Powerful Moderator Variables in Behavioral Science? Don't Bet on Them

Ryne A. Sherman¹ & Harold Pashler²

¹Hogan Assessments

²University of California, San Diego

Authors' Note

We are grateful to Robert Hogan for granting us access to the Hogan Data Archive used in this research. We are grateful to David C. Funder for granting us access to the RAP-I, RAP-II, and RSP data sets and for feedback on a previous draft of this article. All errors and omissions remain our own. The RAP-I and RAP-II data sets used in this research were supported by NIMH grants R01-MH42427 and R01-MH40808 while the RSP data set was supported by NSF grant NCS-064223 all to David C. Funder. Any opinions, findings, conclusions, or recommendations expressed in this article are those of the individual researchers and do not necessarily reflect the views of the National Institutes of Health or the National Science Foundation. All statistical analyses were conducted using R (R Core, Team, 2016).

Correspondence regarding this article may be sent to Ryne A. Sherman, Hogan Assessments, 11 S Greenwood, Tulsa, OK 74120. Email: rsherma@hoganassessments.com

Abstract

The current upsurge of interest in research replicability (and the exposure of many failures of reproducibility) has led to a much discussion about the possible role of statistical moderation (i.e., variable $\times$ variable interactions) in behavioral and social science. These interactions are so widespread and powerful, it is often argued, that we should hardly be surprised when attempts to reproduce important findings frequently lead to failure. Prior literature provides little empirical evidence about how common powerful moderation is. Using five large-scale behavioral research datasets we sought to shed light on the issue. The data reflected several thousand people engaging in a variety of behaviors over considerable periods of time, and included hundreds of demographic and psychological independent variables (IVs). For each outcome variable, we measured the interaction of every pair of IVs. Many IVs had sizable main effects on behavior, but interactions were usually very small in magnitude. While there is no doubt that interactions can occur in behavioral science contexts, the priors revealed here suggest that this should be postulated only as a last, not a first resort.

Reproducibility Failures and Moderator Effects

The past few years have seen a great upsurge of interest in reproducibility in many empirical research fields including psychology. Many studies have reported failures to replicate well-known psychological findings (e.g., Open Science Collaboration, 2015; Klein et al., 2015). The interpretation of these failures has generated much controversy. On the one hand, critics of prevailing practices have argued that if many studies have low statistical power (as appears to be the case in many fields, e.g., Button et al., 2013; Rossi, 1990) *and* if a large fraction of the potential effects investigators set out to find do not exist (with any appreciable minimum effect size), then it is a matter of mathematical necessity that Type I errors must be very frequent in the literature (Ioannidis, 2005). Worse yet, if flexibility in data collection/analysis and questionable research practices are as common as it would appear they are (see Simmons, Nelson, & Simonsohn, 2011; John, Loewenstein, & Prelec, 2012), the proportion of unreal findings could be still greater yet. Recent large-scale examinations of replicability have produced dismal findings which appear to validate these concerns (e.g., Open Science Collaboration, 2015).

On the other hand, some critics of replicability concerns have suggested that failures to replicate reported findings may merely reflect the ubiquity and strength of moderation effects, rather than widespread error in the literature. For example, Feldman-Barrett (2015) said “If the studies were well designed and executed, it is more likely that the phenomenon from Study A is true only under certain conditions” [rather than simply being untrue]. Along similar lines, Dijksterhuis, van Knippenberg, and Holland (2014) said “most psychological phenomena tend to have many moderators (including factors that cannot always be easily controlled in a lab)” (p. 202) and Stroebe and Strack (2014) said “replication failures [frequently occur because]

important moderators were not spelled out by a theory” (p. 66). Similar remarks can be found in many papers contributing to the current debates on replication.

As just noted, statistical moderation, commonly referred to in psychology as *interaction between variables*, has often been assumed to be ubiquitous and powerful across the landscape of behavioral science. Indeed, one resolution to the person-situation debate, a major preoccupation of social psychology over several decades, was the assumption of powerful “person $\times$ situation interactions” (Kenrick & Dantchik, 1983; Kenrick & Funder, 1988; Magnusson & Endler, 1977). The person $\times$ situation interaction assumption largely rests on deductive grounds: if person and situation variables only explain a relatively small portion of behavior, the rest must be explained by the interaction. Similarly, and as was just noted, the interaction assumption has been expressed with particular enthusiasm as controversies have erupted in recent years about frequent replication failures in behavioral sciences (a controversy that has sometimes become acrimonious; Meyer & Chabris, 2014; Yong, 2012).

However, such enthusiasm for moderators is starkly at odds with decades of arguments and experience that interaction effects are largely weak (see Chaplin, 1997 for review). In perhaps one of the earliest discussion of the effect size of moderators, Ghiselli (1963) stated “It is quite possible that the time and effort required to develop moderators might be more fruitfully spent in seeking improvements in reliability and validity of the sort that follow from classic psychometric theory” (p. 86). In a discussion of making decisions based on linear regression, Dawes and Corrigan (1974) wrote, “The whole trick is to know what variables to look at then to know how to add” (p. 105). Perhaps most directly of all, Cronbach (1975) wrote, “Once we attend to interactions, we enter a hall of mirrors that extends to infinity. However far we carry our analysis-to third order or fifth order or any other-untested interactions of a still higher order

can be envisioned. When I say something like that, some colleague is likely to reply: ‘In my experience, interaction effects are not large.’” (p. 119).

Realistic Priors for Moderation in Behavioral Data

Perhaps the most critical part about Cronbach’s statement is the phrase “in my experience.” Indeed, the scientific literature offers very little in terms of empirical investigations as to the typical size, base rate, or prior probabilities, for moderators. A serious problem for comparing these conflicting views of replication failures is the lack of any empirical basis for deciding what *a priori* expectations we should have for the size of moderation effects. Very few datasets offer any clues. The recent Many Labs Project harnessed 36 labs to work simultaneously in an effort to reproduce the results of 13 well-known studies from the fields of social psychology and judgment and decision making (Klein et al., 2015). The rate of replication of original studies was reasonably good overall but varied across fields: a number of judgment and decision making findings were reproduced with effect sizes about as large as the original findings, while two social priming effects showed essentially zero effect size. Of greatest interest for the point discussed here, however, is the fact that between-lab differences in results proved to be quite minimal and consistent in magnitude with the possibility they reflected nothing more than sampling error (i.e., datasets reflecting a given condition as measured in the various labs could all have been drawn from the very same population). While such a result might be indicative of a lack of cross-lab (or so-called “hidden”) moderators, the Many Labs Project was not designed to estimate the size and frequency of moderators writ large.

