

PRE-PROOF VERSION TO APPEAR IN

Sommet, N., Weissman, D. L., Cheutin, N., & Elliot, A. (2023). How many participants do I need to test an interaction? Conducting an appropriate power analysis and achieving sufficient power to detect an interaction. *Advances in Methods and Practices in Psychological Science*

5 **How Many Participants Do I Need to Test an Interaction? Conducting an Appropriate
Power Analysis and Achieving Sufficient Power to Detect an Interaction**

Nicolas Sommet^{1*} [ID 0000-0001-8585-1274](https://orcid.org/0000-0001-8585-1274),

David L. Weissman² [ID 0000-0002-0778-1992](https://orcid.org/0000-0002-0778-1992),

Nicolas Cheutin³, &

10 Andrew J. Elliot² [ID 0000-0002-1664-6426](https://orcid.org/0000-0002-1664-6426)

¹LIVES Center, University of Lausanne, Switzerland

²Department of Psychology, University of Rochester, NY, USA

³Unaffiliated software development engineer, Lausanne, Switzerland

15 *Correspondence to: nicolas.sommet@unil.ch

20

Submission: December 12, 2022

Revision: March 18, 2023

Acceptance: April 17, 2023

Word count: 7,911 words (excl. abstract, equations, tables/figures, disclosure, and references)

25

Nicolas Sommet, LIVES Center, SSP, ISS, University of Lausanne. Andrew J. Elliot, David L. Weissman, Department of Psychology, University of Rochester. Nicolas Cheutin, Unaffiliated.

Correspondence concerning this article should be addressed to Nicolas Sommet, Bâtiment Géopolis, Bureau #5785, Quartier UNIL-Mouline, Université de Lausanne, Switzerland; +41

30 (0)21 692 32 55; nicolas.sommet@unil.ch

Abstract

Power analysis for first-order interactions poses two challenges: (i) Conducting an appropriate power analysis is difficult, because the typical expected effect size of an interaction depends on its shape; (ii) Achieving sufficient power is difficult, because interactions are often modest in size. This paper consists of three parts. PART 1 addresses the first challenge. We first use a fictional study to explain the difference between power analyses for interactions and main effects. Then, we introduce an intuitive taxonomy of 12 types of interactions based on the shape of the interaction (reversed, fully attenuated, partially attenuated) and the size of the simple slopes (median, smaller, larger), and we offer mathematically-derived sample size recommendations to detect each interaction with a power of .80/.90/.95 (for two-tailed tests in between-participant designs). PART 2 addresses the second challenge. We first describe a preregistered meta-study (159 studies from recent articles in influential psychology journals) showing that the median power to detect interactions of a typical size is .18. Then, we use simulations ($\approx 900,000,000$ datasets) to generate power curves for the 12 types of interactions, and test three approaches to increase power without increasing sample size: (i) preregistering one-tailed tests (+21% gain), (ii) using a mixed design (+75% gain), and (ii) preregistering contrast analysis for a fully attenuated interaction (+62% gain). PART 3 introduces INT \times Power, a web-application that enables users to draw their interaction and determine the sample size needed to reach the power of their choice, with the option of using/combining these approaches: www.intxpower.com.

Keywords: Statistical power; first-order interaction; effect size; meta-study; simulations; web application.

A fundamental characteristic of good research practice is testing hypotheses with adequate statistical power (in the Null Hypothesis Significance Testing [NHST] framework; Cohen, 1988). Running studies with high power increases the likelihood of documenting not only true effects, but also replicable effects (Altmejd et al., 2019; Button et al., 2013; Fraley & Vazire, 2014). Whereas power analysis for main effects is relatively straightforward (Kovacs et al., 2022), power analysis for first-order (i.e., two-way) interactions (hereinafter simply referred to as “interactions”) poses two important challenges.

First, conducting appropriate power analysis for interactions cannot be easily carried out with software like G*Power (Faul et al., 2007) or packages like pwr (Champely et al., 2017) because the expected effect size of an interaction depends on its shape (Maxwell & Delaney, 2004). This can lead scholars to base their power analyses on an incorrect effect size. Second, achieving sufficient power to detect interactions is not an easy task either, because the effect size of interactions is often small and, by extension, the required sample size is often very large (Brysbaert, 2019). This can prevent scholars with limited resources (e.g., from low-income countries; Bredan, 2020) from achieving adequate power.

Although there has been much discussion about the issue of statistical power for interactions in popular blogs (Gelman, 2018; Giner-Sorolla, 2018; Simonsohn, 2014) and scientific articles (for classic work, see Aiken et al., 1991; Cronbach, 1987; McClelland & Judd, 1993; for recent work, see Blake & Gangestad, 2020; Lakens & Caldwell, 2021; Perugini et al., 2018), there are still no practical and easy-to-implement solutions to the two challenges mentioned above. This paper aims to propose such solutions for dichotomous and continuous predictors. It is divided into three parts.

In PART 1, we address the challenge of *conducting appropriate power analysis for interactions*: First, we offer a description of why decisions regarding power analyses differ between main and interaction effects, and then we introduce an intuitive taxonomy of 12 types of interactions along with the required N s to detect each of them with power = .80/.90/.95.

In PART 2, we address the challenge of *achieving sufficient power to detect interactions*: First, we report the findings from a meta-study showing that most studies testing interactions are underpowered, and then we describe the results from simulations testing three approaches to increase power for detecting interactions without increasing sample size.

In PART 3, we introduce [INT×Power](#), a user-friendly web application that enables researchers to draw a figure of their expected interaction to obtain the required N to achieve their desired power of (using the calculation described in PART 1), with the possibility of optimizing power without increasing sample size (using the approach described in PART 2).

PART 1. How to Conduct an Appropriate Power Analysis for Interactions

1.1. Calculating Effect Size and Required Sample Size for Interactions

In this part, we present a simple illustration of why expectations about effect size and sample size requirements differ between main effects (in a two-group design) and interaction effects (in a 2×2 designs). This illustration uses a dichotomous predictor and moderator, but applies to continuous predictors as well. Only basic statistical knowledge is required, and the formulas provided can be found in most textbooks on power analysis (e.g., Aberson, 2019; Cohen, 1988; Maxwell & Delaney, 2004).

95 *Calculating the Required Sample Size to Detect a Median-Sized Main Effect*

Imagine you are planning a study to estimate the effect of a new intervention aiming to improve people's well-being. Specifically, you intend to use a two-group experimental design to test the following hypothesis: "Compared to participants in the control group, participants in the intervention group report higher levels of well-being." To test your hypothesis, you will use a simple linear regression:

$$Well-being_i = B_0 + B_1 \times Condition_i + e_i \quad (1)$$

...with $i = 1, 2, 3, \dots, N$ (number of participants), where $Condition_i = -0.5$ for the control group and $Condition_i = +0.5$ for the intervention group, and e_i represents the error¹

To plan your study efficiently, you must first determine the expected effect size of $Condition_i$. As your intervention is new, you do not have a clear sense about the expected magnitude of its effect, and you decide to use the median effect size in psychology, namely, Cohen's $d = 0.35$ (for the rationale on how we choose this value, see section 1.2). Assuming equal sample sizes and a standard deviation of 1 for both groups, you expect the mean well-being score to be 0.35 points higher in the intervention than in the control group:

$$\begin{aligned} d_{(Condition_i)} &= \frac{M_{intervention} - M_{control}}{Pooled\ SD} \\ &= \frac{M_{intervention} - M_{control}}{1} \\ &= M_{intervention} - M_{control} \\ &= .35 \end{aligned} \quad (2)$$

...where $M_{intervention}$ and $M_{control}$ represents the standardized mean in the intervention group and control group, respectively.

¹ In this and the subsequent equations, we use sample notation for simplicity.

You now wonder how many participants are needed to detect such a median-sized main effect with a statistical power of $1 - \beta = .80$ (i.e., if the effect exists, there is an 80% chance of detecting a true positive) using a two-tailed test with an α of .05 (i.e., if the effect does *not* exist, there is a 5% risk of detecting a false positive). To calculate the required N , you enter the three
 115 key parameters (expected d , target power, α) in G*Power, which uses a formula very similar to the following:²

$$\begin{aligned}
 N &= \frac{4 \times (Z_{1-\alpha/2=.025} + Z_{1-\beta=.80})^2}{d_{(Condition_i)}^2} \\
 &= \frac{4 \times (1.96 + 0.84)^2}{0.35^2} && (3) \\
 &= 256
 \end{aligned}$$

...where $Z_{1-\alpha/2=.025} = 1.96$ and $Z_{1-\beta=.80} = 0.84$ are the critical Z-values associated with a two-tailed test with $\alpha = .05$ and $1 - \beta = .80$, respectively.

Simply put, assuming that your intervention works as expected, a sample size of $N =$
 120 256 participants will give you an 80% probability to observe a median-sized (or larger) effect of $Condition_i$ with $p < .05$.

Calculate the Required Sample Size to Detect an Interaction Effect of a Typical Size

Now imagine you believe that your intervention should only benefit women, not men.³
 You are planning to use a 2×2 factorial design to test the following hypothesis: “Compared to
 125 women in the control group, women in the intervention group report higher levels of well-being; for men, there is no difference between the two conditions.” To test your hypothesis, you will use the multiple linear regression below:

$$\begin{aligned}
 Well-being_i &= B_0 + B_1 \times Condition_i + B_2 \times Gender_i + \\
 &B_3 \times Condition_i \times Gender_i + e_i && (4)
 \end{aligned}$$

...where $Gender_i = -0.5$ for men and $+0.5$ for women.

To plan your study efficiently, you must again begin by determining the expected effect
 130 size of $Condition_i \times Gender_i$. In this situation, you may *feel* it is reasonable to use the same generic value of $d = 0.35$ used above to describe an interaction effect of typical size. Then, since

² We approximate the t -distribution using the Z -distribution, which simplifies calculation as the distribution does not vary based on degrees of freedom. As G*Power uses critical values of the t -distribution, there might be slight discrepancies between the values generated by the software and the values produced by Eq. 3. These discrepancies will disappear as degrees of freedom increase.

³ In this example, we treat gender as if it was a binary variable for simplicity.

you now have a 2×2 instead of a two-group design, it may seem logical to double the sample size (for examples of recent papers following this reasoning, see Majer et al., 2022; Tepe & Byrne, 2022; Wang & Xie, 2021). However, this would be a mistake.

135 The reason why this is a mistake is simple: Contrary to the Cohen's d of a main effect in a two-group design, the Cohen's d of an interaction in a 2×2 design should not be seen as a difference between means, but as *a difference between subdifferences*. Assuming equal sample sizes and a standard deviation of 1 for each of your subgroups, the calculation corresponds to the difference between (i) the subdifference between $M_{intervention\text{♀}}$ and $M_{control\text{♀}}$ for women (the simple slope $d_{(Condition_1)\text{♀}}$) and (ii) the subdifference between $M_{intervention\text{♂}}$ and $M_{control\text{♂}}$ for men (the simple slope $d_{(Condition_1)\text{♂}}$):

$$\begin{aligned}
 d_{(Condition_1 \times Gender_1)} &= \frac{(M_{intervention\text{♀}} - M_{control\text{♀}}) - (M_{intervention\text{♂}} - M_{control\text{♂}})}{2 \times Pooled\ SD} \\
 &= \frac{d_{(Condition_1)\text{♀}} - d_{(Condition_1)\text{♂}}}{2}
 \end{aligned} \tag{5}$$

Because the effect size of an interaction is derived from the effect sizes of its simple slopes, it is *not* reasonable to expect the effect size of $Condition_1 \times Gender_1$ to be as high as $d = 0.35$. Indeed, given that the simple slope for men is null, such an interaction would involve an unusually large simple slope for women of $d_{(Condition_1)\text{♀}} = 0.70$.

$$\begin{aligned}
 d_{(Condition_1 \times Gender_1)} &= \frac{d_{(Condition_1)\text{♀}} - d_{(Condition_1)\text{♂}}}{2} \\
 &= \frac{0.70 - 0.00}{2} \\
 &= 0.35
 \end{aligned} \tag{6}$$

In this case, it would be more reasonable to expect the effect size of $Condition_1 \times Gender_1$ to be $d = 0.175$, because such an interaction would this time involve a median simple slope for women of $d_{(Condition_1)\text{♀}} = 0.35$.

$$\begin{aligned}
 d_{(Condition_1 \times Gender_1)} &= \frac{d_{(Condition_1)\text{♀}} - d_{(Condition_1)\text{♂}}}{2} \\
 &= \frac{0.35 - 0.00}{2} \\
 &= 0.175
 \end{aligned} \tag{7}$$

From there, you can use Eq. 3 to calculate the required N to detect an interaction effect of $d = 0.175$ with a statistical power of .80 using a two-tailed test with $\alpha = .05$. You will realize

that you do not need a sample twice as large but *four times as large* as the sample used in the first case, that is, $N = 1,024$ participants.

$$\begin{aligned}
 N &= \frac{4 \times (Z_{1-\alpha/2=.025} + Z_{1-\beta=.80})^2}{d(\text{Condition}_1 \times \text{Gender}_1)^2} \\
 &= \frac{4 \times (1.96 + 0.84)^2}{0.175^2} \\
 &= 1024
 \end{aligned} \tag{8}$$

155 While this example is only one of many possible interactions, it illustrates the danger of relying on a generic value to define the expected effect size of an interaction (e.g., believing that an interaction of a typical size will always be $d = 0.35$). A less error-prone approach is to define the expected effect sizes of the simple slopes (e.g., a median simple slope of $d = 0.35$ combined with a null simple slope of $d = 0.00$), and work from there to determine the expected effect size of the interaction and the required sample size. In the next section, we propose a taxonomy of interactions that will allow us to provide comprehensive sample size recommendations.

