, Olivier Corneille, Jay Hull, Dave Kenny, Christoph Klauer, Olivier Klein, and Christophe Labiouse for their comments on previous versions.

* Corresponding author.
E-mail address: vincent.yzerbyt@psp.ucl.ac.be (V.Y. Yzerbyt).

0022-1031/$ - see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/j.jesp.2003.10.001

in the analysis of quasi-experimental and correlational designs. The second case is when the covariate is suspected to be a mediator of the effects of the independent variables. One then controls for the covariate in order to examine whether the partialled effect estimates of the independent variables are reduced in magnitude, compared to the zero-order effects, thereby supporting the mediational conjecture. The distinction between these two cases is a theoretical rather than analytic one. In the first the covariate is simply confounded with the independent variables; in the latter it is assumed to be caused by the independent variables.

An interesting and common use of the adjustment function occurs when the covariate is related to one independent variable that is thought to interact with another independent variable with which the covariate is uncorrelated. This situation often arises when the first independent variable is a measured variable, for instance a stable disposition of research participants, while the second one is a manipulated one, to which participants have been randomly assigned. Designs that cross these independent variables, one measured and one manipulated, are used routinely in social psychology where a central theoretical concern has focused on situation by person interactions, i.e., the impact of situational variables depends on some stable individual difference. Typically this has been done by exposing participants to different levels of a situational manipulation and examining whether the impact of that manipulation depends on some measured individual difference.

In such designs, it is often the case that there exists a covariate that is related to the measured variable, either as a confounded variable or a potential mediator, and one seeks to adjust for its effect as a covariate in the analytic model. Here the theoretical focus is on the predicted interaction between the measured independent variable and the manipulated one. And one wishes to estimate the impact of that interaction while controlling for the covariate that is correlated with the measured independent variable.

In general, the analysis is conducted by including the covariate in the model while testing the predicted interaction between the manipulated independent variable and the measured one. The point of the present paper is to suggest that this analytic model is frequently incomplete and results in a biased effect estimate for the interaction. Whenever the covariate is confounded with the measured independent variable and additionally, participants have been randomly assigned to levels of the manipulated independent variable, then the predicted interaction between the measured and manipulated independent variables will be confounded with the interaction between the covariate and the manipulated independent variable. When this is the case, then the interaction between the manipulated and measured independent variable will be estimated with bias unless both the covariate and the interaction between the covariate and the manipulated variable are controlled.

We illustrate the problem with an example drawn from a highly cited recent article in social psychology. Steele and Aronson (1995) argued that when a test is defined as diagnostic of an ability and when the absence of that ability is stereotypically associated with a particular ethnic group, stereotype threat ensues and performance on the test is impaired for members of that group. To support this argument, they recruited participants from different ethnic groups (the measured independent variable) and randomly assigned them to conditions where a test is defined to be diagnostic of an ability or not (the manipulated independent variable). Their central prediction concerned the interaction of these two independent variables: diagnosticity will lower performance only among that group for whom the ability in question is stereotypically threatening. Because the ethnic groups that are used typically differ in prior performance in the ability domain (for example, in SAT performance), the analytic model includes that prior performance measure as a covariate when testing the predicted interaction.

But there really are two alternative interactions that compete here. On the one hand, there is the one that is theoretically preferred by Steele and Aronson (1995): diagnosticity lowers performance only among members of the ethnic group for whom the (lack of) ability is stereotypically threatening. On the other hand, maybe this effect really has nothing to do with ethnicity per se. Maybe it is a more general performance induced effect: diagnosticity lowers performance among those people who perform less well in the domain for which the test is supposedly diagnostic. To argue for the first (and preferred) of these two interactions, Steele and Aronson should have controlled for the second. They did not. And simply including prior performance as a covariate does not adequately control for its interaction with diagnosticity. If in fact prior performance does interact with diagnosticity, then the estimated effect of the ethnicity by diagnosticity interaction will be biased unless the prior performance by diagnosticity interaction is included in the analytic model. Including the covariate alone does nothing to reduce this bias.