The best—and perhaps only—direct attempt to estimate the frequency and size of moderators in the behavioral sciences was conducted over 25 years ago by Chaplin (1991). On the basis of a variety of analyses of $N = 112$ subjects, Chaplin concluded that moderators were

rare and frequently small in effect size. In the present work, we go even further than Chaplin by testing for all possible pairwise interactions in five large behavioral datasets to which we have access. In doing so, we not only examine the frequency of interactions reaching statistical significance, but also the typical effect size associated with such interactions. This is critical because as McClelland and Judd (1993) showed, interactions may only be statistically significant and large in situations where variables making up the interaction are non-normally distributed. Of course, in naturalistic settings (i.e., real life outside the laboratory), variables are typically normally distributed. The data sets upon which we draw come from studies of daily life or, in the case of two laboratory studies, settings intended to be as realistic as possible.

In the present investigation, we set out to shed light on the a priori distribution of interaction effect sizes by drawing on archival data from both laboratory and naturalistic settings. The five studies reported below involve parallel analyses quantifying the distribution of moderator effect sizes for all the potential person(ality) by situation interactions that can be assessed in each of five large datasets. Because sex is often analyzed as a potential moderator, we also examine all possible personality $\times$ sex and situation $\times$ sex effects in each of these studies. In the next section we, in turn, describe the basic study design of each of the five studies. After that, we present the estimated distributions of all the moderation effect sizes, using several different measurement approaches.

The approach here has a major advantage over previous efforts at estimation of reasonable priors for effect sizes. These previous efforts have utilized published data to estimate the priors (chiefly focusing on main effects). The problem with such efforts, of course, is publication bias: the published effects are the ones that authors and editors viewed as most interesting (often using statistical significance as an indicator of interestingness), and thus they

are guaranteed to overestimate effect sizes. By contrast, the five datasets used here are used in their entirety, and these datasets did not draw our attention due to the size of the interaction terms. Later we will discuss limitations in the breadth and coverage of the datasets we are using, but readers should also notice that lack of susceptibility to publication bias is a major virtue.¹

Finally, the issue examined here has an obvious connection to recent debates about replicability in the social sciences. However, the findings also have broader significance for some less well-known, but still important, questions about proper methodology in behavioral research (for example, the question of the relative merits of convenience samples versus probability samples, and the causes for different outcomes often observed in field experiments and laboratory studies). These connections are explored in the General Discussion.[RS1]

The results indicate that moderators proved to arise far more rarely than the interaction assumption would leave researchers to believe, with even smaller effect size than suggested by Chaplin. As we will discuss later, the results we present here are in close agreement with some findings from the recent “Many Labs” project (Klein, Ratliff, Vianello, Adams, Bahník, Bernstein, ... & Woodzicka, J. A., 2015).[RS2]

Study 1 Design

The Riverside Accuracy Project – Phase I (RAP-I) dataset was collected from October 1990 through June of 1993 and was designed to identify and understand mechanisms involved in accurate personality judgment. It includes a large battery of self-reported personality measures, peer reports of personality, real time mood data collected using beepers, and behavioral data

from as many as seven different situations (see Furr & Funder, 1998 for full details). Of import here, participants engaged in three video recorded 5 min situations with another participant of the opposite sex with whom they were not previously acquainted. In the first (unstructured) situation, participants were told they could talk about whatever they wanted for the next five minutes. In the second (cooperative) situation, participants were given a set of Tinker Toys and instructed to work together to build a particular model. In the third (competitive) situation, participants played the memory game SIMON; the winner of three out of five games received an extra dollar.

For each of these three situations, four independent coders viewed the videotaped interaction and completed the Riverside Behavioral Q-sort (RBQ version 2; Funder, Furr, & Colvin, 2000) to quantify each participant's behavior. This version of the RBQ contained 64 behavioral outcome measures (e.g., "Dominates interaction," "Expresses agreement frequently"). No coder rated a participant in more than one of the three situations, to avoid potential rater $\times$ target biases. Because this archival dataset contains only composite scores, we are unable to recalculate the reliability of these behavioral ratings. However, a previous publication reported the median item reliability to be .55 (Furr & Funder, 1998).

At a later session participants completed self-reports of the California Adult Q-set (CAQ: Block, 1978; as modified for non-professional use by Bem and Funder, 1978) using the Q-sort procedure. The CAQ contains 100 diverse personality characteristics (e.g., "Is genuinely dependable and responsible"; "Is guileful and deceitful, manipulative, opportunistic") and is designed to broadly cover the personality domain, making it well-suited for examining potential person $\times$ situation interactions.

Study 1 Results²

Because each subject's behavior was measured three times on each of 64 items, mixed effects modeling was employed to account for the nesting. We began by constructing a model to predict each of the 64 behaviors from each of the 100 personality characteristics and two dummy-coded variables corresponding to the situation condition. We estimated both fixed and random intercept terms, thereby allowing individuals to differ in their average amount of each behavior across the three situations. All slopes were estimated as fixed effects only. We used the method described by Nakagawa and Schielzeth (2013) to estimate the overall model fit for the fixed effects of each model. Across all 6,400 models the average multiple R was .326 ($SD = .261$, median = .244). Next, we re-estimated each of these models this time including the two personality characteristics $\times$ situation dummy-code interaction terms. The average multiple R across all 6,400 models including the interaction effects was .333 ($SD = .258$, median = .251).

We also computed chi-square difference tests between those models without and with interaction effects included as a test of the statistical significance of the interaction. The total number of statistically significant ($p < .05$) chi-square tests was 458 out of 6,400, or about 7%. Of course, 5% is the number one would expect by chance alone under the assumption that each test is independent. The independence assumption is not valid here because the items of both the RBQ and CAQ are sometimes correlated. To account for this, we adapted a randomization procedure devised by Sherman and Funder (2009) to empirically estimate the distribution of statistically significant interaction effects under the true null of no effect.³ Based on 100 randomizations, the mean number of statistically significant findings expected by chance was 407.68 ($SE = 206.36$), which is greater than the nominal figure of 320. Indeed, 27% of the simulations yielded a number of statistically significant interaction effects greater than the observed figure of 458. To summarize, then, it would appear that there is a reasonable likelihood

that *virtually all* of the 458 observed statistically significant interaction effects reflect nothing more than random sampling error.

Increase in model fit R_{Δ}

To better quantify the incremental gain in model fit by adding the interaction term to a model that previously lacked one, we computed the R_{Δ} as the square root of the difference between the models ($R_{\Delta} = \sqrt{R_{mod2}^2 - R_{mod1}^2}$; Cohen, Cohen, West, & Aiken, 2003).⁴ The average R_{Δ} across all 6,400 models was .044. Figure 1 displays the full distribution of R_{Δ} in this dataset. In contrast to the interaction assumption, the R_{Δ} are quite small on average with the majority falling below .05.