160 1.2. A Taxonomy of 12 Types of Interactions

Now, we introduce an intuitive taxonomy of 12 types of interactions based on two criteria: (1) the expected shape of the interaction (based on the signs of the simple slopes) and (2) the expected sizes of the simple slopes (based on the values of the simple slopes). We give the required sample size to detect each interaction with adequate power.

165 *Criterion 1. What is the Expected Shape of My Interaction?*

The literature typically defines three basic shapes of interaction: (i) a reversed interaction, (ii) a fully attenuated interaction, and (iii) a partially attenuated interaction (for related distinctions, see Baranger et al., 2022; Blake & Gangestad, 2020; Brysbaert, 2019; Giner-Sorolla, 2018; Lakens & Caldwell, 2021; Ledgerwood, 2019; Perugini et al., 2018).

170 First, a *reversed interaction* (also known as a “cross-over interaction” or “reversal of the effect”) involves simple slopes that go in opposite directions (see **Figure 1**, left panel). For instance, joining (vs. leaving) a group that one holds in high esteem increases life satisfaction, whereas joining (vs. leaving) a group that one holds in low esteem decreases life satisfaction (DeMarco & Newheiser, 2019).

175 Second, a *fully attenuated interaction* (also known as a “knockout interaction” or “elimination of the effect”) involves a simple slope that goes in one direction, and a simple slope that is null (see **Figure 1**, middle panel). For instance, students from lower socioeconomic backgrounds benefit from having a growth mindset, whereas students from higher socioeconomic

backgrounds do not (Sisk et al., 2018).

180 Third, a *partially attenuated interaction* (also known as a “spreading interaction” or
“attenuation of the effect”) involves simple slopes that go in the same direction, but for which
one slope is steeper than the other (see **Figure 1**, right panel). For instance, middle-aged indi-
viduals who are less educated are more likely to be depressed, whereas in older individuals, the
link between education and depression persists but is weaker (a trend known as the “age-as-
185 leveler pattern”; Abrams & Mehta, 2019).

Criterion 2. What Are the Expected Effect Sizes of My Simple Slopes?

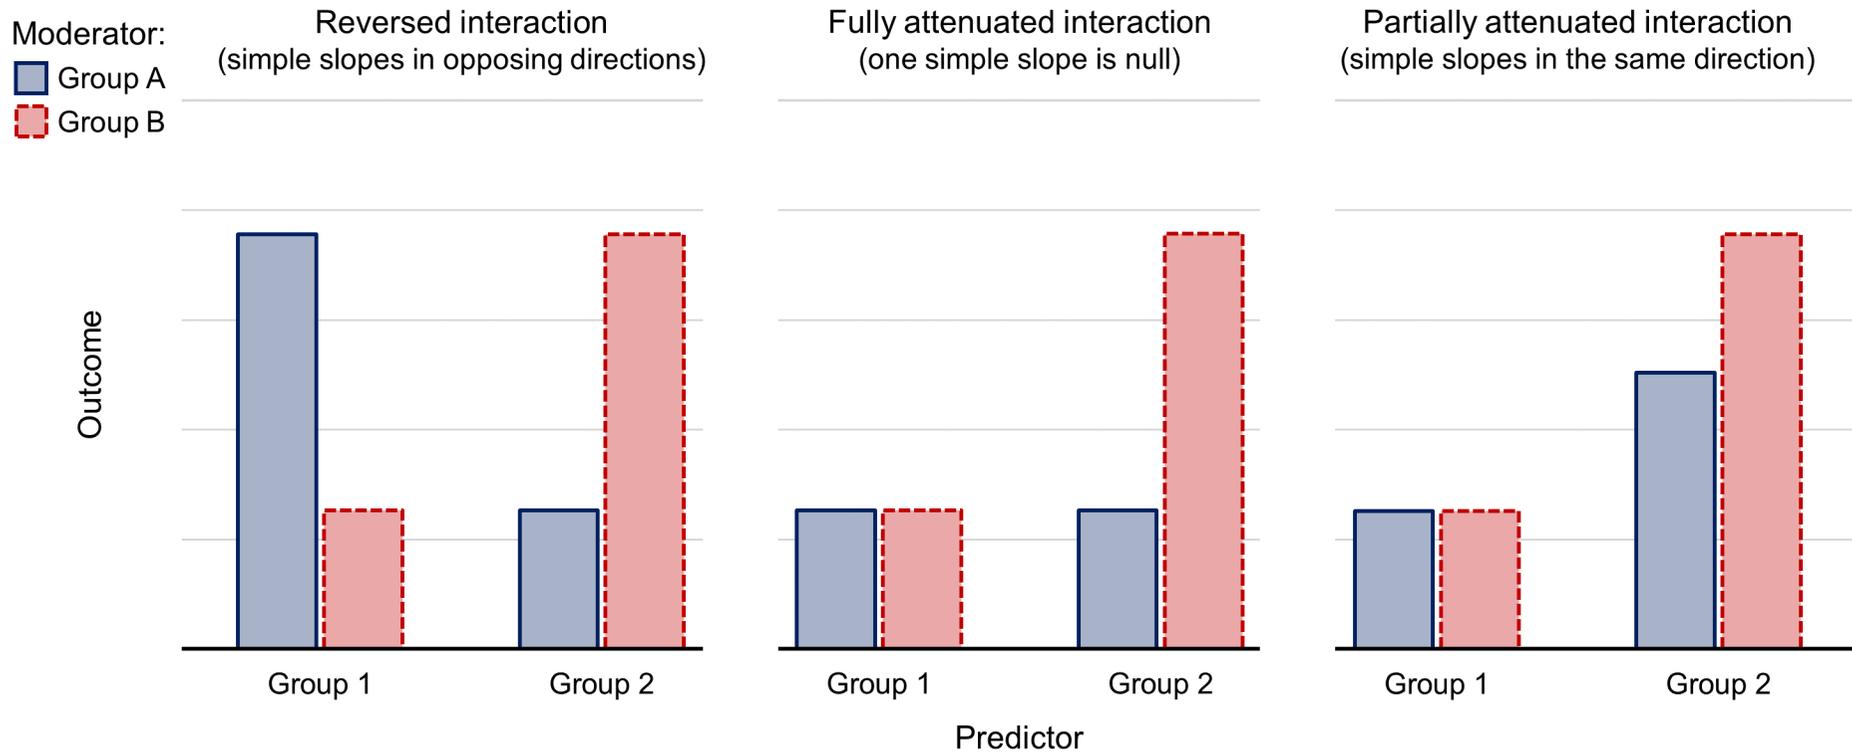
In an ideal world, researchers would always base their expectations for effect size on
prior, high-powered, preregistered studies. However, researchers often test new hypotheses for
which no such studies are available, leaving them with only a vague idea of the effect size in
190 the population. Consequently, researchers often resort to using Cohen’s (1965) benchmarks
(e.g., when using G*Power), which equate small, medium, and large effects to standardized
mean differences of $d = 0.20$, 0.50 , and 0.80 , respectively (see Eq. 2). Although these bench-
marks are widely used today, Cohen (1988) himself recognized that they offer “no more reliable
a source than intuition” and argued that “[w]ith the accumulation of experience, [these bench-
195 marks] may well require revision (I suspect downward)” (p. 478).

Many scholars have discussed substituting Cohen’s classic benchmarks for empirically
derived benchmarks (Brysbaert, 2019; Funder & Ozer, 2019; Hemphill, 2003). To this end,
Gignac and Szodorai (2016) gathered correlation coefficients from 87 meta-analyses in so-
cial/personality psychology, and established that the 25th, 50th and 75th percentiles roughly
200 corresponded to Cohen’s d s of 0.20 , 0.40 , and 0.60 , respectively (see also Fraley & Marks,
2007; Lovakov & Agadullina, 2021; Richard et al., 2003). However, effect sizes from non-
preregistered studies are inflated by two to five times due to publication bias (Camerer et al.,
2018; Ebersole et al., 2020; Klein et al., 2018), and these meta-analytically derived benchmarks
are likely overestimations (Correll et al., 2020).

205 More recently, Schäfer and Schwarz (2019) gathered effect sizes from 93 randomly se-
lected preregistered studies in psychology and found a slightly lower median effect size of about
 $d = 0.35$. Importantly, the authors documented significant variation in effect sizes among sub-
disciplines and study designs, indicating that there are no universally applicable benchmarks
(see also Adachi & Willoughby, 2015; Flora, 2020; Lakens, 2013). Despite this limitation, we
210 suggest that the following empirically derived benchmarks may be useful heuristic descriptions
of relatively small, median, and relatively large effects: $d = 0.20$, 0.35 and $d = 0.50$, respectively

Figure 1

Example of a reversed interaction (left panel), a fully attenuated interaction (middle panel), and a partially attenuated interaction (right panel) in the context of a 2 (Predictor: Group 1 vs. Group 2) × 2 (Moderator: Group A vs. Group B) design.



215

Note: These are just three examples of reversed, fully attenuated, and partially attenuated interactions; other configurations are possible. For the reversed interaction, the main effects are set to zero, resulting in a symmetrical pattern; if one of the main effects was nonzero, the pattern would be asymmetrical. As another example, for the fully attenuated interaction, the two main effects are equal in size and positive, resulting in an ordinal interaction (the cross-over of predicted values is at the boundary); if one of the main effects was negative, the interaction would become dis-ordinal (Widaman et al., 2012).

220

Table 1.

Empirically derived benchmarks used in the present paper.

	Cohen's d	Cohen's f	Pearson's r	η_p^2
Smaller	0.20	0.10	.10	.01
Median	0.35	0.175	.17	.03
Larger	0.50	0.25	.24	.06

Note: These value are based on former work (in particular, Gignac & Szodorai [2016] and Schäfer & Schwarz [2019]). For conversion formulas, see **Appendix**, Eq A₁-A₃, on the OSF.

225

(for correspondences with other types of estimates, see **Table 1**). These benchmarks can be used to quantify archetypal effect sizes or to define the smallest effect size of interest (Anvari & Lakens, 2021). For instance, $d = 0.20$ might sometimes be considered as the minimum effect size of theoretical or practical significance.

230 **12 Types of Interactions**

Description of the Taxonomy. As can be seen in **Table 2**, our taxonomy covers interactions involving any combination of median ($d = 0.35$), smaller ($d = 0.20$), and larger ($d = 0.50$) simple slopes. This encompasses:

• Six reversed interactions. The *typical* reversed interaction is the “+0.35 | -0.35 reversed interaction,” which involves a median simple slope that goes in one direction (e.g., $d = +0.35$) and a median simple slope that goes in another direction (e.g., $d = -0.35$). Note that the order of the signs is arbitrary, as a +0.35 | -0.35 and a -0.35 | +0.35 interaction have the same overall effect size of $|d| = 0.35$, which is equivalent to a median-sized main effect.

• Three fully attenuated interactions. The *typical* fully attenuated interaction is the “+0.35 | 0.00 fully attenuated interaction,” which involves a median simple slope that goes in either direction (e.g., $d = +0.35$) and a null simple slope ($d = 0.00$). Note that the sign of the nonnull simple slope is arbitrary, as a +0.35 | 0.00 and a -0.35 | 0.00 interaction have the same overall effect size of $|d| = 0.175$, which is half the size of a median-sized main effect.

• Three partially attenuated interactions. The *typical* partially attenuated interactions is the “+0.20 | +0.35” or the “+0.35 | +0.50 partially attenuated interaction,” which involve a simple slope that goes in one direction (e.g., $d = +0.20$) and a larger simple slope that goes in the same direction (e.g., $d = +0.35$). Note that the common sign of the simple slopes is arbitrary, as

245

a $-0.20 \mid -0.35$ and a $-0.35 \mid -0.50$ interaction have the same overall effect size of $|d| = 0.075$, which is nearly 5 times smaller than a median-sized main effect.

250 **Sample Size Recommendation in Between-Participant Designs and Assumptions.**

In **Table 2**, we provide the required sample sizes to achieve power of .80, 90, or .95 when using a two-tailed test with $\alpha = .05$ to detect each of the 12 interaction in 2×2 designs. The values of .80 and .90 are commonly used as lower limits for acceptable statistical power (Cohen, 1988), whereas .95 is a more stringent standard that may be useful in certain circumstances, such as
255 when planning an exact replication (Hedges & Schauer, 2019).

These required sample sizes can be used as case-sensitive recommendations, provided that the usual assumptions of linear regression are met: approximate multivariate normality (O'Connor, 2006), homogeneity of variance across subgroups (Overton, 2001), independence of residual error (Arend & Schäfer, 2019), and lack of severe multicollinearity (Shieh, 2010).
260 However, in the context of interaction, two additional assumptions deserve particular attention (for relevant research, see Aguinis, 1995; Aguinis & Gottfredson, 2010; Jaccard et al., 2003; McClelland & Judd, 1993).