In 1992, Hull, Tedlie, and Lehn published a short article in *Personality and Social Psychology Bulletin* in which they argued a very similar point to the one we are making. Three issues motivate us to believe that the point needs to be made again, more forcefully. First, as the example we have just given illustrates (and, as we review later, there are many other recent cases), the lesson from the Hull et al. paper has not been attended to. Second, this lack of attention has largely been caused, we think, by the fact that Hull et al. article confined its discussion of the issue to research in personality psychology. Indeed the title of the article is

"Moderator variables in personality research." As the example from Steele and Aronson (1995) illustrates, the issue they (and we) point to deserves attention from the broader social psychological audience. Moreover, the problem is not confined to cases where the covariate is related to a measured independent variable but is also relevant for situations in which all independent variables are manipulated and the covariate assumes the role of a potential mediator. Third, Hull et al. discuss the issue in general terms. They did not actually demonstrate the bias nor explore factors that affect its magnitude. In what follows, we do both.

In the following section we demonstrate the bias analytically. We then report some simulations that examine factors that affect its magnitude. Following this, we report our survey of articles published in 2002 in four leading journals in social psychology. This survey documents the frequency of the problem and the general failure to appropriately deal with it. In a final section, we address the issue of the use of covariates in mediational analyses.

Demonstration of bias

Let us assume that X_{1i} is a measured independent variable, representing some stable disposition of participants. We assume it is normally distributed in the population with an expected value of zero and a variance of $\sigma^2_{X_1}$.[1]

Next, assume that X_{2i} is a manipulated independent variable, with two levels to which participants are randomly assigned with equal probabilities. We will adopt a contrast-code convention, so that the levels of X_{2i} are coded as +1 and 1. Accordingly, the expected value of X_{2i} is also zero and its variance equals 1.0.

Third, we assume there is some measured covariate, C_i, again normally distributed with an expected value of zero and a variance of σ^2_C. Additionally, we assume it covaries with X_{1i} in the population, $\sigma_{X_1,C} \neq 0$. This covariance may be due to a variety of functional relationships between the two: one may cause the other or some other variable may be responsible for variation in both. Although the distinctions among the functional relationships between C_i and X_{1i} are theoretically important, analytically all that matters is that the two covary.[2]

Finally, assume there is some variable Y_i that is considered the dependent variable, and that its values are a function of the previous variables, interactions among those variables, and residual normally distributed random errors:

$$Y_i = \beta_{11}X_{1i} + \beta_{12}X_{2i} + \beta_{13}C_i + \beta_{14}X_{1i}X_{2i} + \beta_{15}C_iX_{2i} + \varepsilon_{1i}, \quad (1)$$

where $X_{1i}X_{2i} = X_{1i} \cdot X_{2i}$ and $C_iX_{2i} = C_i \cdot X_{2i}$. This is assumed to be the "true" population model, meaning that it specifies the factors responsible for variation in Y_i.

The expected values of all variables are zero (including the residual) and hence in the population there is no intercept. Notice we are assuming that the manipulated independent variable, X_{2i}, interacts with both C_i and X_{1i}, but that these two variables, while correlated, do not interact with each other in affecting Y_i.

The question of bias in the estimation of the coefficient of the $X_{1i}X_{2i}$ interaction focuses on whether its parameter estimate is biased if one estimates a model in which the C_iX_{2i} interaction is not included as a predictor. In other words, assuming that one estimates the following misspecified model (i.e., the classic ANCOVA model):

$$Y_i = \beta_{21}X_{1i} + \beta_{22}X_{2i} + \beta_{23}C_i + \beta_{24}X_{1i}X_{2i} + \varepsilon_{2i} \quad (2)$$

the question is whether β_{24} differs from its true value, β_{14}.

As we demonstrate in the Appendix A, from these two expressions and the assumptions we made, we can derive

$$\beta_{24} = \beta_{14} + \beta_{15}\frac{\sigma_{C,X_1}}{\sigma_{X_1^2}}.$$

Accordingly, only in two conditions will the parameter estimate in the misspecified model (β_{24}) equal the parameter in the correct model (β_{14}). The first is when C_i and X_{1i} do not covary. The second condition is when β_{15}, the effect of the C_iX_{2i} interaction, equals zero. In other words, assuming that C_i and X_{1i} are related and that the effect of the C_iX_{2i} interaction is not zero, then the coefficient for the $X_{1i}X_{2i}$ interaction will be biased unless one includes the C_iX_{2i} interaction in the model. Note also that if C_i and X_{1i} have equal variances, then the ratio $\sigma_{C,X_1}/\sigma^2_{X_1}$ equals the correlation between the two variables. In this case, the degree of bias will be a linear function of that correlation. When the correlation is zero, there will be no bias. With a correlation approaching 1.00, the biased coefficient will approach $\beta_{14} + \beta_{15}$.