Figure 1. Distribution of R_{Δ} in model fit (R) by adding a person $\times$ situation interaction term to a model with main effects for person and situation.

Sex as a Moderator

The analyses just described were repeated with sex as the moderating variable of the person-behavior relationship and sex as the moderating variable of the situation-behavior relationship. For the person-behavior links, 6,400 mixed effects regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .097 ($SD = .058$; median = .09) for the main effects only. With the sex $\times$ personality interaction term included, the average model R was .110 ($SD = .056$, median = .102). 364 four out of the 6,400 models (5.7%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .040$ ($SD = .032$, median = .033). Thus, personality $\times$ sex was both rare and weak.

For the situation-behavior links, 64 mixed effects regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .333 ($SD = .262$; median = .249) for the main effects of situation and sex only. With the situation $\times$ sex interaction term included, the average model R was .340 ($SD = .258$, median = .252). Six out of the 64 models (9.4%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .043$ ($SD = .029$, median = .039). Thus, situation $\times$ sex moderation was both rare and weak.

Study 2 Design

The Riverside Accuracy Project – Phase II (RAP-II) dataset was collected from the spring of 2000 through the spring of 2002 and was also designed to identify and understand the mechanisms involved in accurate personality judgment. This dataset includes a large battery of

self-reported personality measures, peer-reports of personality, clinician reports of personality, a life history interview with a trained clinician (see Letzring, Wells, & Funder, 2006 for full details). Of import here, participants also engaged in one of four randomly assigned video recorded situations with two other participants with whom they were not previously acquainted.

The first video recorded situation (trivia condition) lasted 50 minutes and participants were instructed to work together to complete a set of trivia questions. The second situation (short unstructured condition) lasted for 50 minutes and participants were instructed to talk about whatever they would like. The third situation (get to know you condition) lasted 50 minutes and participants were instructed to get to know each other as well as possible. The fourth situation (long unstructured) condition lasted for three hours and participants were instructed to talk about whatever they would like. Participants assigned to this condition were given a short break midway through the session.

For each of these four situations, four independent coders viewed the videotaped interaction and completed the RBQ (Funder, Furr, & Colvin, 2000), using the same Q-sorting procedure described previously, to quantify each participant's behavior. As a means of quality control, codes were compared with those from other raters of the same participant using intra-class correlations. A coding was retained if it agreed at least .30 with one other coding or at least .20 with two others. When these criteria were not met the coding was discarded and an additional one was completed. Behavioral composites were formed by averaging the ratings on the 64 RBQ items across the four codes. The average ICC for a single rater across the 64 items was .31 ($SD = .15$) with an average reliability for the composite of .61 ($SD = .17$).

Participants completed as many as three take home packets of personality surveys at different times throughout the course of their participation. The first packet included the CAQ

and participants rated themselves on each of the 100 items using a 1 (*extremely uncharacteristic*) to 9 (*extremely characteristic*) Likert-type rating format.

Study 2 Results

Because a between-subjects design was employed, ordinary linear regression could be used to examine the question of interest. We began by constructing a model to predict each of the 64 behaviors from each of the 100 personality characteristics and three dummy-coded variables corresponding to the situation conditions. We used the model R as the indicator of model fit. Across all 6,400 models the average multiple R was .290 ($SD = .150$, median = .253). Next, we re-estimated each of these models this time including the three personality characteristics $\times$ situation dummy-code interaction terms. The average multiple R across all 6,400 models including the interaction effects was .327 ($SD = .142$, median = .296), a small gain in model fit, but a bit larger than that found in Study 1.

We also computed chi-square difference tests between those models without and with interaction effects included. The total number of statistically significant ($p < .05$) chi-square tests was 350 out of 6,400, or about 5.5%. Of course, 5% is the number one would expect by chance alone under the assumption that each test is independent. As before, the independence assumption is not valid here and to account for this, we adapted a randomization procedure devised by Sherman and Funder (2009) to empirically estimate the distribution of statistically significant interaction effects under the true null of no effect. Based on 100 simulations, the mean number of statistically significant findings expected by chance was 319.92 ($SE = 43.54$), which is nearly identical to the nominal figure of 320. Still, 21% of the simulations yielded a number of statistically significant interaction effects greater than the observed figure of 350. To

summarize, then, it would appear that there is a reasonable likelihood that virtually *all* of the 350 observed statistically significant interaction effects can be explained by sampling error.

Proportional Increase in model fit (R)

To visualize the distribution of the amount of gains achieved by adding an interaction term in a model that previously lacked one, we plotted the improvement in model fits (R) by adding the interaction term ($R_{\Delta} = \sqrt{R_{mod2}^2 - R_{mod1}^2}$; Cohen, Cohen, West, & Aiken, 2003) for all 6,400 models. This increase in model fit has a lower bound of zero and Figure 2 shows the distribution for Study 2. As can be seen, the distribution of multiple R_{Δ} in Study 2 was shifted a bit further right (i.e., larger on average) than the distribution in Study 1. Still, the effects were nowhere near as large as one might hope on the basis of the interaction assumption.

Figure 2. Distribution of R_{Δ} in model fit (R) by adding a person $\times$ situation interaction term to a model with main effects for person and situation.

Sex as a Moderator

We once again repeated the same analyses just described this time with sex as the moderating variable of the person-behavior relationship and sex as the moderating variable of the situation-behavior relationship. For the person-behavior links, 6,400 regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .146 ($SD = .077$; median = .138) for the main effects only. With the sex $\times$ personality interaction term included, the average model R was .171 ($SD = .076$, median = .163). Three-

hundred seventy-two out of the 6,400 models (5.8%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .069$ ($SD = .052$, median = .059). Thus, personality $\times$ sex moderators were both rare and weak.

For the situation-behavior links, 64 regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .301 ($SD = .148$; median = .256) for the main effects of situation and sex only. With the situation $\times$ sex interaction term included, the average model R was .344 ($SD = .133$, median = .318). Five out of the 64 models (7.8%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .14$ ($SD = .056$, median = .144). Thus, situation $\times$ sex moderators were both rare and relatively weak.

Study 3 Design

The Riverside Situation Project dataset was collected from the fall of 2007 through the spring of 2009 and was designed to identify and measure the psychological characteristics of situations in real-world settings. The dataset includes a large battery of self-reported personality measures, peer reports of personality, self-reports of properties of four different situations participants experienced in their daily lives, self-reports of behavior in those situations, and self-reports of how participants felt in those situations. Of import here, the dataset includes self-reported personality on the 100-item CAQ from 221 undergraduate participants and self-reported behavior from four situations experienced by these participants.