First, equal sample sizes are expected across subgroups. In our introductory example, this mean having the same number of women and men in each condition. When this assumption
265 is violated, the power to detect the interaction decreases, albeit only slightly. For instance, when the ratio is 2:1 instead of 1:1 (e.g., twice the number of women compared to men), power is reduced by $\approx 5\%$. However, when the ratio is 9:1, power is reduced by 33% (Stone-Romero et al., 1994).

Second, measurement error should be zero. In our introductory example, with regard to
270 the moderator, this means that 100% of the participants are expected to report their gender without error (e.g., selecting the wrong response option). When this assumption is violated, the power to detect the interaction decreases in proportion to the degree of error: For instance, if 95% of the participants correctly reported their gender, the expected effect size used in the power analysis should be adjusted with $d_{adj} = d \times .95$ (Blake & Gangestad, 2020). Note that our
275 benchmarks, which are empirically derived from actual studies with measurement error, already take this into account (meaning if you manage to enhance reliability, you can use a larger effect size).

Anticipating Two Criticisms. Two potential arguments can be made against our taxonomy. First, one could argue that interactions with different shapes may have the same overall
280 effect size, calling into question the relevance of the concept of “shape.” For example, reversed,

fully attenuated, and partially attenuated interactions could all theoretically have an effect size of $d = 0.35$ and require the same N to reach $power = 0.80$. However, we find this argument misleading as reversed interactions with $d = 0.35$ involve median-sized simple slopes and are therefore relatively common, while attenuated interactions with $d = 0.35$ involve huge simple slopes and are almost never encountered.

Second, one can argue that the 12 interactions correspond to specific combinations of interactive and main effects, calling into question the importance of simple slopes. For instance, a “fully attenuated interaction $+0.35 \mid 0.00$ ” corresponds to a combination of interactive and main effects of $ds = .175$. However, we also find this argument irrelevant as researchers testing interactions do not typically think in terms of specific combinations of interactive and main effects. Rather, they hypothesize interaction patterns, making it practical to think about the relative sizes and signs of their expected simple slopes.

In PART 1, we saw that the required sample size to detect interactions is often larger than one would intuitively estimate, especially for attenuated interactions (for related research, see Bakker et al., 2016). In PART 2, we report the findings of a preregistered meta-study showing that most studies testing interactions are indeed underpowered, and then we describe the results from simulations testing three approaches to increase power without increasing sample size.

Table 2.

300 Taxonomy of 12 possible types of interactions as a function of shape (Criteria 1) and expected effect sizes of simple slopes (Criteria 2); we provide the total required sample sizes to achieve a power of .80, .90 or .95 using a two-sided test with dichotomous/continuous predictors in a between-participant design ($\alpha = .05$).

Criterion 1. Shape	Criterion 2.				Overall interaction effect size	Total required sample size		
	Expected effect sizes (Cohen's <i>d</i>)					Power = .80	Power = .90	Power = .95
	Simple Slope #1	Simple Slope #2						
Reversed	+0.20 (smaller)	-0.20 (smaller)			0.20	<i>N</i> = 784	<i>N</i> = 1,051	<i>N</i> = 1,300
Reversed	+0.20 (smaller)	-0.35 (median)			0.28	<i>N</i> = 415	<i>N</i> = 556	<i>N</i> = 687
Reversed	+0.20 (smaller)	-0.50 (larger)			0.35	<i>N</i> = 256	<i>N</i> = 343	<i>N</i> = 424
Reversed	+0.35 (median)	-0.35 (median)			0.35	<i>N</i> = 256	<i>N</i> = 343	<i>N</i> = 424
Reversed	+0.35 (median)	-0.50 (larger)			0.43	<i>N</i> = 174	<i>N</i> = 233	<i>N</i> = 288
Reversed	+0.50 (larger)	-0.50 (larger)			0.50	<i>N</i> = 125	<i>N</i> = 168	<i>N</i> = 208
Fully attenuated	+0.20 (smaller)	0.00 (null)			0.10	<i>N</i> = 3,136	<i>N</i> = 4,204	<i>N</i> = 5,198
Fully attenuated	+0.35 (median)	0.00 (null)			0.18	<i>N</i> = 1,024	<i>N</i> = 1,373	<i>N</i> = 1,697
Fully attenuated	+0.50 (larger)	0.00 (null)			0.25	<i>N</i> = 502	<i>N</i> = 673	<i>N</i> = 832
Partially attenuated	+0.20 (smaller)	+0.35 (median)			-0.08	<i>N</i> = 5,575	<i>N</i> = 7,474	<i>N</i> = 9,241
Partially attenuated	+0.20 (smaller)	+0.50 (larger)			-0.15	<i>N</i> = 1,394	<i>N</i> = 1,869	<i>N</i> = 2,310
Partially attenuated	+0.35 (median)	+0.50 (larger)			-0.08	<i>N</i> = 5,575	<i>N</i> = 7,474	<i>N</i> = 9,241

305 *Note:* Typical interactions (i.e., involving at least one median simple slope) are highlighted; the *Ns* were calculated using Eq. 3. For the reversed interactions, the signs of simple slopes can be switched (e.g., a +0.20 | -0.35 and a -0.20 | +0.35 reversed interaction require the same *N*); for the fully attenuated interactions, the signs of simple slope #1 can either be positive or negative (e.g., a +0.20 | 0.00 and a -0.20 | 0.00 fully attenuated interaction require the same *N*); for the partially attenuated interactions, the signs of the simple slopes can either be both positive or both negative (e.g., a +0.20 | +0.35 and a -0.20 | -0.35 partially attenuated interaction require the same *N*).

310

PART 2. How to Achieve Sufficient Power to Detect Interactions

2.1. A Preregistered Meta-Study on Power Analysis Practices

In this part, we describe the findings from a preregistered meta-study that examined the power analysis and research practices used in testing interaction hypotheses. We aimed to build a sample of relevant studies from approximately 100 articles⁴ published in 10 influential psychology journals. We sought to answer three research questions: RQ₁. What is the proportion of studies testing attenuated interactions (the most difficult to detect)? RQ₂. What is the proportion of studies using an adequate power analysis? RQ₃. What is the median statistical power of the studies?

Method

The study was preregistered (any deviation from the preregistered protocol is noted). The eligibility criteria, the preregistration, the coding sheets, the raw dataset, and the Stata scripts to reproduce the findings can be found at <https://osf.io/xh5tc/>

Eligibility Criteria and Rationales

We focused on articles that met five eligibility criteria (for the detailed list, see **Table S1**). First, we focused on articles in which an interaction hypothesis was formulated a priori (Criterion 1), as NHST is mainly appropriate for confirmatory and not exploratory analysis (Wagenmakers et al., 2012; but see Rubin, 2017). Moreover, we focused on articles testing first-order interactions (Criterion 2), as second-order interactions correspond to a difference between differences of subdifferences and are notoriously harder to detect (Dawson & Richter, 2006). Finally, we focused on articles using a between-participants design (Criterion 3) and a regular regression framework (including ANOVA; Criterion 4) with a one-degree-of-freedom test (Criterion 5). The reason for these foci was that within-participant designs, nonlinear functions, multilevel modeling, polytomous variables, etc., all involve different formulas for statistical power calculations (see Brysbaert, 2019; Demidenko, 2008; Domingue et al., 2022; Mathieu et al., 2012). However, we will later see how mixed designs and planned contrast analysis can be used to increase power.

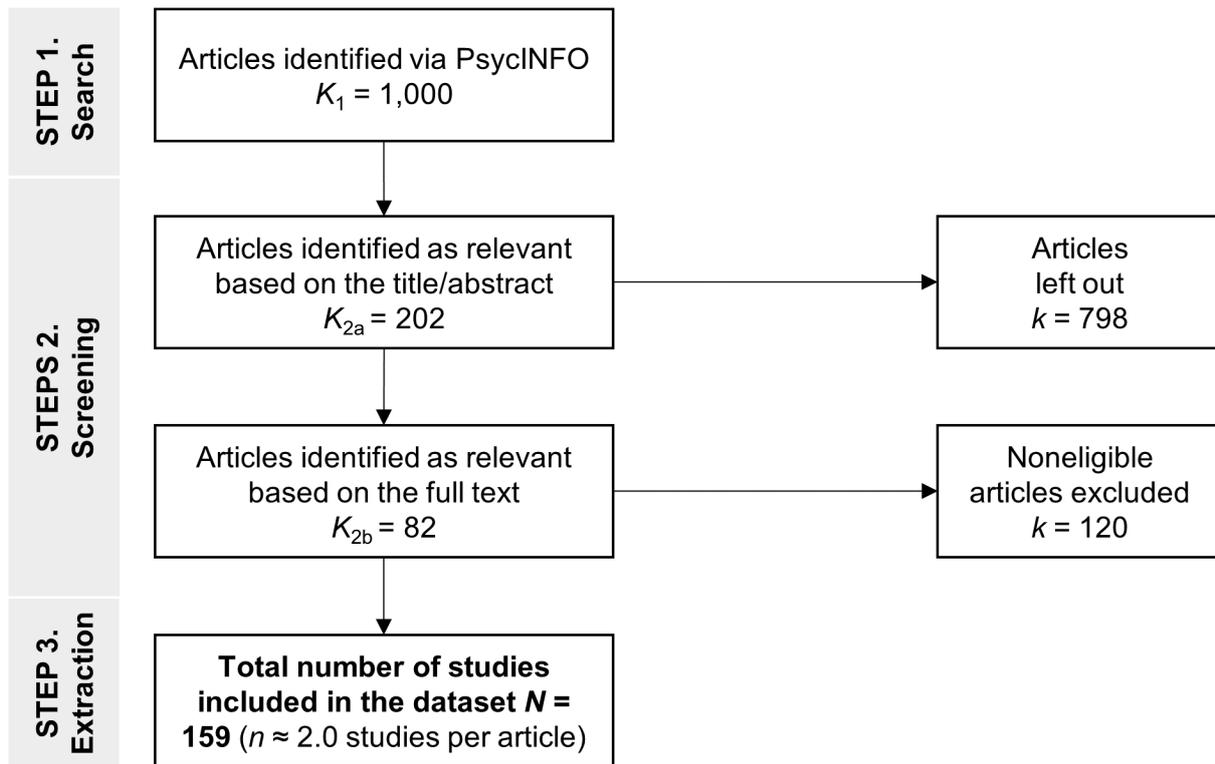
Article Selection and Data Extraction

Our study procedure consisted of three steps (for a flowchart detailing these steps and the number of articles and studies, see **Figure 2**).

⁴ The target number of articles was primarily based on the funds available for paying coders. However, the analytical sample ($N = 154$ studies) was deemed representative of the contemporary psychological literature testing interaction hypotheses. Because our meta-study is merely descriptive, power/sensitivity analysis does not apply to this case.

Figure 2.

Flowchart depicting the three steps of the article selection and data extraction procedure.



345

STEP #1. Search Strategy. In June 2021, we sampled articles from 10 top-tier empirical journals in general, social, and/or personality psychology (for the list of these journals and their characteristics, see **Table 3**). The median Journal Impact Factor Percentile was 85th, meaning that the typical journal in our dataset was in the top 15% of its category (Clarivate Analytics, 2021). For each journal, we used PsycINFO to identify the 100 most recent articles containing the words “moderat*” or “interact*” (in any field). At the end of Step #1, our sample comprised $K_1 = 10$ (journals) \times 100 (articles) = 1,000 articles.

350

STEP #2a. Screening Titles/Abstracts. In July 2021, we gave two coders the titles and abstracts of the $K_1 = 1,000$ articles. The coders were unaware of the specific purposes of the study. For each journal, they were asked to identify the 20 most recent papers that *likely* satisfied our five eligibility criteria. After training with articles from *Social Psychological and Personality Science (SPPS)*, they coded the remaining articles independently from one another using the following coding scheme: 0 = *nonrelevant article*, 1 = *potentially relevant article*. The inter-rater agreement, measured using Cohen’s Kappa, was $\kappa = .70$ (above our preregistered threshold of $\kappa = .60$), with a substantial percentage of agreement of 88.2% (Belur et al., 2021). Then, the two coders independently reviewed their disagreements and could change their responses (the

360

remaining disagreements were resolved by discussion). At the end of Step #2a, our sample comprised $K_{2a} = 202$ articles.⁵

365 **STEP #2b. Screening Full Texts.** In August 2021, we gave the two coders the full texts of the $K_2 = 202$ articles. For each journal, the coders were asked to first identify the 10 most recent papers that satisfied Criteria 2-5, and then Criterion 1. After training with the articles from *SPPS*, they coded the remaining articles independently from one another using the following coding scheme: 0 = *nonrelevant article*, 1 = *relevant article*. The inter-rater agreement was $\kappa = .79$, with a percentage of agreement of 90.7%. Disagreements were resolved by discussion. At the end of Step #2b, our sample comprised $K_{2b} = 82$ articles.