One additional result is of interest. Suppose a misspecified model were estimated with both C_i and the C_iX_{2i} interaction omitted as predictors

$$Y_i = \beta_{31}X_{1i} + \beta_{32}X_{2i} + \beta_{33}X_{1i}X_{2i} + \varepsilon_{3i}. \quad (3)$$

Because the expected covariance between C_i and $X_{1i}X_{2i}$ equals zero (as shown in the Appendix A), the coefficient for the $X_{1i}X_{2i}$ interaction in this second

[1] The derivation that follows applies even when the independent variable is a dichotomous and potentially manipulated one as well, so long as the distribution of the independent variable is symmetrical (i.e., equal numbers in both categories).

[2] We will assume that the two variables have been scaled so that their covariance is positive.

misspecified model (β_{33}) would exactly equal the biased coefficient in the first misspecified model (β_{24}). In other words, including the covariate in the model, but failing to include its interaction with X_{2i} does nothing to eliminate the bias in the estimate of the effect of the $X_{1i}X_{2i}$ interaction.

Simulations

In order to demonstrate the bias more concretely, we conducted Monte Carlo simulations. The results of these are necessarily consistent with the analytic results just presented. However, the simulation results provide a fairly dramatic illustration of the problems to be encountered.

In these simulations, we focused in particular on Type I error rates. Accordingly in the true model

$$Y_i = \beta_{11}X_{1i} + \beta_{12}X_{2i} + \beta_{13}C_i + \beta_{14}X_{1i}X_{2i} + \beta_{15}C_iX_{2i} + \varepsilon_{1i} \quad (1)$$

we constrained β_{14} to exactly equal zero while fixing β_{11}, β_{12}, and β_{13} to 1. The variance of the residuals, $\sigma^2_{\varepsilon_1}$, was fixed at 5 and the variances of both C_i and X_{1i} were set at 1. We then varied the magnitude of β_{15} (between .00 and 1.00 in increments of .20) and the magnitude of the correlation between C_i and X_{1i} (between .00 and .80 in increments of .20). We allowed these two factors to vary because the derivation we have just presented shows that they determine the degree to which β_{24} is biased. Since β_{14} was fixed at zero, finding a significant β_{24} constitutes a Type I error. Accordingly, in our simulations we examined not only the mean values of β_{24} but also the relative frequency with which β_{24} was significant. At each combination of the parameters that varied, we conducted 1000 simulation trials, randomly sampling 100 cases each time.

Reported in Table 1 are the mean values (and standard errors) of β_{24}, the coefficient for the $X_{1i}X_{2i}$ interaction in the misspecified ANCOVA model (Eq. (2)). Given our simulation specifications, the analytic results we gave earlier mandate that the expected value for β_{24} equals the product of the two parameters that vary across the simulations, i.e., β_{15} and the correlation between C_i and X_{1i}. Within sampling error, this is true in every case.

Reported in Fig. 1 are the probabilities of Type 1 errors from the simulation as a function of the two factors that varied. Recall that in the simulations there was no effect of the $X_{1i}X_{2i}$ interaction in the correctly specified model. The probabilities reported in Fig. 1 are the probabilities that the coefficient for that interaction was significant in the incorrectly specified model that omitted the C_iX_{2i} interaction. As should be the case, these Type I error rates equal .05 when either the effect for the C_iX_{2i} interaction equals zero or when C_i and X_{1i} are uncorrelated. However, they increase as these factors depart from zero. Although this is the general result, the specific values given in Fig. 1 are of course contingent on the specifications used in these simulations.