During an initial visit to the laboratory participants completed the aforementioned CAQ. In the four weeks following the initial visit, each participant returned four more times (at least 48 hours after any previous visit). During these visits the participant was given a 3 $\times$ 5 (in.) index card and asked to write down what he or she was doing at one of four pre-specified times within

the previous 24 hours (10am, 2pm, 5pm, or 9pm). A Latin-square design was used such that approximately 1/4th of the participants completed the study using each of the following time sequences across the four measurement occasions: 10am-2pm-5pm-9pm; 2pm-5pm-9pm-10am; 5pm-9pm-10am-2pm; 9pm-10am-2pm-5pm. Participants were instructed to specify only one situation, and that if they were sleeping they should write down what they were doing right before they went to sleep or right after they woke up. The situations described by participants were typical of an undergraduate population (e.g. going to class, eating lunch with friends, at work; more detail on the methods used is provided by Sherman and colleagues [2010]).

Immediately following this, participants rated the situation they just described using the Riverside Situation Q-sort (RSQ version 2.0: Wagerman & Funder, 2009). Later in the session participants completed an RBQ (version 3.0; Furr, Wagerman, & Funder, 2010) describing their behavior in that situation. This version of the RBQ contains 67 items and is designed to measure behavior in more naturalistic (i.e., non-laboratory) situations.

Study 3 Results

Because each subject's behavior was measured four times on each of the 67 items, we employed mixed effects modeling to account for the nesting. We began by predicting each of the 67 behaviors from each of the 100 personality characteristics and 81 situation characteristics, a total of 542,700 mixed effects regression models. As before, in the first step we estimated only main effects of a single personality and situation variable. Intercepts and situation slopes were estimated as both fixed and random effects (to allow individuals do differ in means and slopes), and personality was only estimated as a fixed effect. We again used the method described by Nakagawa and Schielzeth (2013) to estimate the overall model fit for the fixed effects of each model. Across all 542,700 models the average multiple *R* was .091 (*SD* = .061, median = .079).

Next, we re-estimated each of these models this time including the personality $\times$ situation characteristic interaction term. The average multiple R across all 542,700 models including the interaction effects was .098 ($SD = .059$, median = .086).

As before, we also computed chi-square difference tests between those models without and with interaction effects included. The total number of statistically significant ($p < .05$) chi-square tests was 27,698 out of 542,700, or about 5.1%. Of course, 5% is the number one would expect by chance alone under the assumption that each test is independent. Although the independence assumption is not valid here, we did not account for this (as before) for two reasons. First, computing these statistics for 542,700 mixed effects models was computationally intensive taking over 16 hours to complete on a standard university desktop computer. Repeating this procedure 100 times via a randomization test would clearly take far longer. Second, because the total number of statistically significant findings is so close to the nominally expected amount (under the invalid independence assumption) it is virtually impossible to reasonably conclude that there is something more than chance effects here. Thus, it appears that there is a reasonable likelihood that virtually *all* of the 27,698 observed statistically significant interaction effects can be explained by random sampling error alone.

Following the same method as in Study 1, Figure 3 shows the distribution of increases in model fit (R) by adding a person $\times$ situation interaction term to a model with main effects for person and situation.⁵ As can be seen, this distribution leans strongly towards zero with very few R_{Δ} exceeding .05. This is again in sharp contrast to what one would expect on the basis of the interaction assumption.

Figure 3. Distribution of R_{Δ} in model fit (R) by adding a person $\times$ situation interaction term to a model with main effects for person and situation.

Sex as a Moderator

As in the previous two studies, we repeated the same analyses just described this time with sex as the moderating variable of the person-behavior relationship and sex as the moderating variable of the situation-behavior relationship. For the person-behavior links, 6,700 mixed effects regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .063 ($SD = .034$; median = .058) for the main effects only. With the sex $\times$ personality interaction term included, the average model R was .077 ($SD =$

.034, median = .073). Three-hundred sixty-two out of the 6,700 models (5.4%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .035$ ($SD = .027$, median = .030). Thus, personality $\times$ sex moderators were both rare and weak.

For the situation-behavior links, 5,427 mixed effects regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .085 ($SD = .06$; median = .071) for the main effects of situation and sex only. With the situation $\times$ sex interaction term included, the average model R was .093 ($SD = .058$, median = .080). Two-hundred eighty-two out of the 5,427 models (5.2%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .027$ ($SD = .021$, median = .023). Thus, situation $\times$ sex moderators were both rare and weak.

Study 4 Design

The Florida Atlantic Experience Sampling dataset was collected in the fall of 2013. It was designed to measure the simultaneous impact of participants' personality and situational characteristics on real-time behavior and emotion. The study consisted of two phases. In the first phase, participants completed a brief (~5 min) video recorded interview and a number of personality measures including the HEXACO-60 (Ashton & Lee, 2009) and the Subjective Happiness Scale (Lyubomirsky & Lepper, 1999). In the second phase, participants were sent a text message eight times per day over the next seven days. The text message contained a link to a personalized survey. The first page of the survey asked participants to describe their current situation using an 8-item measure of the DIAMONDS situation characteristics (Rauthmann & Sherman, 2016; see also Rauthmann et al., 2014) using a 1 (*extremely uncharacteristic*) to 7 (*extremely characteristic*) Likert-type scale. The second page of the survey asked participants to describe their current behavior / emotion 10 bipolar, seven-point rating scales pertaining to

different personality characteristics (e.g., outgoing, sociable—reserved, quiet). More details are provided by Sherman, Rauthmann, Brown, Serfass, & Jones, 2015.

Study 4 Results

In modeling these data, we used the 7 personality dimensions assessed in the first phase of the study and crossed these with the 8 situation measures recorded by the subject in the second phase, to predict one of the 10 behavioral/emotional outcome variables. This yielded a total of 560 interaction terms. Consistent with Sherman et al. (2015), we within-person centered situation variables and included each person's mean score as a level-2 predictor. We estimated fixed and random effects for the intercept and the within-person situation variable and fixed effects for personality dimensions and situation means. We again used the method described by Nakagawa and Schielzeth (2013) to estimate the overall model fit for the fixed effects of each model. Across all 560 models the average multiple R was .221 ($SD = .113$, median = .196). Next, we re-estimated each of these models this time including the personality $\times$ situation characteristic interaction term. The average multiple R across all 560 models including the interaction effects was .222 ($SD = .114$, median = .200).