370 **STEP #3. Data Extraction.** In September 2021, we gave the two coders a data extraction spreadsheet (**Table S2**). For each of the $K_{2b} = 82$ articles, the coders were asked to: (i) identify the interaction hypothesis/es and specify its/their types (reversed, fully attenuated, or partially attenuated [RQ₁]), (ii) identify the power analysis/es and specify its/their characteristics (e.g., type, focus, whether the shape of the interaction was taken into account [RQ₂]), and (iii) collect the analytical sample size (for us to calculate power [RQ₃]). After training with three articles from *Psychological Science*, they completed the spreadsheet independently from one another. The mean inter-rater agreement for the categorical/numeric responses was $\bar{\kappa} = .87$, with an overall percentage of agreement for all questions of 82.0%. At the end of Step #3, our sample 380 comprised $N = 159$ studies from $K = 82$ articles (≈ 2 studies per article). The sample size was somewhat below our target number of 100 articles. However, the articles were deemed representative of the current literature in that they covered a large proportion, if not most, of the target population of studies published in the 10 chosen journals between 2017 and 2021.

⁵ The total number of studies was different from 10 [journals] \times 20 [articles] = 200 articles, and varied from one journal to another (see **Table 3**) because the coders were sometimes unable to identify 20 articles or—following the resolution of disagreements—ended up identifying more. The same applied to the number of studies identified in Step #2b.

385 **Table 3.**

List of the 10 journals, with their Journal Impact Factor Percentile, the number of articles *k* at each intermediary step (Step #1. Search Strategy; Step #2a. Screening abstract and title; Step #2b. Data extraction), and the number of articles *k* and studies *n* at the end of the final step (Step #3. Data extraction).

	Years	JIF perc.	Step #1. <i>k</i>	Step #2a. <i>k</i>	Step #2b. <i>k</i>	Step #3. <i>n</i>
<i>JPSP</i>	2021-2022	94 th	100	20	9	21
<i>PS</i>	2019-2021	93 rd	100	20	6	18
<i>SPPS</i>	2019-2021	92 nd	100	20	9	12
<i>EJP</i>	2017-2019	91 st	100	11	3	8
<i>JP</i>	2018-2020	85 th	100	21	10	14
<i>JESP</i>	2020-2021	85 th	100	23	9	20
<i>JEP:G</i>	2020-2021	82 nd	100	26	6	8
<i>PSPB</i>	2021	84 th	100	21	11	21
<i>BJSP</i>	2020-2021	80 th	100	20	9	15
<i>EJSP</i>	2019-2021	73 rd	100	20	10	22
Total	2017-2022	Median = 85	<i>K</i> = 1,000	<i>K</i> = 202	<i>K</i> = 82	<i>N</i> = 159

390 Note: JIF perc. = Journal Impact Factor Percentile (JIF perc. transforms the rank in category by
 JIF into a percentile value, thereby allowing us to take the category of the journal into account);
JPSP = *Journal of Personality and Social Psychology*; *PS* = *Psychological Science*; *SPPS* =
 395 *Social Psychological and Personality Science*; *EJP* = *European Journal of Personality*; *JP* =
Journal of Personality; *JESP* = *Journal of Experimental Social Psychology*; *JEP:G* = *Journal*
of Experimental Psychology: General; *PSPB* = *Personality and Social Psychology Bulletin*;
BJSP = *British Journal of Social Psychology*; *EJSP* = *European Journal of Social Psychology*.

Results

RQ1. What is the Proportion of Studies Testing Attenuated Interactions?

400 **Short Answer.** About 85% of the studies tested attenuated interaction hypotheses.

Long Answer. The 159 studies tested a total of 194 interaction hypotheses. Seven studies were excluded from the analysis because the coders could not determine the shape of the hypothesized interactions. Among the 152 remaining studies, only 15% tested at least one *reversed* interaction hypothesis, whereas 36% tested at least one *fully* attenuated interaction, and
405 52% tested at least one *partially* attenuated interaction (**Figure 3**, left panel). Note that the percentages do not add up to 100% because one study could test multiple types of interaction hypotheses.

RQ2. What is the Proportion of Studies Using an Adequate Power Analysis?

Short Answer. Less than 5% of the studies.

410 **Long Answer.** Out of the 159 studies, 45% did not report a power/sensitivity analysis—specifically, $n = 65$ did not report any analysis, $n = 3$ did not specify the type of analysis, and $n = 3$ reported an incorrect “post-hoc” power analysis (see Hoenig & Heisey, 2001). Another 20% reported a power/sensitivity analysis *not* focused on the interaction—specifically, $n = 23$ did not specify the focal effect and $n = 9$ focused on a main effect or another type of effect. Finally,
415 another 32% reported a power or sensitivity analysis focused on the interaction, but did *not* take the shape of the interaction into account. In sum, only 4% of the studies reported an adequate power analysis (**Figure 3**, right panel).

RQ3. What Is the Median Statistical Power of the Studies?

Short Answer. The overall median power to detect interactions of a typical size is .18.

420 **Long Answer.** For each hypothesis of each study, we used Eqs. 5 and A₄ to calculate the power afforded by the analytical sample size to detect the *typical* version of the interaction expected by the authors.⁶ For the studies testing a reversed interaction ($n = 23$)⁷, the power to detect a +0.35 | -0.35 reversed interaction was at or above .80 in 65% of the cases and the median was $\overline{power} = .87$. For the studies testing a fully attenuated interaction ($n = 54$), the
425 power to detect a +0.35 | 0.00 fully attenuated interaction was at or above .80 in 19% of the cases and the median was $\overline{power} = .36$. For the studies testing a partially attenuated interaction

⁶ In the preregistration, we stated that we would use the power estimates derived from our simulations (rather than from mathematical formulas). Both approaches led to the exact same conclusions.

⁷ When a study tested interaction hypotheses having different shapes (e.g., a reversed interaction and a partially attenuated interaction), we focused on the largest power estimate.

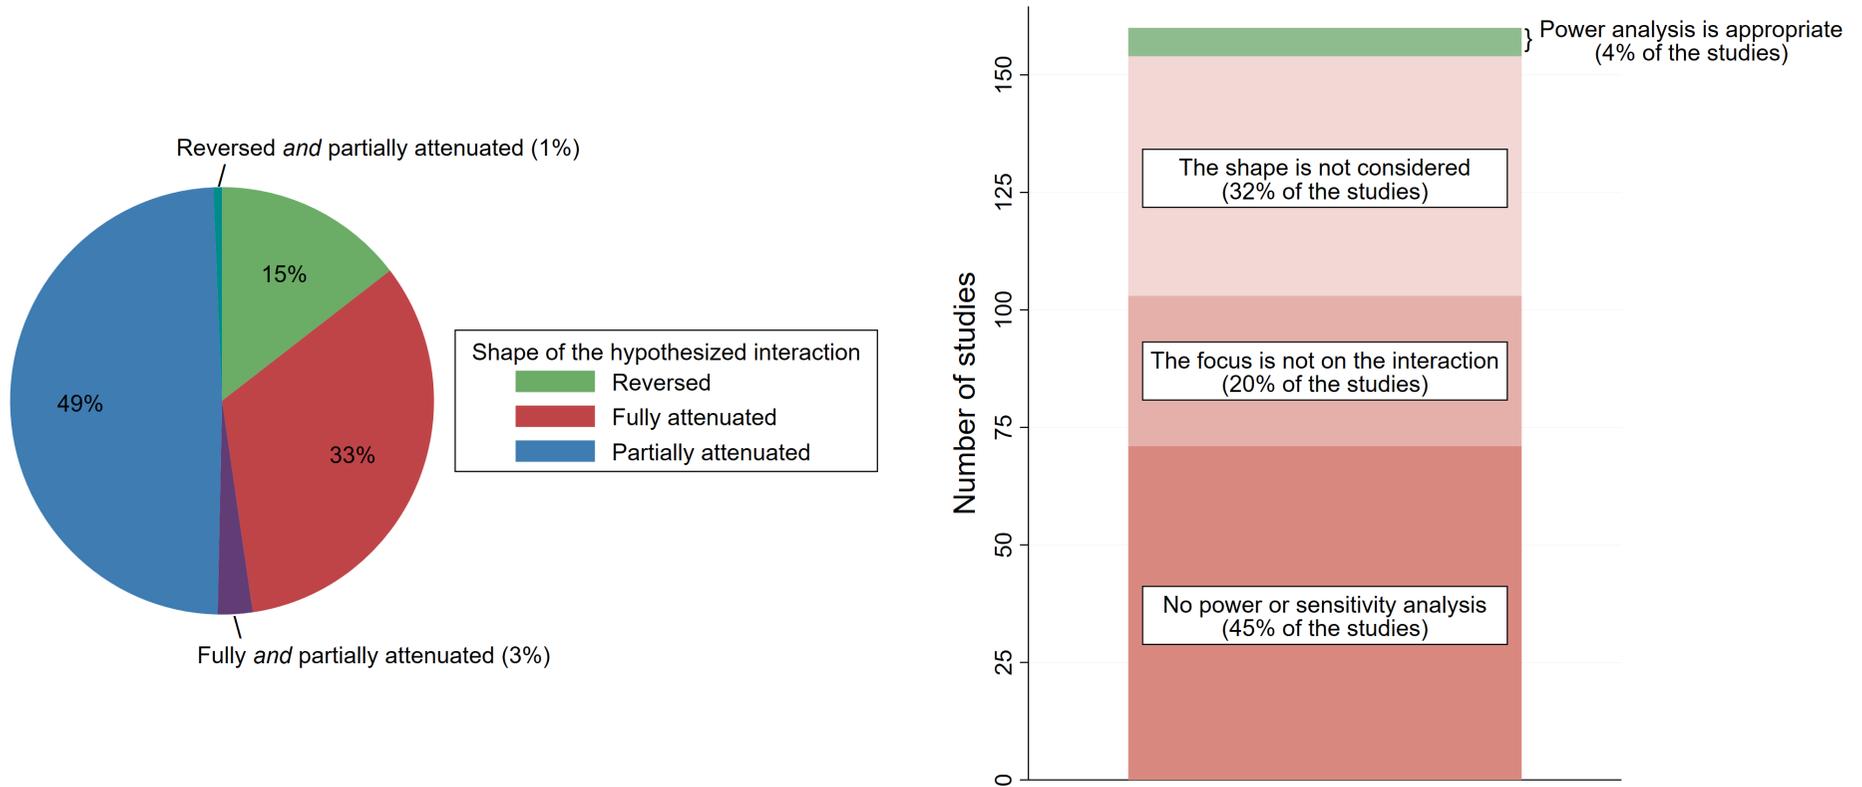
($n = 74$), the power to detect a $+0.35 \mid +0.20$ (or, equivalently, a $+0.35 \mid +0.50$) partially attenuated interaction was at or above .80 in *none* of the cases and the median was $\overline{power} = .11$. Overall, only 17% of the studies had a power at or above .80 and the median was $\overline{power} = .18$ (Figure 4, lower panel). If we repeat the analysis while focusing on larger versions of the hypothesized interactions (i.e., $+0.50 \mid -0.50$, $+0.50 \mid 0.00$, and $+0.50 \mid +0.20$), 27% of the studies had a power at or above .80, and the median was $\overline{power} = .25$.

Discussion

Our meta-study shows that the majority of recently published studies testing interactions focus on attenuated interactions. Only a minority reported an adequate power analysis, and less than one in five studies has a sample size sufficient to detect an interaction of a typical size with a power of .80. This illustrates the fact that studies testing interactions are often underpowered and, in the next part, we present simulations aiming to test three approaches to increase power without increasing sample size.

440 **Figure 3**

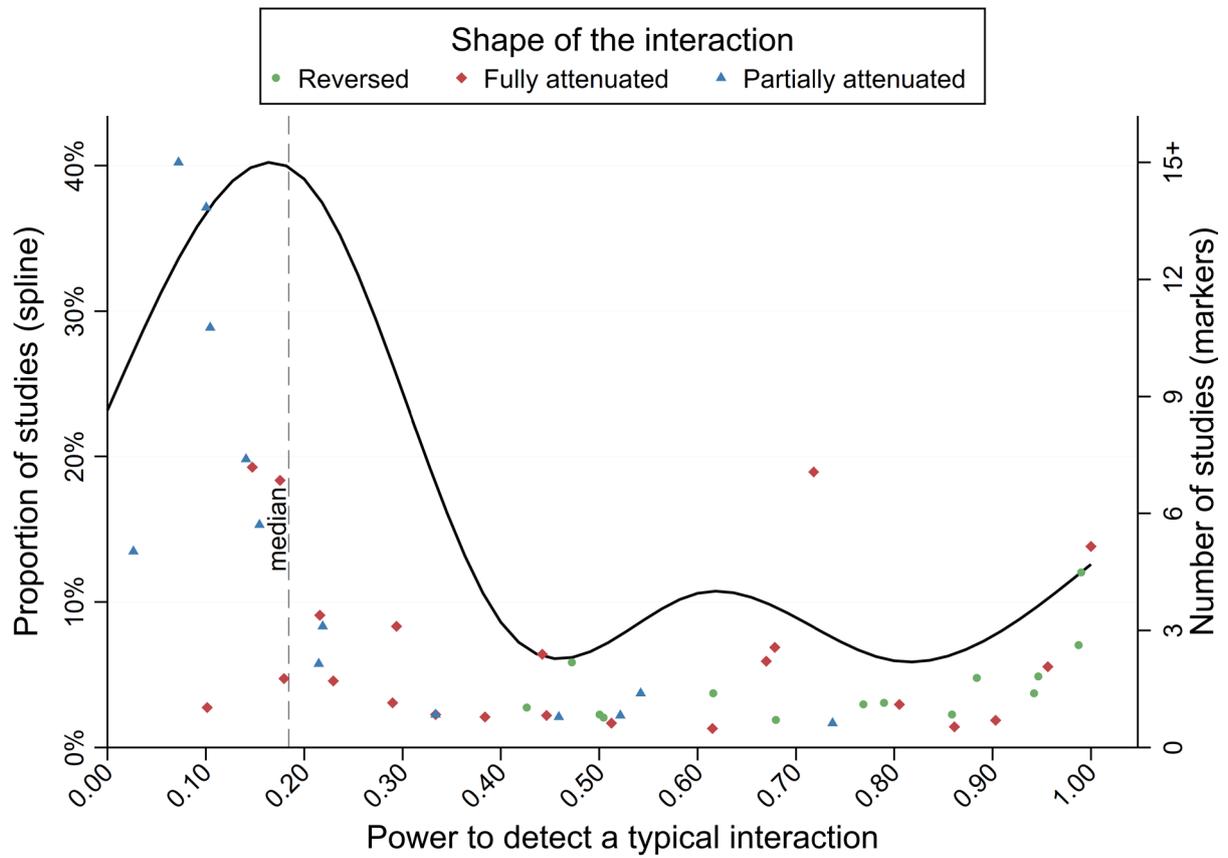
Proportion of studies as a function of the shape of the hypothesized interaction (pie chart, left panel) and number of studies as a function of the appropriateness of the power analysis (stacked bar chart, right panel).