One additional set of results from these simulations is of considerable interest. In the correctly specified model, we could assess the probability that the effect of the C_iX_{2i} interaction emerged as significant, assuming in fact its parameter β_{15} was something other than zero. This amounts to the statistical power to detect the C_iX_{2i} interaction effect in the correctly specified model. These power results are given in Fig. 2. What is remarkable here is that it is frequently the case that there will be substantial bias in the estimate of the $X_{1i}X_{2i}$ interaction effect in the misspecified model (i.e., β_{24}) even when the power to detect the presence of the C_iX_{2i} interaction effect in the correctly specified model is relatively low. For instance, when the true effect of the C_iX_{2i} interaction is .6 and the correlation between C_i and X_{1i} is .80, the amount of bias in estimating the $X_{1i}X_{2i}$ interaction effect in the misspecified model (i.e., β_{24}) is considerable. Instead of the true effect of zero, the average estimated coefficient for the interaction will be .48 in the misspecified model and the probability of a Type I error in that model equals .508. Yet, the power to detect the actual effect of the C_iX_{2i} interaction effect in the correctly specified model is quite low (.349). What this means is that a significance test of the C_iX_{2i} interaction in the correctly specified model should not be used to decide whether or not to retain it in the model. Even if the C_iX_{2i} interaction is found to be non-significant, the test of the $X_{1i}X_{2i}$ interaction effect in the misspecified model may well be seriously biased.

Table 1
β_{24} in the incomplete model as a function of the effect of β_{15} and the correlation between C and X_1 (standard error in parentheses)

	$r(C, X_1)$				
β_{15}	0	.2	.4	.6	.8
0	−0.018 (0.0162)	−0.017 (0.0138)	0.002 (0.0128)	0.017 (0.0107)	−0.004 (0.0073)
.2	−0.026 (0.0168)	0.020 (0.0152)	0.067 (0.0126)	0.110 (0.0107)	0.163 (0.0071)
.4	0.025 (0.0167)	0.068 (0.0145)	0.167 (0.0130)	0.263 (0.0105)	0.320 (0.0070)
.6	0.001 (0.0166)	0.132 (0.0151)	0.241 (0.0127)	0.365 (0.0105)	0.479 (0.0074)
.8	0.007 (0.0165)	0.175 (0.0146)	0.305 (0.0127)	0.487 (0.0111)	0.638 (0.0074)
1	−0.032 (0.0168)	0.214 (0.0150)	0.421 (0.0133)	0.594 (0.0109)	0.789 (0.0077)

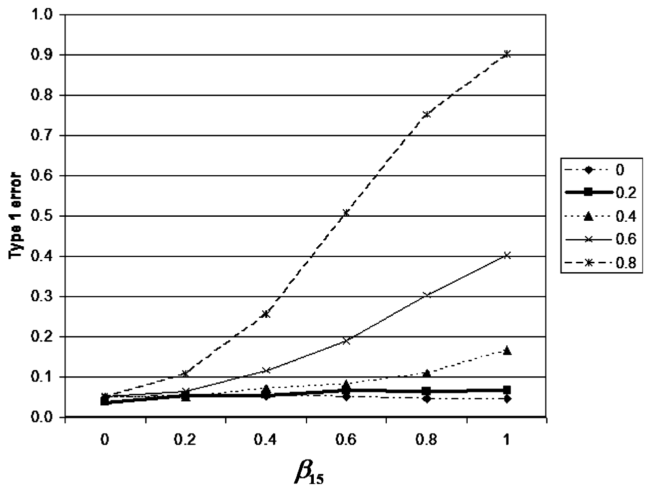

Fig. 1. Type 1 error in the test of β_{24} in the incomplete model as a function of β_{15} and the correlation between C and X_1.

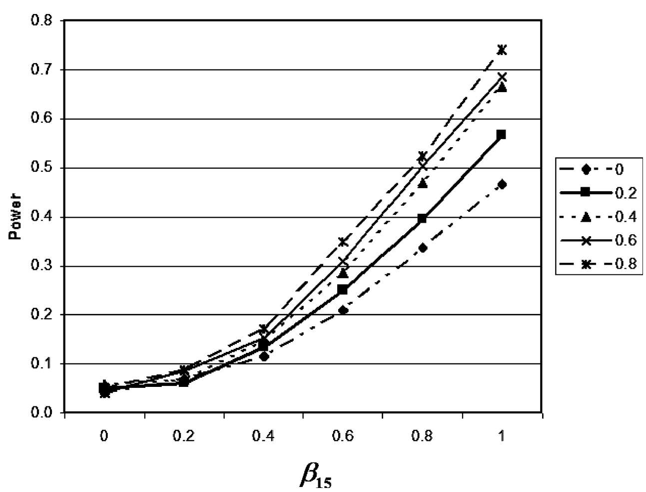

Fig. 2. Power of the test of β_{15} in the complete model as a function of β_{15} and the correlation between C and X_1.