As before, we also computed chi-square difference tests between those models without and with interaction effects included. The total number of statistically significant ($p < .05$) fits was 46 out of 560, or about 8.2%. Of course, 5% is the number one would expect by chance alone under the assumption that each test is independent. Unfortunately, the structure of the data in this study do not permit a proper randomization test (as in the other studies) to determine the likelihood that these 46 observed statistically significant interactions could be due to chance. This is because the nested structure of the data requires randomization at both the between- and within-subject levels. This is possible when all subjects have an equal number of repeated

assessments (as in Study 1), but not in studies such as this one where subjects vary widely in their number of assessments. At issue is the inability to assign values from one subject who may have had 20 reports randomly to another subject who may have had 30 reports, and vice-versa. As a result, we cannot specify to degree to which the 8.2% of statistically significant effects may be due entirely to sampling error, but we do know that we would expect a minimum of 5% by chance alone.

Following the same method as in Study 1, we calculated the R_{Δ} in model fit by adding a person $\times$ situation interaction term to a model with main effects of person and situation.⁶ The average R_{Δ} was .019 and Figure 4 shows the full distribution of these values. In contrast with the interaction assumption, the vast majority of the interaction terms added very little in the way of predictive value.

Figure 4. Distribution of R_{Δ} in model fit (R) by adding a person $\times$ situation interaction term to a model with main effects for person and situation.

Sex as a Moderator

As in the previous studies, we repeated the same analyses just described this time with sex as the moderating variable of the person-behavior relationship and sex as the moderating variable of the situation-behavior relationship. For the person-behavior links, 70 mixed effects regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .133 ($SD = .066$; median = .125) for the main effects only. With the sex $\times$ personality interaction term included, the average model R was .142 ($SD = .068$,

median = .130). Seven out of the 70 models (10%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .04$ ($SD = .032$, median = .032). Thus, personality $\times$ sex moderators were both relatively rare and weak.

For the situation-behavior links, 80 mixed effects regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .191 ($SD = .122$; median = .145) for the main effects of situation and sex only. With the situation $\times$ sex interaction term included, the average model R was .198 ($SD = .118$, median = .177). Six out of the 80 models (7.5%) indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .034$ ($SD = .022$, median = .032). Thus, situation $\times$ sex moderators were both rare and weak.

Study 5 Design

The Hogan Data Archive contains personality assessment results from millions of job applicants for nearly every occupation in the United States. Of import here, the archive contains both personality—as measured by the Hogan Personality Inventory (HPI; Hogan & Hogan, 1995)—and job performance scores for 7,314 people from 60 different jobs in 7 different job families.⁷ Further, each of these jobs has been measured on the six RIASEC dimensions of occupations identified by Holland (1966). As such, these data can be used to examine person (HPI personality trait scores) by situation (RIASEC scores) interactions on job performance.

Study 5 Results

Because the data in the Hogan Archive come from many studies conducted over many years, different versions of the HPI have been used. Moreover, overall job performance ratings varied from study to study (e.g., some used binary scales, others used raw sales figures, etc.). In order to standardize HPI scores, we first converted all HPI scores to percentiles using the norms

for each instrument. Next, we converted these percentiles into Z-scores after trimming persons scoring in the zeroth or one-hundredth percentile to the first and ninety-ninth respectively (to avoid Z-scores of infinity). Lastly, we standardized job performance by Z-scoring within each study (job).

To model these data we used ordinary least squares regression for each combination of the seven personality characteristics measured by the HPI and the six dimensions of the RIASEC to predict overall performance, creating 42 models in total. We again used model *R* as our indicator of model fit. Across the 42 models the average multiple *R* was .056 (*SD* = .031, median = .051). Next, we re-estimated each of these models this time including the personality × RIASEC interaction term. The average multiple *R* across the 42 models including the interaction effect was .063 (*SD* = .029, median = .057).

As before, we also computed chi-square difference tests between those models without and with interaction effects included. The total number of statistically significant ($p < .05$) chi-square tests was 16 out of 42, or about 38.1%. Once again, 5% is the number one would expect by chance alone under the assumption that each test is independent. As before, the independence assumption is not valid here and to account for this, we adapted a randomization procedure devised by Sherman and Funder (2009) to empirically estimate the distribution of statistically significant interaction effects under the true null of no effect. Based on 100 simulations, the mean number of statistically significant findings expected by chance was 2.08 (*SE* = 2.11), which is nearly identical to the nominal figure of 2.1. None (0%) of the simulations yielded a number of statistically significant interaction effects greater than the observed figure of 16 suggesting that the observed interaction effects are unlikely to be due entirely to random sampling error.

Proportional Increase in model fit (R)

Following the procedure used in Study 2, we calculated the incremental improvement in model fit by adding the interaction term for each of these 42 models. The average incremental model fit was $R = .02$ and Figure 5 displays their distribution. As can be seen, and in contrast with the interaction assumption, the typical R_{Δ} was very small with very few extending beyond .05.

Figure 5. Distribution of R_{Δ} in model fit (R) by adding a person $\times$ situation interaction term to a model with main effects for person and situation.

Sex as a Moderator

As in the previous studies, we repeated the analyses just described this time with sex as the moderating variable of the person-performance relationship and sex as the moderating variable of the situation-performance relationship. For the person-performance links, 7 regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .054 ($SD = .039$; median = .044) for the main effects only. With the sex $\times$ personality interaction term included, the average model R was .056 ($SD = .038$, median = .050). Zero out of the 7 models indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .009$ ($SD = .007$, median = .009). Thus, personality $\times$ sex moderators of performance were nil and weak.

For the situation-behavior links, 6 regression models, with main effects only and then main effects plus interaction, were conducted. The average model R was .016 ($SD = .006$; median = .013) for the main effects of situation and sex only. With the situation $\times$ sex interaction term included, the average model R was .029 ($SD = .013$, median = .029). Three out of the 6 models indicated a statistically significant increase in model fit with the average increase being $R_{\Delta} = .022$ ($SD = .016$, median = .022). Thus, situation $\times$ sex moderators were common by statistical significance standards, but very weak.

General Discussion

The notion that much of human behavior can be explained by person $\times$ situation interactions is common among social-personality psychologists. Popular undergraduate textbooks often emphasize such interactions (e.g., Kenrick, Neuberg, & Cialdini, 2015). However, in the five studies described here, we found very little empirical evidence for such person $\times$ situation interactions. In studies 1-4, the interaction effects reaching conventional levels

of statistical significance were likely entirely due to random sampling error. More damning, in all five studies, the explanatory power added by the interaction was far lower than researchers might expect under the interaction assumption--and often proved trivial in magnitude. Table 1 summarizes these key findings from all five studies. While it may be tempting to look at the average R_{Δ} in Study 2 (.131) or the percentage of statistically significant interactions in Study 5 (38.1%) to hold out some hope for person $\times$ situation interactions, this would be folly. The results in Study 2, though somewhat larger on average, are virtually all Type I errors and should be ignored. The results of Study 5, on the other hand, are unlikely to be Type I errors. Thus, Study 5 provides the best evidence for the existence of person $\times$ situation interaction effects and simultaneously the best estimate of their added predictive value, which is only $R = .02$. Hardly the effect size one might expect given the interaction assumption.