Note: In the pie chart, the overlapping slices pertain to studies that include interaction hypotheses of different shapes.

445 **Figure 4**

Proportion of studies as a function of the power to detect the hypothesized interaction, assuming a typical size.



Note: We added spherical noise to the plotted markers to avoid overlap.

2.2. Simulations Testing Ways to Improve Power for Interactions

Now, we report the results from 877,500,000 simulations that served two aims: (1) producing the power curve for each of the 12 types of interactions (offering an empirical replication of the mathematically derived required sample size displayed in **Table 2**) and (2) testing three approaches to increase power without increasing sample size: one-tailed testing, mixed designs, and planned contrast analysis.

Existing Simulations

Several studies have used simulations to investigate the question of power when testing interactions (for pioneering work, see Champoux & Peters, 1987). In existing simulations, however, the size of the interaction in the population is often fixed across conditions (for exceptions, see Durand, 2013; Shieh, 2009), and the emphasis is put on important but very specific issues such as the use of median split in multicollinearity contexts (Iacobucci et al., 2015; see also McClelland et al., 2015), the calculation of the product term in latent variable modeling (Chin et al., 2003; see also Goodhue et al., 2007), or the violation of the homoscedasticity assumption when sample sizes differ among subgroups (Alexander & DeShon, 1994; see also Aguinis & Stone-Romero, 1997). In our own simulations, we aimed to produce the power curve of the 12 interactions identified earlier, and determine the extent to which these power curves could be flattened by the use of one-tailed testing, mixed designs, and planned contrast analysis.

Approach #1. Preregistering the Use of One-Tailed Testing

Most scientists using NHST use two-tailed tests (also called two-sided or nondirectional tests) with a canonical alpha of .05. This means that they place one half of their alpha ($\alpha/2 = .025$) in the lower tail of the t distribution (testing whether an effect is less than zero) and the other half ($\alpha/2 = .025$) in the upper tail (testing whether an effect is greater than zero; Wiley & Pace, 2015). To put it simply, if a scientist is looking for an effect that does *not* exist in the population, they have a 2.5% chance of observing a negative effect with $p < .05$ and a 2.5% chance of observing a positive effect with $p < .05$ (the overall type I error [false positive] rate of 5%).

However, scientists using NHST often have an a priori hypothesis and could run a more powerful *one-tailed test* (also called one-sided or directional tests) with the same alpha of .05 (Knottnerus & Bouter, 2001). This involves placing *all* of the alpha ($\alpha = .05$) in the tail of the t distribution corresponding to their hypothesis (the lower *or* upper tail, depending on whether they hypothesized a negative or positive effect), and refrain from interpreting any effect in the opposite direction (even if $p < .05$). In simpler terms, if a hypothesis is incorrect, there is a 5%

chance of observing the hypothesized effect with $p < .05$ (maintaining an overall type I error rate of 5%).

Statisticians have long recommended using one-tailed testing when a directional hypothesis is formulated and when an effect in the opposite direction would be psychologically meaningless and theoretically uninformative (Kimmel, 1957). Despite these conditions often being met, journal editors and reviewers tend to frown upon one-tailed testing. This reluctance may be due to the high prevalence of HARKing in the field (Motyl et al., 2017), which involves formulating a hypothesis *after* results are known (Kerr, 1998) and is likely influenced by hindsight bias (i.e., the “I-knew-it-all-along” effect; Nosek et al., 2018; see also Giner-Sorolla, 2012; Rubin, 2020). In such a context, one-tailed testing may be seen as overly liberal and a threat to cumulative science.

However, with the advent of preregistration (Van't Veer & Giner-Sorolla, 2016; Wagenmakers et al., 2012) and registered reports (Chambers, 2013), it has become possible for researchers to record their hypothesis *before* running a study, thus effectively preventing HARKing (Lakens, 2019). We believe that, *whenever possible*, researchers should create a preregistration in which they (i) clearly formulate their interaction hypothesis, (ii) vow to not interpret an interaction in the opposite direction, (iii) plan to use one-tailed testing. We emphatically encourage this practice, believing it is both conceptually and pragmatically the best approach. A researcher registering an interaction hypothesis and the use of one-tailed rather than two-tailed testing (with $\alpha = .05$) would need 21% less participants to reach a power of .80 (for the mathematical demonstration, see **Appendix**, Eqs. A_{5a-c}).

Approach #2. Using Mixed-Participants Designs

Despite the difficulty in comparing and calculating effect sizes across different designs (Lakens, 2013; Morris, 2008; Olejnik & Algina, 2003), combining between-participant and within-participant measures is usually considered an efficient way to increase statistical power (Lakens, 2022). This is because each participant from a study using two repeated measures does not provide one but two data points, reducing the error term by controlling for consistent individual differences (for related research, see Baker et al., 2021; Goulet & Cousineau, 2019). As a result, fewer participants are needed to detect the same type of effect while maintaining the same level of power, especially when the correlation between the within-participant measures is positive and large (Maxwell & Delaney, 2004).

Mixed-participant designs may have potential drawbacks. Participants may be more likely to discern the hypothesis, experience fatigue due to the increased length of the study, or

be influenced by the order of presentation of the scales/conditions (Myers & Hansen, 2011). However, these drawbacks are arguably often outweighed by the benefits of increased statistical power. To illustrate, imagine a mixed-design study in which (i) the correlation between the measurements is $\rho = .50$ (a conservative estimate and less than the average correlation from existing [replication] studies; Brysbaert, 2019) and (ii) the simple slope sizes of the within-participant variable is similar to what it would be in a between-participants design (despite the tendency for within-participants variables to be stronger; e.g., see Murphy et al., 2009). Given these assumptions, a researcher testing an interaction and using a mixed design rather than a 2 × 2 between-participants design (with $\alpha = .05$) would need 75% fewer participants to reach a power of .80 (for the mathematical demonstration, see **Appendix**, Eq. A6_{a-f}). If they combined a mixed design with preregistered one-tailed testing, the researcher would need 80% fewer participants $(1 - (1 - .75) \times (1 - .21))$.

Approach #3. Preregistering the Use of Planned Contrast Analysis

The factorial approach is the default approach to testing interaction hypotheses. It involves regressing the outcome on the predictor, the moderator, and the product term, whose weights are as follows for the example used in PART 1:

	Control group		Intervention group	
	Men	Women	Men	Women
Predictor (-1 = control, +1 = intervention)	-1	-1	+1	+1
Moderator (-1 = men, +1 = women)	-1	+1	-1	+1
Predictor × Moderator	+1	-1	-1	+1

As one can see, the weights of the product term form a cross $(\begin{smallmatrix} +1 & -1 \\ -1 & +1 \end{smallmatrix})$, which makes it optimal for testing reversed interactions involving a negative simple slope for men $(\begin{smallmatrix} +1 \\ -1 \end{smallmatrix})$ and a positive simple slope for women $(\begin{smallmatrix} -1 \\ +1 \end{smallmatrix})$. However, these weights are suboptimal for testing other shapes of interaction, in particular, fully attenuated interactions.

An alternative approach is *the planned contrast approach* (Rosenthal & Rosnow, 1985). In the case of a fully attenuated interaction, this approach involves concatenating (i.e., combining) the predictor and moderator variables, and creating one planned and two orthogonal contrasts using Helmert coding (Rosnow & Rosenthal, 1991). The contrast weights could be as follows for the example used in PART 1⁸.

⁸ Other weights are possible, provided that (i) the sum of the weights for each contrast is zero, and (ii) the sum of the products of the weights for each pair of contrasts is zero. Note that we used integers for

	Control group		Intervention group	
	Men	Women	Men	Women
Planned contrast	-1	-1	-1	+3
Orthogonal contrast #1	-1	+2	-1	0
Orthogonal contrast #2	-1	0	+1	0

As one can see, the weights of the planned contrast form a left-angled triangle $(-1, -1, +3)$, corresponding to hypothesized pattern and comparing women in the intervention group (+3) to participants in the three other subgroups (-1). The two orthogonal contrasts ensure that there is no residual difference between women in the control group and men (Orthogonal contrast #1) or between men in the control group and men in the intervention group (Orthogonal contrast #2).

To reject H_0 , a significant planned contrast *and* nonsignificant orthogonal contrasts are needed (a LR χ^2 could test for the joint significance of the orthogonal contrasts; Abelson & Prentice, 1997; Brauer & McClelland, 2005; Guggenmos et al., 2018). However, absence of evidence is not evidence of absence, and cautious analysts might want to use equivalence testing to ensure that the overall effect of the orthogonal contrasts is smaller than the smallest effect of interest (Lakens et al., 2018; Richter, 2016).

Importantly, some statisticians argue that planned contrast analysis offers excessive flexibility in data analysis (Ravenscroft & Buckless, 2017), and may not be suitable to test specific interaction patterns (see the debate between Abelson [1996] and Rosnow & Rosenthal [1996]). Thus, we recommend that authors preregister the use of contrast analysis, and use it only for testing fully attenuated interactions in which three of the 2×2 means are expected to be the same. In such a case, a researcher using planned contrast analysis with Helmert coding rather than the orthodox factorial approach (with $\alpha = .05$) would need 62% less participants to reach a power of .80 (for relevant work, see Perugini et al., 2018; for the mathematical demonstration, see **Appendix**, Eqs. A7a-d). If they combined planned contrast analysis with one-tailed testing, the researcher would need 70% fewer participants $(1 - (1 - .62) \times (1 - .21))$.

Simulations Testing the Three Approaches to Increase Power Without Increasing the N

Method

We used functions from the *stats* and *car* packages in R (Fox & Weisberg, 2018; R Core

clarity. However, as in our simulation, it is better to use a $-0.5 \mid +0.5$ coding scheme (for the factorial approach), and $^{-1/4 \mid -1/4 \mid -1/4 \mid +3/4}$, $^{-1/3 \mid +2/3 \mid -1/3 \mid 0}$, and $^{+1/2 \mid 0 \mid +1/2 \mid 0}$ contrasts (for the planned contrast approach) to estimate meaningful one-unit differences.

Team, 2013) to simulate a total of 877,500,000 datasets with one continuous outcome variable, one categorical predictor, and one categorical moderator, and generated power curves for the 12 interactions of our taxonomy while using (i) default conditions (between-participants design and a factorial approach) with two-tailed and one-tailed testing, (ii) a mixed-participants design
 570 with two-tailed and one-tailed testing, and (iii) planned contrast analysis with two-tailed and one-tailed testing. The R scripts to reproduce the results are available at <https://osf.io/xh5tc/>

For each interaction, we simulated 100K datasets with a sample size of n per condition, and we used an increment of $n + 1$ to identify the tipping points at which 80% and 90% of datasets returned significant tests with $p < .05$ (e.g., simulating 100K datasets with $n = 100$,
 575 another 100K with $n = 101$, another 100K with $n = 102$, and so on). In each case, this enabled us to identify the required Ns to achieve a power of .80 and .90, respectively (i.e., the most commonly used lower limits for acceptable power).