A survey of the literature

The statistical literature that is most frequently encountered on analysis of covariance is that contained in the classic books on factorial analysis of variance (e.g., Keppel & Zedeck, 1989; Kirk, 1995; Maxwell & Delaney, 1990; Myers & Well, 1991; Winer, Brown, & Michels, 1991). Although these treatments talk about a related issue in analysis of covariance, the homogeneity of regression assumption, they do not discuss the issue we are addressing. What we have shown is that when a covariate is confounded with one independent variable, and when that independent variable interacts with a second one, the test of the interaction will be biased unless the covariate interaction is controlled.

As we mentioned earlier, one previously published paper we are aware of explicitly addresses the issue (Hull, Tedlie, & Lehn, 1992). However, it seemed to us that the issue continues to be problematic in the literature. To demonstrate this, we searched the recent social psychological literature to determine how often the problem occurred. We surveyed all articles that appeared in the 2002 issues of the *Journal of Personality and Social Psychology*, the *Journal of Experimental Social Psychology*, the *Personality and Social Psychology Bulletin*, and the *European Journal of Social Psychology*.

Table 2
Survey of the 2002 literature

	JESP	PSPB	JPSP	EJSP
One-way	0	3	5	2
Repeated-measures analysis	1	2	3	1
Power analysis	0	1	5	1
Irrelevance	1	2	4	1
Misspecification	5	7	8	1
Total # of papers	7	15	25	6
% Relevant	71%	47%	32%	17%
% Wrong/relevant	100%	100%	100%	100%

Note. JESP, Journal of Experimental Social Psychology; PSPB, Personality and Social Psychology Bulletin; JPSP, Journal of Personality and Social Psychology; EJSP, European Journal of Social Psychology.

We compiled a list of all articles that mentioned using analysis of covariance or ANCOVA. As shown in Table 2, we found 53 such articles. We examined whether the covariate was used in conjunction with one or several independent variables. We eliminated those cases where only a single independent variable was used.

We next eliminated those studies where all but one of these varied within subjects. Although the issue we are raising is potentially important when using covariates in the analysis of repeated measures, typically covariates in such models vary only between subjects. In such situations, the standard analytic packages used to conduct repeated measures analyses automatically include the covariate by within-subject interactions in the models.

Next, we eliminated those cases in which the authors used a covariate that was uncorrelated with all independent variables for the sole purpose of increasing power. Our concerns are relevant only when the covariate is correlated with one of the independent variables.

Another category of studies involved cases in which the authors made clear predictions for and indeed only found main effects of their independent variables on their dependent variables. Our focus is on the use of covariates in contexts where one is primarily interested in the interaction between two independent variables, with one of them correlated with a covariate.

Finally, we were left with those papers (21 in total; 40% of all papers using analysis of covariance) where there was a clear prediction of an interaction between two independent variables and where a covariate was included that likely was correlated with at least one of them. In all of these papers, only the covariate was included in the analytic model. None of them tested the predicted interaction while including the relevant covariate interaction. Thus, all of these papers seem to have made the error that we are pointing to, with the result that their reported interaction effects may well be biased.

Covariate models

A covariate and an independent variable may be correlated because of three possible causal processes: the independent variable causes the covariate, the covariate causes the independent variable, and they are spuriously related because some third variable or variables cause them both. Since a covariate is a measured, rather than manipulated variable, the second and third processes can rarely be teased apart. However, in some cases the first process can be distinguished from the other two. For instance, suppose the independent variable was manipulated rather than measured, with participants randomly assigned to its levels. In this case, any correlation between it and a subsequently measured covariate must be due to the causal effect of the independent variable on the covariate.