Table 1. Summary of Personality $\times$ Situation Results

Study #	Sample	N	N _{obs}	N _{effects}	Avg. R		Avg. R	Avg. R
					ME	ME + Int	Incr	% Sig
1	RAP-I	169	490	6400	.326	.333	.044	7.0%
2	RAP-II	143	143	6400	.290	.327	.131	5.5%
3	RSP	204	809	542,700	.091	.098	.025	5.1%
4	FAU	186	7,298	560	.221	.222	.019	8.2%
5	Hogan	7,314	7,314	42	.056	.063	.020	38.1%

Note. ME = Models with main effects only. ME + Int = Models with main effects and interaction term(s) included. Incr = Incremental Model Fit after adding interaction term. Sig = Statistically Significant.

The results were equally dismal when we examined sex as a moderator of both personality and situation variables on behavior and workplace performance. As sex is often a go-to moderator variable (e.g., “I wonder if it is moderated by sex?”) for researchers, this finding may be even more surprising than the lack of personality $\times$ situation interactions. To be clear, we are not implying that sex had no relationship to behavior. Most studies examining sex or gender differences focus on main effects (e.g., Schmitt, Realo, Voracek, & Allik, 2008; Zell, Krizan, & Teeter, 2015), and it is indeed the case that, in the studies reported here, sex often had sizable main effect associations with the outcome variables. However, sex virtually never had a sizable interactive effect on the relationship between other variables, suggesting that sex is not a particularly potent moderator of relationships between other variables.

Under what sort of conditions might these results arise? To answer this question, we conducted two illustrative simulations, each with 10,000 simulated data sets of $N = 100$. In the first, we modeled two uncorrelated main effects predicting an outcome variable at $r = .30$ and a third interaction variable correlated $r = .00$ with the outcome. In the second simulation, we changed this latter value to $r = .10$. In the first simulation, the number of statistically significant interaction terms was, of course, 5% and the average R_{Δ} was .07 (adjusted $R_{\Delta} = .03$). In the second, the number of statistically significant interaction terms was 19% (approximately what one would expect based on a power analysis) and the average R_{Δ} was .11 (adjusted $R_{\Delta} = .07$). When the sample sizes were increased to 1,000, these values were 5% and 94% statistically significant with mean R_{Δ} of .02 and .10 (R_{Δ} adjusted .01 and .10), respectively. Thus, the simulation with results most similar (indeed, nearly identical) to those found in these 5 studies was the simulation in which the interaction term had zero relationship to the outcome variable.

We would be remiss if we did not emphasize the strengths of the 5 studies examined here. First, by social-personality psychology standards, the study sizes are relatively large, with the smallest study gathering data from 143 participants and the largest having over 7,000. Further, three of the studies employed repeated measures designs, greatly enhancing statistical power. Second, the total number of interactions examined in these studies is enormous and could easily be considered quite comprehensive of the person $\times$ situation interaction variable space. That is, Studies 1-3 used the CAQ, a broad and rich measure of personality. Studies 4 and 5 used well-validated factor-level measures of personality (the HEXACO and HPI), which also broadly cover the personality trait taxa. On the situation side, Studies 3 and 4 used actual measures of situations (the RSQ and the DIAMONDS), that broadly cover the landscape of situational variables. Likewise, Study 5 relied on the Holland RIASEC model of workplace situations, which broadly covers the workplace landscape. While Studies 1 and 2 did not measure psychological situations, they were manipulated via changes in laboratory tasks. Indeed, a third strength of this array of studies is the use of both laboratory and real-world study designs. A fourth strength of these studies is that they allowed us to examine interactions in both within-person (Studies 1, 3, and 4) and between-person designs (Studies 2 and 5). Finally, while two of these studies did use self-reported behavioral outcomes (Studies 3 and 4) three of the studies did not. The outcomes in Studies 1 and 2 were directly observed behaviors during social interactions, while the outcome of Study 5 was supervisor rated job performance. Thus, the studies presented here do not rely merely on self-reports of behavior.

Implications and Connections to Prior Research

In the Introduction, we highlighted the importance of the topic investigated here from the perspective of contemporary debates about replication. When replications fail, should the default

assumption be that the original result was correct in the context in which it was measured, but the replication differs in the level of some moderating variable? A high level of disagreement about this point seems to exist in the field at present. The results described here suggest that interaction terms are only rarely substantial. This implies that the assumption that replication failures represent subtle moderators (e.g., the clothes the research assistant wore) is quite suspect.

Another potentially important but indirect implication relates to the question of how participants should be selected for behavioral research. Many statisticians have been critical of the tendency of behavioral researchers to make heavy use of “convenience samples,” i.e., groups such as college students or users of particular websites who happen to be conveniently available to the researcher (e.g., Berk & Freedman, 2001). These critics have argued that the use of convenience samples seriously undermine the validity of research. In their view, credible research requires nothing less than probability samples (samples in which every member of the population has an approximately equal likelihood of being recruited). What do the present results say about this issue? If a researcher aims to determine what percent of the population holds a particular attitude or behavior, as in political polling, there is obviously no substitute for a probability sample. On the other hand, especially if the goal is merely to assess the direction of the effect of a given variable, as is often the case in behavioral science (see, e.g., Loftus, 1978) then given the rarity of crossover interactions revealed in the current study, the use of convenience samples may generally be adequate. (For interesting discussions of the issue from varying disciplines, see also Druckman and Kam, 2011, from political science and Weinberg, Freese, & MacEhatten, 2014, from sociology, and Hultsch, MacDonald, Hunter, Maitland, and Dixon, 2002, from psychology.)

Counter-Arguments and Limitations

Why would experienced investigators assume that powerful interaction effects are commonplace, if in fact they are relatively rare? For one thing, the flexibility and context-dependence of human behavior clearly *allow for* dramatic interaction effects to exist. One can see this in the strong and replicable interaction effects that are widely believed to exist in the literature. (Unfortunately, such beliefs may be misguided given the reality of publication bias.) But even if one is skeptical about interactions that appear consistently in the literature, as a thought experiment it is easy to envision interventions that, one has little doubt, could produce dramatic interactions “on cue.” For example, suppose an experimenter spoke with half the participants en route to a laboratory to participate in a behavioral experiment and told them that it was important that they do the opposite of whatever was asked of them in the study (an intervention that some innovative hypnosis researchers came close to doing; see Orne, 1971). If the intervention was carried out effectively, is there any doubt it could invert effects of some independent variables, and thus produce a strong crossover interaction?