The population simple slope sizes for the 12 interactions were $d_1 | d_2 = +0.20 | -0.20, +0.20 | -0.35, +0.20 | -0.50, +0.35 | -0.35, +0.35 | -0.50, +0.50 | -0.50, +0.20 | 0.00, +0.35 | 0.00,$
 580 $+0.50 | 0.00, +0.20 | +0.35, +0.20 | +0.50, \text{ and } +0.35 | +0.50$. To enable direct comparison across the four approaches, each dataset was first sampled from a standard multivariate normal distribution and then adjusted using the mean vector $\mu = [0 \ 0 \ d_1 \ d_2]'$ and the covariance matrices Σ relevant to each condition:

$$\begin{array}{cc}
 \textit{Fully between design} & \textit{Mixed design} \\
 \Sigma = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} & \Sigma = \begin{bmatrix} 1 & 0 & \rho & 0 \\ 0 & 1 & 0 & \rho \\ \rho & 0 & 1 & 0 \\ 0 & \rho & 0 & 1 \end{bmatrix} \quad (9)
 \end{array}$$

...where within-subject correlations in the mixed design were set to $\rho = .50$, a conservative
 585 estimate that was well below the median value of $\rho = .75$ from replication studies (Brybaert, 2019).⁹

In the *default approach*, for each $d_1 | d_2$ combination and each n , we calculated the proportion of the 100K datasets for which a regression analysis using a factorial approach returned a significant interaction between the between-participants predictor and moderator. We re-
 590 peated this for both two-tailed and one-tailed tests.

⁹ For example, when $d_1 | d_2 = +0.50 | 0.00$ and $n = 125$, we (1) simulated 10K datasets of size 125×4 from a standard multivariate normal distribution, (2) adjusted each dataset by the mean vector μ and the covariance matrices Σ described above, (3) conducted the relevant significance test for each approach using the adjusted datasets, and (4) calculated the proportion of significant interactions from the 100K datasets for each approach.

In the *mixed designs approach*, for each $d_1 | d_2$ combination and each N , we calculated the proportion of the 100K datasets for which a mixed analysis of variance returned a significant interaction between the between-participants predictor and the within-participants moderator. We repeated this for both two-tailed and one-tailed tests.

595 In the *planned contrast analysis approach*, we focused only on fully attenuated interactions, and we created a planned contrast (assigning weights of $^{-1/4} | ^{-1/4} | ^{+3/4} | ^{-1/4}$ to the 2×2 cells of the interaction) and two orthogonal contrasts (using weights of $^{-1/3} | ^{-1/3} | ^{-2/3} | 0$ and $^{-1/2} | ^{+1/2} | 0 | 0$). For each of the three $d_1 | d_2$ combinations and each N , we calculated the proportion of the 100K datasets for which a regression analysis returned a significant planned contrast and a non-
600 significant joint test of the orthogonal contrasts (using an omnibus post-estimation two-tailed Wald test). We repeated this for both two-tailed and one-tailed tests.

Results

The simulation revealed two set of findings. First, the mathematically derived required sample sizes for each of the 12 interactions calculated in PART 1 were replicated using simula-
605 tions: The required N s to detect the interactions with a power of .80 ranged from 128 (to detect a $+0.50 | -0.50$ reversed interaction) to 5,632 (to detect a $+0.35 | +0.50$ partially attenuated interaction. Second, compared to the power curves of the default approach with two-tailed testing: (a) preregistering one-tailed tests requires 21.3% and 18.5% fewer participants to reach a power of .80 and .90, respectively, (b) using a mixed design with $\rho = .50$ requires 74.6% and
610 74.7% fewer participants when using two-tailed testing, and 80.1% and 79.4% when using one-tailed testing, and (c) preregistering planned contrast analysis requires 62.6% and 59.4% fewer participants when using two-tailed testing, and 70.1% and 66.3% when using one-tailed testing (see **Table 4** for the required N s and **Figure 5** for the simulation-based power curves).

Other Ways to Maximize Power

615 In our simulation, we have chosen to focus on one-tailed testing, mixed designs, and planned contrast analysis because they present minimal potential drawbacks. Although there are other approaches to increase power without increasing sample size, they often carry greater drawbacks. For instance, increasing sample homogeneity (Heidel, 2016), using instructional manipulation checks (Oppenheimer et al., 2009), or controlling for relevant covariates (Her-
620 nández et al., 2004) can enhance power, but may also pose a threat to generalizability, alienate participants, or create spurious suppression effects, respectively.

However, as briefly mentioned in PART 1, an approach to maximizing power with little

or no drawbacks is to increase the reliability of the outcome. Basically, when reliability increases, the effect size increases, resulting in increased power (implying that if your outcome is measured with higher-than-average reliability, the benchmarks used in this paper may be underestimations). Heo et al. (2015) demonstrated that by increasing Cronbach's alpha from the minimum acceptable level of .70 to a high value of .90, power to detect an interaction in a mixed design can increase from approximately .40 to almost .90. Consequently, improving reliability by using validated scales instead of made-up scales, using multiple-item measures instead of single-item measures, or using multiple trials rather than just one can each be an extremely efficient way to increase power for detecting interactions.

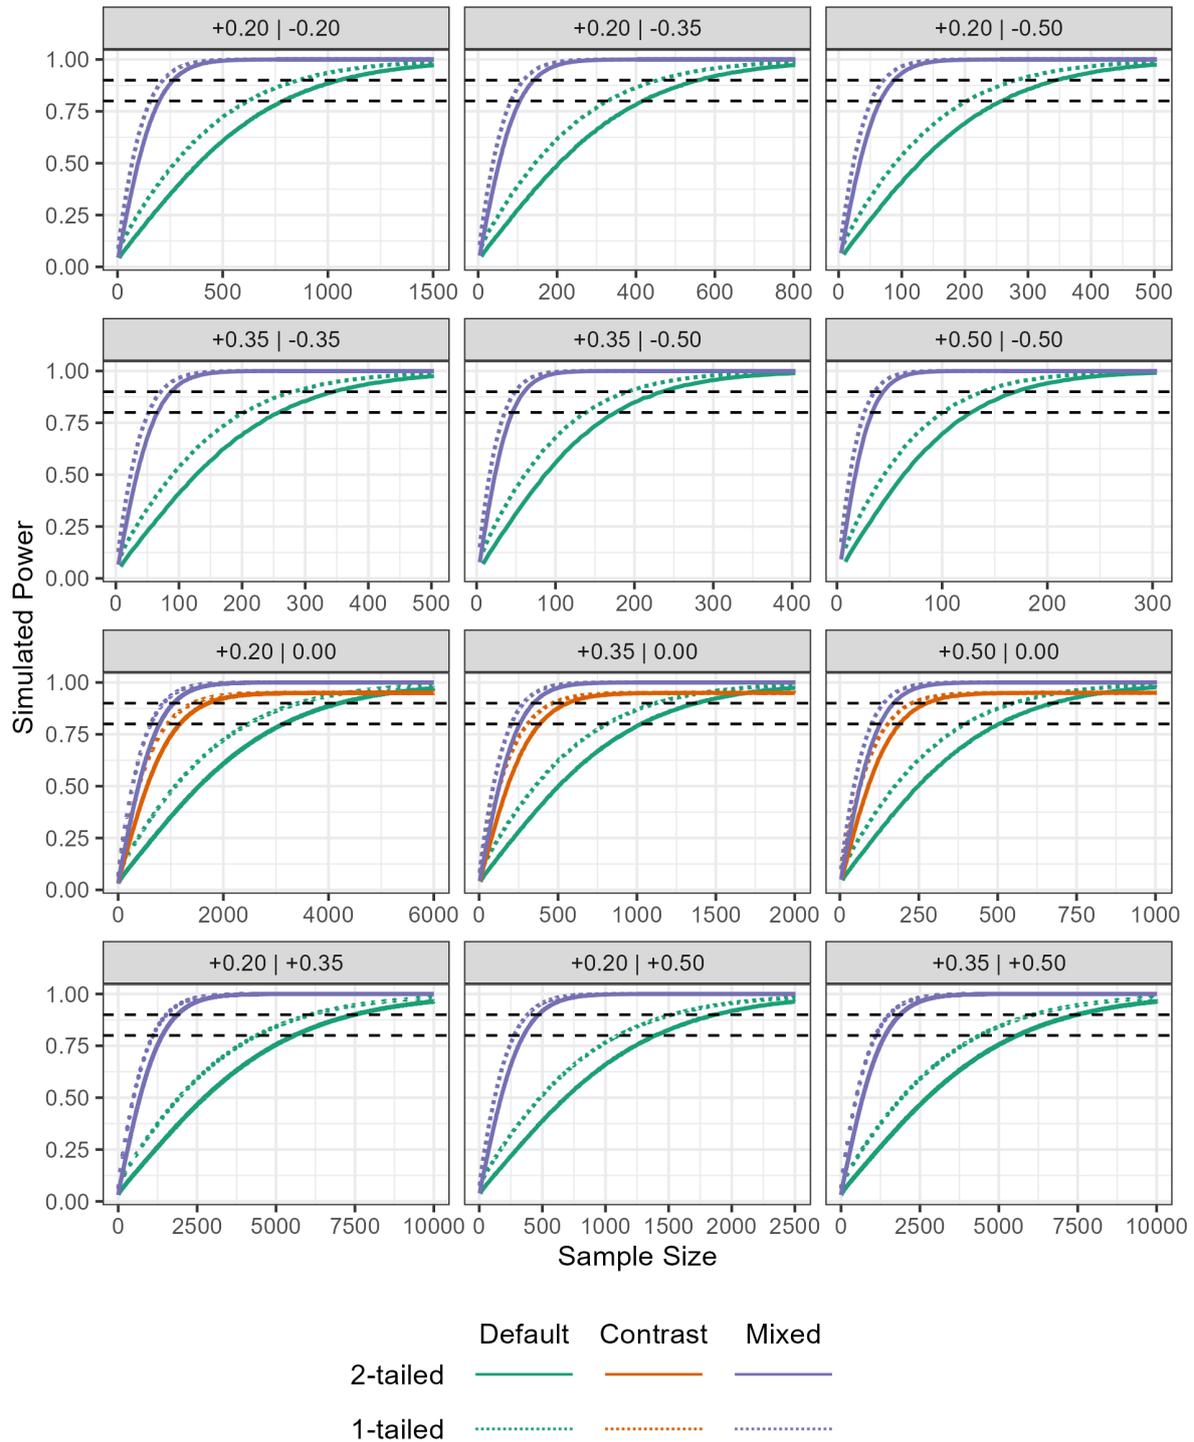
Table 4

635 Overall required sample sizes for the 12 types of interactions as a function of approach.

	Simple Slope #1	Simple Slope #2	Default		Mixed Design		Planned Contrast	
			Two-tailed	One-tailed	Two-tailed	One-tailed	Two-tailed	One-tailed
<i>Power = .80</i>	+0.20	-0.20	788	620	198	154		
	+0.20	-0.35	416	328	106	84		
	+0.20	-0.50	260	204	66	52		
	+0.35	-0.35	260	204	66	52		
	+0.35	-0.50	176	140	46	36		
	+0.50	-0.50	128	100	34	26		
	+0.20	0.00	3,124	2,476	786	624	1,172	940
	+0.35	0.00	1,028	812	258	202	384	304
	+0.50	0.00	504	396	128	100	188	152
	+0.20	+0.35	5,588	4,384	1,390	1,096		
	+0.20	+0.50	1,396	1,100	352	276		
	+0.35	+0.50	5,632	4,412	1,392	1,096		
<i>Power = .90</i>	+0.20	-0.20	1,056	860	266	216		
	+0.20	-0.35	556	456	142	114		
	+0.20	-0.50	344	280	88	72		
	+0.35	-0.35	344	280	88	72		
	+0.35	-0.50	236	192	60	48		
	+0.50	-0.50	172	140	44	36		
	+0.20	0.00	4,224	3,436	1,060	854	1,716	1,424
	+0.35	0.00	1,376	1,120	344	284	556	464
	+0.50	0.00	676	548	170	138	276	228
	+0.20	+0.35	7,448	6,088	1,868	1,524		
	+0.20	+0.50	1,856	1,524	468	382		
	+0.35	+0.50	7,516	6,108	1,866	1,518		

Figure 5

Simulation-based power curves for the 12 types of interactions as a function of approach.



640 *Note:* Power curves for the second half of the Mixed Design approach were extrapolated.

PART 3. A User-Friendly Web App to Conduct Power Analyses for Interactions

There are countless existing software and applications to run power analyses. Some are general and focus on the most common statistical tests (e.g., [G*Power](#); Faul et al., 2007; [PAN-GEA](#); Bartlett & Charles, 2021; [the jpower module in Jamovi](#); Bartlett & Charles, 2021). Others
645 are specific to certain statistical tests, such as multilevel regression ([summary-statistics-based power for mixed-effects modelling](#); Murayama et al., 2022), structural equation modeling ([pwrSEM](#); Wang & Rhemtulla, 2021), and fMRI study analyses ([NeuroPower](#); Durnez et al., 2016).

Some applications enable one to run power analyses for interactions such as [Interaction-Power](#) (Baranger et al., 2022), [Superpower](#) (Lakens & Caldwell, 2021), or [Power Analysis for 2x2 Factorial Interaction](#) (White II, 2018). Although these applications are useful, they may not
650 always be intuitive. For instance, *InteractionPower* requires users to input the effect size of the interaction term as a correlation coefficient. While some researchers will manage to calculate the expected size of their interaction correctly (e.g., using Eqs. 5 and A_2), others may use generic benchmarks and enter an inappropriate value. *Superpower* may be more intuitive as it asks
655 users to input the most common standard deviation and the raw means for each cell in the ANOVA. However, an even simpler approach would be to allow users to draw their interaction. Therefore, we have developed an application to run power analyses for interactions that only requires basic statistical knowledge and relies on a user-friendly graphical interface.