When a potential covariate is caused by an independent variables, then the inclusion of the covariate in the model, and the resulting adjustment function, is routinely labeled a mediational analysis, in which it is thought that the covariate mediates the effect of the independent variable on the dependent one. One includes the covariate to examine whether the adjusted effect of the independent variable is reduced once the mediator is controlled (Baron & Kenny, 1986; Judd & Kenny, 1981).

A consideration of this case thus points us to a situation where the issue we have raised may well be a problem even when both independent variables, whose interaction is of theoretical interest, have been experimentally manipulated. Suppose, for instance that both X_{1i} and X_{2i} have been experimentally manipulated and it is their interaction that is of theoretical interest. And suppose further that C_i is affected by X_{1i}, as a potential mediator. Then it will be the case that the $X_{1i}X_{2i}$ interaction will be confounded with the C_iX_{2i} interaction.

More formally, when C_i is measured after the independent variables have been manipulated, then one can specify a model for their effects on the covariate

$$C_i = \beta_{41}X_{1i} + \beta_{42}X_{2i} + \beta_{43}X_{1i}X_{2i} + \varepsilon_{4i}. \tag{4}$$

In the present context, we assume that only X_{1i} affects C_i (i.e., that both β_{42} and β_{43} equal zero). If we take the correct model for Y_i given earlier

$$Y_i = \beta_{11}X_{1i} + \beta_{12}X_{2i} + \beta_{13}C_i + \beta_{14}X_{1i}X_{2i} + \beta_{15}C_iX_{2i} + \varepsilon_{1i} \tag{1}$$

and substitute first for C_i (according to Eq. (4)) in the covariate interaction, we get

$$Y_i = \beta_{11}X_{1i} + \beta_{12}X_{2i} + \beta_{13}C_i + (\beta_{14} + \beta_{15}\beta_{41})X_{1i}X_{2i} + \varepsilon_{5i}. \tag{5}$$

This shows that the estimate of the interaction under the correct model β_{14} will be the same as the one found in the incomplete model β_{24} whenever either β_{15} or β_{41} is equal to 0. Not surprisingly, this conclusion is the same as the one derived earlier because it is the case that β_{41} equals $\sigma_{C,X_1}/\sigma_{X_1^2}$, assuming that both β_{42} and β_{43} equal zero.

If we now substitute for the remaining C_i in Eq. (5), we get

$$Y_i = (\beta_{11} + \beta_{13}\beta_{41})X_{1i} + \beta_{12}X_{2i} + (\beta_{14} + \beta_{15}\beta_{41})X_{1i}X_{2i} + \varepsilon_{6i}. \tag{6}$$

This sequential substitution nicely illustrates that, under the current assumptions we have made about the effect of the independent variables on C_i, the inclusion of the covariate interaction in the analytic model affects the interaction between the independent variables (see Eq. (5)) whereas the inclusion of the covariate only affects the coefficient associated with X_{1i} (see Eq. (6)).

Thinking of the covariate in this manner, as a potential mediator, allows the specification of different mediational cases that may be of interest. For instance, it may be the case that it is the interaction between X_{1i} and X_{2i} that affects C_i (i.e., that both β_{41} and β_{42} equal zero but β_{43} does not). In this case, C_i can be seen as a possible candidate to mediate the effect of the interaction of the two independent variables on the dependent variable. If this is the true model for the covariate (i.e., mediator), then it can be shown that an analytic model that only adjusts for the covariate is unbiased.

These two cases that we have just considered, where C_i is a mediator for the effects of the manipulated independent variables, define the alternative cases of what Baron and Kenny (1986) and Wegener and Fabrigar (2000) call "mediated moderation." In the first case, one independent variable, which interacts with a second, affects the covariate. One thus needs to control for the covariate's interaction with the second manipulated variable. In the second case, the interaction between the two independent variables affects the covariate. One then needs to control only for the covariate.

The lesson is that in the case of a covariate as a mediator, one needs to think about the true model of the effects of the independent variables on the mediator. Estimating a model for the mediator can only inform this process. Even outside of the mediational context, researchers should think about how the covariate relates to their independent variables and, in light of that, the specific adjustments that are accomplished by the inclusion of the covariate or covariate interactions in their analyses.