The point being made here is merely that strong interactions are obviously *possible*. Indeed, the published experimental psychology literature frequently reports interactions. How can such results be reconciled with the data reported here? First, experimental studies are frequently *designed* to find interaction effects via extreme manipulations of one independent variable. Such extreme manipulations allow one to detect, via statistical significance testing, such interactions at the extremes of variables. However, in realistic settings such variables rarely reach the extremes. Thus, interaction effects demonstrated in experimental contexts may have no practical value in everyday settings. Second, it is presently difficult to rely on the published literature to gauge the ubiquity of interaction effects given the joint contributions of publication bias and data collection/analytic flexibility. In contrast to published studies, the studies presented

here tested comprehensively (i.e., not selectively) for person by situation interactions and are not subject to such concerns. Ultimately, the data described here suggest that interactions are probably quite uncommon in naturalistic settings.

While the data seem to make a compelling case that interaction effects are generally rare, an additional limitation should be acknowledged. While these studies are large and methodologically diverse, they are limited to kinds of variables they sampled over, which are personality and situation variables. Because personality broadly covers the domain of “person” variables and situations broadly cover the domain of “environmental” variables it is difficult to imagine what categories of variables would yield very different patterns than what we are seeing here. Thus, we believe the findings of this research generalize to interaction effects across the behavioral sciences. Although some may argue that these data only show that personality by situation interactions are rare and have nothing to say about other kinds of interactions (whatever they may be), we have serious doubts about the promise of interactions for providing meaningful explanatory power. Indeed, the evidence for additive main effects of genes is quite compelling for human behavior (Hill, Goddard, & Visscher, 2008).

Conclusion and Recommendations

The presumption of interaction effects appears to be ubiquitous in psychology. However, the studies examined here indicate that substantial interactions between person and situation variables are actually quite rare in broad samples of psychological variables (including sex) that are not selected for publication based on exhibiting interesting patterns of interactions. The implications for psychological science seem more than minor. First, these results suggest that researchers should have strong theoretical rationales for examining moderators. Second, researchers should evaluate statistical evidence of moderators with substantial skepticism.

Indeed, we would go so far as to say that unanticipated findings of interactions should not be considered “real” barring pre-registered studies and/or direct replication. Third, even when the data are persuasive enough to convince us that a moderator is real, we should take a careful look at the explanatory power and ask ourselves about the practical import of such a moderator in real-world contexts. Will the moderator actually matter? Fourth, rather than examining potential moderators of relationships, we suspect researchers would gain more in the way of explanatory power by considering alternative predictor variables (as main effects). Finally, the results of these studies indicate that moderators should not be frequently expected in social/personality psychology. Thus, claims that replication failures rest on unknown moderators seem unlikely to be right very often.

References

- Ashton, M. C., & Lee, K. (2009). The HEXACO–60: A short measure of the major dimensions of personality. *Journal of personality assessment*, 91(4), 340-345.
- Barrett, L. F. (2015). Psychology is not in crisis. *The New York Times*, A65.
- Bem, D. J., & Funder, D. C. (1978). Predicting more of the people more of the time: Assessing the personality of situations. *Psychological Review*, 85(6), 485.
- Berk, R., & Freedman, A. E. (2001). Statistical assumptions as empirical commitments. In T. G. Blomberg and S. Cohen (Eds.), *Law, Punishment, and Social Control: Essays in Honor of Sheldon Messinger*, 2nd ed, (pp. 235–254). Aldine de Gruyter.
- Block, J. (1961). The Q-sort method in personality assessment and psychiatric research. Springfield, IL: Charles C. Thomas
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365-376.
- Chaplin, W. F. (1991). The next generation of moderator research in personality psychology. *Journal of Personality*, 59(2), 143-178.
- Chaplin, W. F. (1997). Personality, interactive relations, and applied psychology. In R. Hogan J. Johnson, & S. Briggs (Eds). *Handbook of Personality Psychology* (pp. 873-890). Orlando, FL: Academic Press.
- Cohen, J., Cohen, P., West, S. G., & Aiken, L. S. (2003). *Applied Multiple Regression / Correlation Analysis for the Behavioral Sciences* (3rd edition). Mahwah, NJ: Lawrence Erlbaum Associates.
- Cronbach, L. J. (1975). Beyond the two disciplines of scientific psychology. *American Psychologist*, 30, 116-127.

- Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin*, 81, 95-106.
- Dijksterhuis, A., Van Knippenberg, A., & Holland, R. W. (2014). Evaluating behavior priming research: Three observations and a recommendation. *Understanding Priming Effects in Social Psychology*, 205.
- Druckman, James N. and Cindy D. Kam. 2011. "Students as Experimental Participants: A Defense of the 'Narrow Data Base'." In *Handbook of Experimental Political Science*, eds. James N. Druckman, Donald P. Green, James H. Kuklinski, and Arthur Lupia. New York: Cambridge University Press, 41-57.
- Ghiselli, E. E. (1963). Moderating effects and differential reliability and validity. *Journal of Applied Psychology*, 47(2), 81-86.
- Funder, D. C., Furr, R. M., & Colvin, C. R. (2000). The Riverside Behavioral Q-sort: A tool for the description of social behavior. *Journal of Personality*, 68, 451-489.
- Furr, R. M., & Funder, D. C. (1998). A multimodal analysis of personal negativity. *Journal of Personality and Social Psychology*, 74(6), 1580-1591.
- Furr, R. M., Wagerman, S., & Funder, D. C. (2010). Personality as manifest in behavior: Direct behavioral observation using the revised Riverside Behavioral Q-sort (RBQ-3.0). In C. R. Agnew, D. E., Carlston, W. G. Graziano, & J. R. Kelly (Eds.), *Then a miracle occurs: Focusing on behavior in social psychological theory and research*. (pp. 186-204). Oxford University Press.
- Hagger, M. S., Wood, C., Stiff, C., & Chatzisarantis, N. L. (2010). Ego depletion and the strength model of self-control: a meta-analysis. *Psychological bulletin*, 136(4), 495.