660 **Introducing INT×Power**

INT×Power is a user-friendly web app that enables researchers to easily draw their expected two-way interaction and produce the required overall sample size to achieve a power of .80. Additionally, INT×Power enables researchers to employ approaches to maximize power (one-tailed testing, mixed designs, and planned contrast analysis).

665 INT ×Power is a JavaScript web-app created using the React framework with Material UI. All libraries used are open-source. The application is hosted on the following web page: www.intxpower.com. The source code of the application is available on GitHub at <https://github.com/ncheutin/INTxPower>. A document presenting the underlying equation can be found on the OSF.

670 **Quick Tour of the Web App**

Figure 6 presents an annotated screenshot of INT×Power, showing the key components of the app:

- ① Users manipulate the 2×2 bars of their expected interaction

② Users see the effect sizes of the simple slopes vary in real-time

675 ③ Users also see the required overall N to reach a power of .80 vary in real-time

④ Users can change the default target power of .80 for any value

⑤ Users who are unsure where to start can automatically produce a bar chart representing an interaction of the expected shape and using the typical size

680 The main equation used to calculate the required overall sample size is Eq. 8, but this equation can be changed to apply to:

Ⓐ One-tailed rather than two-tailed testing,

Ⓑ Mixed-participant design rather than between-participant design, using two-tailed or one-tailed testing,

685 Ⓒ Planned contrast approach with Helmert coding rather than factorial design approach, using two-tailed or one-tailed testing (only for fully attenuated interactions),

Assumptions and Future Version(s)

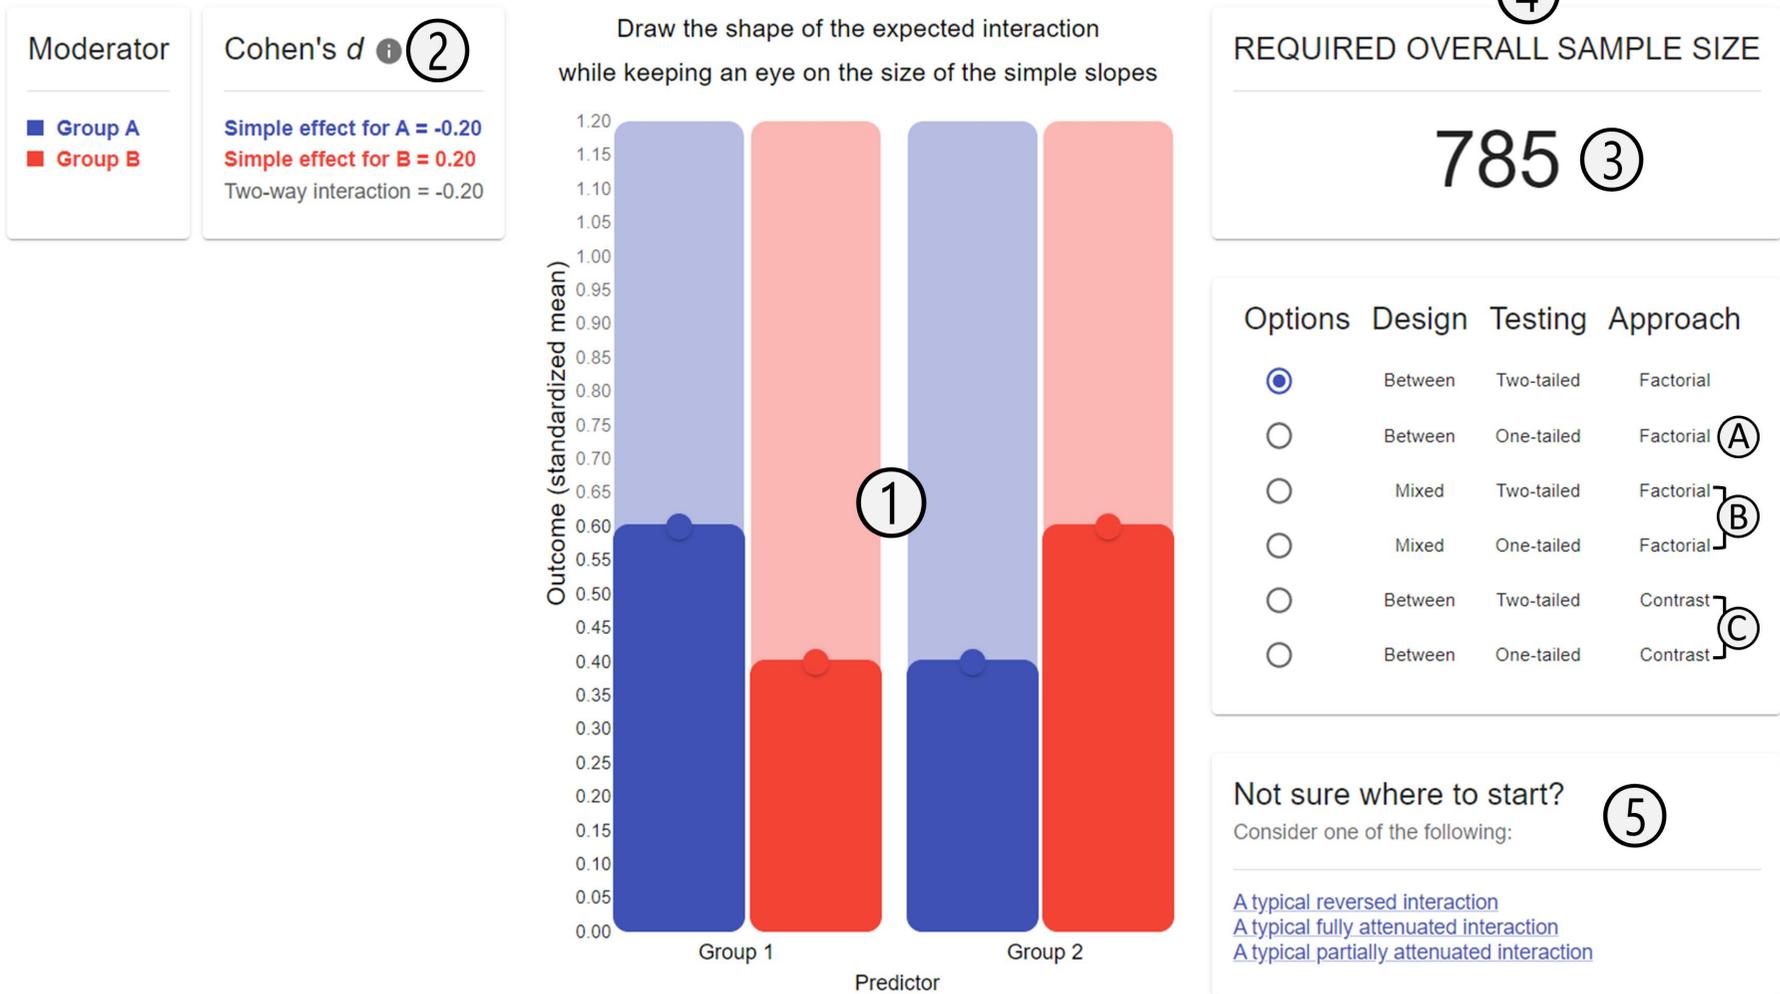
INT×Power focuses on two-way interactions and relies on the same assumptions described in PART 1, as well as sphericity and a common between-measurements correlation of $\rho = .50$ for mixed-participants designs. Although the interface uses dichotomous predictors, 690 INT×Power can be applied to continuous predictors, provided that measurement error does not differ between the two cases.

695 The version of the application is 1.0. We invite users to report bugs and request new features by sending an email to the corresponding author of this article (updates to the source code will be documented on GitHub). Future version(s) of the app may allow users to calculate power for higher-order interactions, predictor/moderator with three categories or more, nonlinear regression such as logistic or Poisson regression, etc. These changes will be the subject of further publications.

Figure 6

Annotated screenshot of INT×Power showing the key components of the Web app.

INT×Power: Finding the target sample size to detect a two-way interaction with power .80 (α = .05)



Disclosures

Author Contributions

Conceptualization: NS, DW, AE; Methodology: NS, DW, AE; Software: NC; Formal
Analysis: NS, DW; Investigation: NS; Data Curation: NS, DW, NC; Writing – Original Draft:
705 NS, DW; Writing – Review & Editing: NS, DW, AE.

Conflicts of Interest

The authors declare that they have no conflict of interest.

Acknowledgements

We thank Adrien Anex and Umberto Cauzo for their help with data collection, as well
710 as David Baranger, Marcello Gallucci, and Dominique Muller for their feedback on the paper.

Funding

This work was funded by a SNSF Ambizione fellowship granted to the first author.

Prior versions

Prior to submission, we uploaded the preprint of the article on the OSF (osf.io/xhe3u).

715 Appendix

The Appendix can be found on the OSF (<https://osf.io/xh5tc/>).

Data, materials, and online resources

The design and analysis plan for the meta-study was pre-registered. The pre-registration,
all materials, and all data were made available on the OSF: <https://osf.io/xh5tc/>

720 Reporting

We report how we determined our sample size, all data exclusions, all manipulations,
and all measures in the meta-study.

References

- Abelson, R. P. (1996). Vulnerability of contrast tests to simpler interpretations: An addendum
725 to Rosnow and Rosenthal. *Psychological Science*, 7(4), 242-246.
- Abelson, R. P., & Prentice, D. A. (1997). Contrast tests of interaction hypothesis.
Psychological Methods, 2(4), 315-328.
- Aberson, C. L. (2019). *Applied power analysis for the behavioral sciences*. Routledge.
- Abrams, L. R., & Mehta, N. K. (2019). Changes in depressive symptoms over age among
730 older Americans: differences by gender, race/ethnicity, education, and birth cohort.
SSM-Population Health, 7, 100399.
- Adachi, P., & Willoughby, T. (2015). Interpreting effect sizes when controlling for stability
effects in longitudinal autoregressive models: Implications for psychological science.
European Journal of Developmental Psychology, 12(1), 116-128.
- 735 Aguinis, H. (1995). Statistical power problems with moderated multiple regression in
management research. *Journal of Management*, 21(6), 1141-1158.
- Aguinis, H., & Gottfredson, R. K. (2010). Best-practice recommendations for estimating
interaction effects using moderated multiple regression. *Journal of Organizational
Behavior*, 31(6), 776-786.
- 740 Aguinis, H., & Stone-Romero, E. F. (1997). Methodological artifacts in moderated multiple
regression and their effects on statistical power. *Journal of Applied Psychology*, 82(1),
192.
- Aiken, L. S., West, S. G., & Reno, R. R. (1991). *Multiple regression: Testing and interpreting
interactions*. Sage.
- 745 Alexander, R. A., & DeShon, R. P. (1994). Effect of error variance heterogeneity on the
power of tests for regression slope differences. *Psychological Bulletin*, 115(2), 308.
- Altmejd, A., Dreber, A., Forsell, E., Huber, J., Imai, T., Johannesson, M., Kirchler, M., Nave,
G., & Camerer, C. (2019). Predicting the replicability of social science lab
experiments. *PLoS One*, 14(12), e0225826.
- 750 Anvari, F., & Lakens, D. (2021). Using anchor-based methods to determine the smallest effect
size of interest. *Journal of Experimental Social Psychology*, 96, 104159.
- Arend, M. G., & Schäfer, T. (2019). Statistical power in two-level models: A tutorial based on
Monte Carlo simulation. *Psychological Methods*, 24(1), 1-19.
- Baker, D. H., Vilidaite, G., Lygo, F. A., Smith, A. K., Flack, T. R., Gouws, A. D., &
755 Andrews, T. J. (2021). Power contours: Optimising sample size and precision in

experimental psychology and human neuroscience. *Psychological Methods*, 26(3), 295-314.

760 Bakker, M., Hartgerink, C. H., Wicherts, J. M., & van der Maas, H. L. (2016). Researchers' intuitions about power in psychological research. *Psychological Science*, 27(8), 1069-1077.

Baranger, D., Finsaas, M., Goldstein, B., Vize, C., Lynam, D., & Olino, T. (2022). InteractionPowerR: Power analyses for interaction effects in cross-sectional regressions. <https://doi.org/10.31234/osf.io/5ptd7>.

765 Bartlett, J. E., & Charles, S. (2021). Power to the people: A beginner's tutorial to power analysis using Jamovi. <https://doi.org/10.31234/osf.io/bh8m9>.

Belur, J., Tompson, L., Thornton, A., & Simon, M. (2021). Interrater reliability in systematic review methodology: exploring variation in coder decision-making. *Sociological Methods & Research*, 50(2), 837-865.

770 Blake, K. R., & Gangestad, S. (2020). On attenuated interactions, measurement error, and statistical power: guidelines for social and personality psychologists. *Personality and Social Psychology Bulletin*, 46(12), 1702-1711.

775 Brauer, M., & McClelland, G. (2005). L'utilisation des contrastes dans l'analyse des données: Comment tester les hypothèses spécifiques dans la recherche en psychologie? [The use of contrasts in data analysis: How to test specific hypotheses in psychological research]. *L'Année Psychologique*, 105(2), 273-305.

Bredan, A. (2020). Conducting publishable research under conditions of severely limited resources. *Libyan Journal of Medicine*, 15(1), 1688126.