Conclusions

Whenever one's attention is focused on the interaction between two independent variables, an interesting situation emerges when one also wishes to control for some covariate(s) known to correlate with one of the independent variables. This state of affairs commonly arises when the design includes a measured independent variable that is crossed with a manipulated independent variable. The covariate is then likely to correlate with the measured variable and the goal is to ascertain the interaction between the independent variables adjusting for the covariate. It can also occur when the study involves two manipulated independent variables and the researcher is concerned about a covariate that is affected by one of them and may be a potential mediator.

The standard strategy is simply to include the covariate in the model. As we have shown analytically, this strategy leads to biased estimates of the effect of the interaction between the two independent variables unless one also controls for the interaction between the covariate and the independent variable with which it is uncorrelated.

Our simulations illustrated the harmful consequences of model misspecification. We systematically varied the two factors that our derivation showed to affect the magnitude of the bias: the relation between the covariate and the independent variable with which it is correlated, on the one hand, and the true effect of the covariate interaction on the other. Whenever these factors departed from 0, the inclusion of the covariate in the analytic model without the covariate interaction resulted in biased estimates and a substantial increase in Type I errors. An interesting case emerged when the power to detect the true effect of the covariate interaction was low. In such a case, one might find the covariate interaction to be non-significant and omit it from the final model. Nevertheless, our simulations show that even in this case, substantial bias can result. In general, we would recommend including the covariate interaction in the model, regardless of its significance.

In spite of the potential seriousness of the problem that we have pointed to, we found no mention of this issue in the classic manuals on experimental design. We did find a journal article (Hull et al., 1992) that addressed the issue explicitly. However, perhaps because these authors limited their scope to research in personality psychology, the larger social psychological audience seems to have overlooked the problem altogether. A survey of four major outlets of our discipline revealed that in none of the situations that called for the inclusion of the covariate interaction was the issue addressed.

The experimental design books do discuss the homogeneity of regression assumption in the analysis of covariance, a somewhat related issue. This assumption requires that the relationship between the covariate and the dependent variable be the same across the different cells of an experimental design. In the context of a design involving two independent variables and one covariate, this assumption would be tested by the comparison of two models. A full model includes the covariate, the main effects of the two independent variables, all possible two-way interactions, and the triple interaction. This model would be compared with a restricted one that omits all of the covariate interactions (i.e., the ANCOVA model). If one rejects the restricted model in favor of the full one, then that means that the covariate interacts with at least one of the independent variables that define that design, thereby making interpretation of the adjusted means problematic.

The issue we have raised and the assumption of homogeneity of regression in ANCOVA are both focused on covariate by independent variable interactions. However, the problem we emphasize here is more focused since we argue that the test of an interaction between a measured and a manipulated independent variable will be biased unless the interaction between the covariate and the manipulated independent variable is included. This is only one of the terms that would be added and tested, with multiple degrees of freedom, in examining the homogeneity of regression assumption. The difference between the two issues is perhaps clearest when one considers a design with only a single independent variable: the homogeneity of regression assumption still is important whereas the issue we are concerned about is irrelevant.

A final issue should be mentioned. Throughout, we have acted as if all our variables were measured without error. Clearly, this is never the case. Our derivations would be considerably more complicated had we assumed measurement error. Our major conclusions, however, hold regardless of the presence or absence of errors of measurement.[3]

Does the present warning mean that some findings widely accepted in the literature, such as the stereotype threat results discussed in our introduction, ought to be reevaluated? Perhaps. On the other hand, it is also possible that the covariate by manipulated independent variable interaction (i.e., SAT by the diagnosticity interaction) in that work would have proven to have no effect, had it been included. In fact, it is conceivable (although unlikely) that the effect of the predicted ethnicity by diagnosticity interaction would actually have been larger had the researchers controlled for the SAT

[3] A reviewer of this paper suggested that since measurement error generally attenuates effects, the presence of error may lead to less bias in the misspecified model that omits the covariate interaction, since it is the errorful covariate that has been omitted rather than the perfectly measured one. However, in models with multiple predictors, measurement error may either attenuate or exaggerate parameter estimates. Thus, the relative magnitude of additional bias if an errorful covariate interaction is omitted is unclear.

by diagnosticity interaction. The point is that we cannot know for sure the direction of the bias in the test of the crucial interaction when only the covariate (SAT), but not its interaction with diagnosticity, has been included in the model. This is indeed unfortunate.