- Hogan, R., & Hogan, J. (1995). *The Hogan Personality Inventory manual* (2nd ed.). Tulsa, OK: Hogan Assessment Systems.
- Holland, J. L. (1966). *The Psychology of Vocational Choice*. Waltham, Massachusetts: Ginn and Compnay.
- Hill, W. G., Goddard, M. E., & Visscher, P. M. (2008). Data and theory point to mainly additive genetic variance for complex traits. *PLoS Genetics* 4(2): e1000008.
<https://doi.org/10.1371/journal.pgen.1000008>
- Hultsch, D. F., MacDonald, S. W., Hunter, M. A., Maitland, S. B., & Dixon, R. A. (2002). Sampling and generalisability in developmental research: Comparison of random and convenience samples of older adults. *International Journal of Behavioral Development*, 26(4), 345-359.
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2, 696-701.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science*, 0956797611430953.
- Kenrick, D. T., & Dantchik, A. (1983). Interactionism, idiographics, and the social psychological invasion of personality. *Journal of Personality*, 51(3), 286-307.
- Kenrick, D. T., & Funder, D. C. (1988). Profiting from controversy: Lessons from the person-situation debate. *American Psychologist*, 43(1), 23.
- Kenrick, D. T., Neuberg, S. L., & Cialdini, R. B. (2015). *Social Psychology: Goals in Interaction* (6th edition). Boston: Pearson.
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr, R. B., Bahník, Š., Bernstein, M. J., ... &

- Woodzicka, J. A. (2015). Investigating variation in replicability. *Social Psychology*.
- Letzring, T. D., Wells, S. M., & Funder, D. C. (2006). Information quantity and quality affect the realistic accuracy of personality judgment. *Journal of personality and social psychology*, 91(1), 111.
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6(3), 312-319.
- Lyubomirsky, S., & Lepper, H. (1999). A measure of subjective happiness: Preliminary reliability and construct validation. *Social Indicators Research*, 46, 137-155.
- Magnusson, D. S., & Endler, N. S. (1977). *Personality at the Crossroads: Current Issues in Interactional Psychology*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- McClelland, G. H., & Judd, C. M. (1993). Statistical difficulties of detecting interactions and moderator effects. *Psychological Bulletin*, 114(2), 376-390.
- Meyer, M., & Chabris, C. (2014). Why psychology's food fight matters: "Important findings" haven't been replicated, and science may have to change its ways. *Slate Magazine*, July 31, 2014.
- Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R² from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4, 133-142.
- Open Science Collaboration. "Estimating the reproducibility of psychological science." *Science* 349.6251 (2015): aac4716.
- Orne, M. T. (1971). The simulation of hypnosis: Why, how, and what it means. *The International journal of clinical and experimental hypnosis*, 19(4), 183-210.
- Rauthmann, J. F., Gallardo-Pujol, D., Guillaume, E. M., Todd, E., Nave, C. S., Sherman, R. A., Ziegler, M., Jones, A. B., & Funder, D. C. (2014). The situational eight DIAMONDS: A

- taxonomy of major dimensions of situation characteristics. *Journal of Personality and Social Psychology*, 107(4), 677-718.
- Rauthmann, J. F., & Sherman, R. A. (2016). Ultra-brief measures for the situational eight DIAMONDS domains. *European Journal of Psychological Assessment*, 32(2), 165-174.
- Rossi, J. S. (1990). Statistical power of psychological research: What have we gained in 20 years?. *Journal of consulting and clinical psychology*, 58(5), 646-656.
- Schmitt, D. P., Realo, A., Voracek, M., & Allik, J. (2008). Why can't a man be more like a woman? Sex differences in Big Five personality traits across 55 cultures. *Journal of personality and social psychology*, 94(1), 168.
- Sherman, R. A., & Funder, D. C. (2009). Evaluating correlations in studies of personality and behavior: Beyond the number of significant findings to be expected by chance. *Journal of Research in Personality*, 43, 1053-1061.
- Sherman, R. A., Nave, C. S., & Funder, D. C. (2010). Situational similarity and personality predict behavioral consistency. *Journal of Personality and Social Psychology*, 99(2), 330-343.
- Sherman, R. A., Rauthmann, J. F., Brown, N. A., Serfass, D. S., & Jones, A. B. (2015). The independent effects of personality and situations on real-time expressions of behavior and emotion. *Journal of Personality and Social Psychology*, 109(5), 872-888.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359-1366.
- Stroebe, W., & Strack, F. (2014). The alleged crisis and the illusion of exact replication. *Perspectives on Psychological Science*, 9(1), 59-71.

- Wagerman, S. A., & Funder, D. C. (2009). *Personality psychology of situations*.
- Weinberg, J. D., Freese, J., & McElhattan, D. (2014). Comparing Data Characteristics and Results of an Online Factorial Survey between a Population-Based and a Crowdsource-Recruited Sample. *Sociological Science 1*: 292-310. DOI: 10.15195/v1.a19
- Yong, E. (2012). In the wake of high profile controversies, psychologists are facing up to problems with replication. *Nature*, 483, 298-300.
- Zell, E., Krizan, Z., & Teeter, S. R. (2015). Evaluating gender similarities and differences using metasynthesis. *American Psychologist*, 70(1), 10-20.

Footnotes

¹ The first author initially analyzed the data for Study 3 years ago hoping to find evidence for many interactions. A conversation about the (disappointing) results of Study 3 between the two of us sparked our interest in conducting parallel analyses in these other data sets. Thus, Study 3 may be viewed as exploratory while the others viewed as confirmatory. The studies are presented in the paper chronologically based upon time of data collection.

² R scripts and data sets for all analyses reported in this paper are publicly available at <https://osf.io/9p2gx/>. Because the Hogan Data Archive is proprietary, we cannot make it publicly available.

³ In this instance, to estimate the number of statistically significant interaction effects under a null model we randomized scores, for each RBQ variable, first within participants and then between (while holding the within dependencies constant).

⁴ While in ordinary least squares regression zero is the lower bound for improvement in model fit, it is possible to have negative improvements in model fit on the *R* metric suggested by Nakagawa and Schielzeth (2013). Indeed, for 714 models this was the case. We fixed all negative values to zero before plotting.

⁵ As before, these incremental fits were calculated by taking the square root of the difference between the model R^2 values and fixing negative estimates of incremental fits to be zero.

⁶ Again, these incremental fits were calculated by taking the square root of the difference between the model R^2 values and fixing negative estimates of incremental fits to be zero.

⁷ The average number of employees per job was 121.9 ($SD = 130$, median = 82, min = 25, max = 784). The seven job families are as follows: Administrative & Clerical ($n = 236$), Managers & Executives ($n = 1,405$), Operations & Trades ($n = 1,849$), Professionals ($n = 835$), Sales & Customer Support ($n = 1,087$), Service & Support ($n = 1,012$), and Technicians & Specialists ($n = 890$).