Appendix A

Given the assumptions we have made, one can derive the expected values of all variances and covariances of the variables (Aiken & West, 1991, pp. 177–182; see also Kenny, 1979; McClelland & Judd, 1993). These are given in Table A1. From this table and the models we defined, it follows that

$$\beta_{14} = \frac{\sigma_{Y,X_1X_2} - \frac{\sigma_{Y,CX_2}\sigma_{C,X_1}}{\sigma_C^2}}{\sigma_{X_1}^2 - \frac{\sigma_{C,X_1}^2}{\sigma_C^2}}.$$

In the misspecified model, the coefficient for the $X_{1i}X_{2i}$ interaction is given by

$$\beta_{24} = \frac{\sigma_{Y,X_1X_2}}{\sigma_{X_1}^2}.$$

From these two expressions and the results in Table A1, one can show that

$$\beta_{24} = \beta_{14} + \beta_{15}\frac{\sigma_{C,X_1}}{\sigma_{X_1}^2}.$$

Table A1. Variances and covariances in the true model

	X_{1i}	X_{2i}	C_i	$X_{1i}X_{2i}$	C_iX_{2i}	Y_i
X_{1i}	$\sigma_{X_1}^2$					
X_{2i}	0	1				
C_i	σ_{C,X_1}	0	σ_C^2			
$X_{1i}X_{2i}$	0	0	0	$\sigma_{X_1}^2$		
C_iX_{2i}	0	0	0	σ_{C,X_1}	σ_C^2	
Y_i	$\beta_{11}\sigma_{X_1}^2 + \beta_{13}\sigma_{C,X_1}$	β_{12}	$\beta_{11}\sigma_{C,X_1} + \beta_{13}\sigma_C^2$	$\beta_{14}\sigma_{X_1}^2 + \beta_{15}\sigma_{C,X_1}$	$\beta_{14}\sigma_{C,X_1} + \beta_{15}\sigma_C^2$	$\beta_{11}^2\sigma_{X_1}^2 + \beta_{12}^2 + \beta_{13}^2\sigma_C^2 + \beta_{14}^2\sigma_{X_1}^2 + \beta_{15}^2\sigma_C^2 + 2\beta_{11}\beta_{13}\sigma_{C,X_1} + 2\beta_{14}\beta_{15}\sigma_{C,X_1} + \sigma_{\varepsilon_1}^2$

References

Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions*. Newbury Park, CA: Sage.

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*, 1173–1182.

Hull, J. G., Tedlie, J. C., & Lehn, D. A. (1992). Moderator variables in personality research: The problem of controlling for plausible alternatives. *Personality and Social Psychology Bulletin, 18*, 115–117.

Judd, C. M., & Kenny, D. A. (1981). Process analysis: Estimating mediation in treatment evaluation. *Evaluation Review, 5*, 602–619.

Judd, C. M., & McClelland, G. H. (1989). *Data analysis: A model comparison approach*. San Diego, CA: Harcourt, Brace and Jovanovich.

Kenny, D. A. (1979). *Correlation and causality*. New York, NY: Wiley.

Keppel, G., & Zedeck, S. (1989). *Data analysis for research designs: Analysis of variance and multiple regression/correlation approaches*. New York, NY: Freeman.

Kirk, R. E. (1995). *Experimental design: Procedures for the behavioral sciences* (3rd ed.). Pacific Grove, CA: Brooks/Cole.

Maxwell, S. E., & Delaney, H. D. (1990). *Designing experiments and analyzing data. A model comparison perspective*. Belmont, CA: Wadsworth.

McClelland, G. H., & Judd, C. M. (1993). Statistical difficulties of detecting interactions and moderator effects. *Psychological Bulletin, 114*, 376–390.

Myers, J. L., & Well, A. D. (1991). *Research design and statistical analysis*. New York, NY: Harper Collins.

Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology, 69*, 797–811.

Wegener, D., & Fabrigar, L. (2000). Analysis and design for nonexperimental data: Addressing causal and noncausal hypotheses. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social and personality psychology* (pp. 412–450). New York, NY: Cambridge University Press.

Winer, B. J., Brown, D. R., & Michels, K. M. (1991). *Statistical principles in experimental design* (3rd ed.). New York, NY: McGraw-Hill.