780 Brysbaert, M. (2019). How many participants do we have to include in properly powered experiments? A tutorial of power analysis with reference tables. *Journal of Cognition*, 2(1), 16.

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: Neurosciences* 14(5), 365-376.

785 Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B. A., & Pfeiffer, T. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637-644.

Chambers, C. D. (2013). Registered reports: A new publishing initiative at Cortex. *Cortex*, 49(3), 609-610.

- 790 Champely, S., Ekstrom, C., Dalgaard, P., Gill, J., Weibelzahl, S., Anandkumar, A., Ford, C.,
Volcic, R., & De Rosario, H. (2017). pwr: Basic functions for power analysis.
<http://CRAN.R-project.org/package=pwr>.
- Champoux, J. E., & Peters, W. S. (1987). Form, effect size and power in moderated
regression analysis. *Journal of Occupational Psychology*, 60(3), 243-255.
- 795 Chin, W. W., Marcolin, B. L., & Newsted, P. R. (2003). A partial least squares latent variable
modeling approach for measuring interaction effects: Results from a Monte Carlo
simulation study and an electronic-mail emotion/adoption study. *Information Systems
Research*, 14(2), 189-217.
- Clarivate Analytics. (2021). *InCites Journal Citation Reports 2021*.
- 800 Cohen, J. (1965). Some statistical issues in psychological research. In B. B. Wolman (Ed.),
Handbook of clinical psychology (pp. 95-121). McGraw-Hill.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Lawrence Erlbaum
Associates.
- Correll, J., Mellinger, C., McClelland, G. H., & Judd, C. M. (2020). Avoid Cohen's
805 'small', 'medium', and 'large' for power analysis. *Trends in Cognitive Sciences*, 24(3),
200-207.
- Cronbach, L. J. (1987). Statistical tests for moderator variables: Flaws in analyses recently
proposed. *Psychological Bulletin*, 102(3), 414-417.
- Dawson, J. F., & Richter, A. W. (2006). Probing three-way interactions in moderated multiple
810 regression: development and application of a slope difference test. *Journal of Applied
Psychology*, 91(4), 917.
- DeMarco, T. C., & Newheiser, A. K. (2019). When groups do not cure: Group esteem
moderates the social cure effect. *European Journal of Social Psychology*, 49(7), 1421-
1438.
- 815 Demidenko, E. (2008). Sample size and optimal design for logistic regression with binary
interaction. *Statistics in Medicine*, 27(1), 36-46.
- Domingue, B., Kanopka, K., Trejo, S., Rhemtulla, M., & Tucker-Drob, E. (2022). Ubiquitous
bias & false discovery due to model misspecification in analysis of statistical
interactions: The role of the outcome's distribution and metric properties.
820 10.31234/osf.io/932fm. *Psychological Methods*.
- Durand, C. P. (2013). Does raising type 1 error rate improve power to detect interactions in
linear regression models? A simulation study. *PLoS One*, 8(8), e71079.

- 825 Durnez, J., Degryse, J., Moerkerke, B., Seurinck, R., Sochat, V., Poldrack, R. A., & Nichols, T. E. (2016). Power and sample size calculations for fMRI studies based on the prevalence of active peaks. *BioRxiv*, 049429.
- Ebersole, C. R., Mathur, M. B., Baranski, E., Bart-Plange, D.-J., Buttrick, N. R., Chartier, C. R., Corker, K. S., Corley, M., Hartshorne, J. K., & IJzerman, H. (2020). Many Labs 5: Testing pre-data-collection peer review as an intervention to increase replicability. *Advances in Methods and Practices in Psychological Science*, 3(3), 309-331.
- 830 Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175-191.
- Flora, D. B. (2020). Thinking about effect sizes: From the replication crisis to a cumulative psychological science. *Canadian Psychology/Psychologie Canadienne*, 61(4), 318-330.
- 835 Fox, J., & Weisberg, S. (2018). *An R companion to applied regression*. Sage publications.
- Fraley, R. C., & Marks, M. J. (2007). The null hypothesis significance testing debate and its implications for personality research. In R. W. Robins, R. C. Fraley, & R. F. Krueger (Eds.), *Handbook of research methods in personality psychology* (pp. 149-169). The Guilford Press.
- 840 Fraley, R. C., & Vazire, S. (2014). The N-pact factor: Evaluating the quality of empirical journals with respect to sample size and statistical power. *PLoS One*, 9(10), e109019.
- Funder, D. C., & Ozer, D. J. (2019). Evaluating effect size in psychological research: Sense and nonsense. *Advances in Methods and Practices in Psychological Science*, 2(2), 156-168.
- 845 Gelman, A. (2018). You need 16 times the sample size to estimate an interaction than to estimate a main effect. Retrieved from <https://statmodeling.stat.columbia.edu/2018/03/15/need-16-times-sample-size-estimate-interaction-estimate-main-effect/>.
- 850 Gignac, G. E., & Szodorai, E. T. (2016). Effect size guidelines for individual differences researchers. *Personality and Individual Differences*, 102, 74-78.
- Giner-Sorolla, R. (2012). Science or art? How aesthetic standards grease the way through the publication bottleneck but undermine science. *Perspectives on Psychological Science*, 7(6), 562-571.
- 855 Giner-Sorolla, R. (2018). Powering your interaction. Retrieved from <https://approachingblog.wordpress.com/2018/01/24/powering-your-interaction-2>.

- Goodhue, D., Lewis, W., & Thompson, R. (2007). Research note—Statistical power in analyzing interaction effects: Questioning the advantage of PLS with product indicators. *Information Systems Research*, *18*(2), 211-227.
- 860 Goulet, M.-A., & Cousineau, D. (2019). The power of replicated measures to increase statistical power. *Advances in Methods and Practices in Psychological Science*, *2*(3), 199-213.
- Guggenmos, R. D., Piercey, M. D., & Agoglia, C. P. (2018). Custom contrast testing: Current trends and a new approach. *The Accounting Review*, *93*(5), 223-244.
- 865 Hedges, L. V., & Schauer, J. M. (2019). More than one replication study is needed for unambiguous tests of replication. *Journal of Educational and Behavioral Statistics*, *44*(5), 543-570.
- Heidel, R. E. (2016). Causality in statistical power: Isomorphic properties of measurement, research design, effect size, and sample size. *Scientifica (Cairo)*, *2016*, 8920418.
- 870 Hemphill, J. F. (2003). Interpreting the magnitudes of correlation coefficients. *American Psychologist*, *58*(1), 78–79.
- Heo, M., Kim, N., & Faith, M. S. (2015). Statistical power as a function of Cronbach alpha of instrument questionnaire items. *BMC Medical Research Methodology*, *15*(1), 86.
- Hernández, A. V., Steyerberg, E. W., & Habbema, J. D. (2004). Covariate adjustment in
875 randomized controlled trials with dichotomous outcomes increases statistical power and reduces sample size requirements. *Journal of Clinical Epidemiology*, *57*(5), 454-460.
- Hoening, J. M., & Heisey, D. M. (2001). The abuse of power: the pervasive fallacy of power calculations for data analysis. *The American Statistician*, *55*(1), 19-24.
- 880 Iacobucci, D., Posavac, S. S., Kardes, F. R., Schneider, M. J., & Popovich, D. L. (2015). Toward a more nuanced understanding of the statistical properties of a median split. *Journal of Consumer Psychology*, *25*(4), 652-665.
- Jaccard, J., Turrisi, R., & Jaccard, J. (2003). Interaction effects in multiple regression (Vol. 72). *Newcastle upon Tyne: Sage*.
- 885 Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and social psychology review*, *2*(3), 196-217.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams Jr, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., & Bahník, Š. (2018). Many Labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and
890 Practices in Psychological Science*, *1*(4), 443-490.

- Knottnerus, J. A., & Bouter, L. M. (2001). The ethics of sample size:: Two-sided testing and one-sided thinking. *Journal of Clinical Epidemiology*, *54*(2), 109-110.
- Kovacs, M., van Ravenzwaaij, D., Hoekstra, R., & Aczel, B. (2022). Samplesizeplanner: A tool to estimate and justify sample size for two-group studies. *Advances in Methods and Practices in Psychological Science*, *5*(1), 25152459211054059.
- 895
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 863.
- Lakens, D. (2019). The value of preregistration for psychological science: A conceptual analysis. *Japanese Psychological Review*, *62*(3), 221-230.
- 900
- Lakens, D. (2022). Sample size justification. *Collabra: Psychology*, *8*(1), 33267.
- Lakens, D., & Caldwell, A. R. (2021). Simulation-based power analysis for factorial analysis of variance designs. *Advances in Methods and Practices in Psychological Science*, *4*(1), 2515245920951503.
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, *1*(2), 259-269.
- 905
- Ledgerwood, A. (2019). New developments in research methods. In R. F. Baumeister & E. J. Finkel (Eds.), *Advanced social psychology* (pp. 39–61). Oxford University Press.
- Lovakov, A., & Agadullina, E. R. (2021). Empirically derived guidelines for effect size interpretation in social psychology. *European Journal of Social Psychology*, *51*(3), 485-504.
- 910
- Majer, J. M., Zhang, K., Zhang, H., Höhne, B. P., & Trötschel, R. (2022). Give and take frames in shared-resource negotiations. *Journal of Economic Psychology*, *90*, 102492.
- Mathieu, J. E., Aguinis, H., Culpepper, S. A., & Chen, G. (2012). Understanding and estimating the power to detect cross-level interaction effects in multilevel modeling. *Journal of Applied Psychology*, *97*(5), 951-966.
- 915
- Maxwell, S. E., & Delaney, H. D. (2004). *Designing experiments and analyzing data: A model comparison perspective*. Lawrence Erlbaum Associates.
- McClelland, G. H., & Judd, C. M. (1993). Statistical difficulties of detecting interactions and moderator effects. *Psychological Bulletin*, *114*(2), 376-390.
- 920
- McClelland, G. H., Lynch Jr, J. G., Irwin, J. R., Spiller, S. A., & Fitzsimons, G. J. (2015). Median splits, Type II errors, and false-positive consumer psychology: Don't fight the power. *Journal of Consumer Psychology*, *25*(4), 679-689.

- 925 Morris, S. B. (2008). Estimating effect sizes from pretest-posttest-control group designs. *Organizational Research Methods, 11*(2), 364-386.
- Morris, S. B., & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods, 7*(1), 105-125.
- 930 Motyl, M., Demos, A. P., Carsel, T. S., Hanson, B. E., Melton, Z. J., Mueller, A. B., Prims, J., Sun, J., Washburn, A. N., & Wong, K. M. (2017). The state of social and personality science: Rotten to the core, not so bad, getting better, or getting worse? *Journal of Personality and Social Psychology, 113*(1), 34-58.
- Murayama, K., Usami, S., & Sakaki, M. (2022). Summary-statistics-based power analysis: A new and practical method to determine sample size for mixed-effects modeling. 935 *Psychological Methods.*
- Murphy, P. K., Wilkinson, I. A., Soter, A. O., Hennessey, M. N., & Alexander, J. F. (2009). Examining the effects of classroom discussion on students' comprehension of text: A meta-analysis. *Journal of Educational Psychology, 101*(3), 740-764.
- Myers, A., & Hansen, C. H. (2011). *Experimental psychology*. Cengage Learning.
- 940 Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences, 115*(11), 2600-2606.
- O'Connor, B. P. (2006). Programs for problems created by continuous variable distributions in moderated multiple regression. *Organizational Research Methods, 9*(4), 554-567.
- 945 Olejnik, S., & Algina, J. (2003). Generalized eta and omega squared statistics: measures of effect size for some common research designs. *Psychological Methods, 8*(4), 434-447.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology, 45*(4), 867-872.
- 950 Overton, R. C. (2001). Moderated multiple regression for interactions involving categorical variables: A statistical control for heterogeneous variance across two groups. *Psychological Methods, 6*(3), 218-233.
- Perugini, M., Gallucci, M., & Costantini, G. (2018). A practical primer to power analysis for simple experimental designs. *International Review of Social Psychology, 31*(1).
- R Core Team. (2013). R: A language and environment for statistical computing.
- 955 Ravenscroft, S. P., & Buckless, F. A. (2017). Contrast coding in ANOVA and regression. In T. Libby (Ed.), *The Routledge companion to behavioural accounting research* (pp. 349-372). Routledge.

- Richard, F. D., Bond Jr, C. F., & Stokes-Zoota, J. J. (2003). One hundred years of social psychology quantitatively described. *Review of General Psychology*, 7(4), 331-363.
- 960 Richter, M. (2016). Residual tests in the analysis of planned contrasts: Problems and solutions. *Psychological Methods*, 21(1), 112-120.
- Rosenthal, R., & Rosnow, R. L. (1985). *Contrast analysis: Focused comparisons in the analysis of variance*. Cambridge University Press.
- Rosnow, R. L., & Rosenthal, R. (1991). If you're looking at the cell means, you're not looking
965 at only the interaction (unless all main effects are zero). *Psychological Bulletin*, 110(3), 574-576.
- Rosnow, R. L., & Rosenthal, R. (1996). Contrasts and interactions redux: Five easy pieces. *Psychological Science*, 7(4), 253-257.
- Rubin, M. (2017). Do p values lose their meaning in exploratory analyses? It depends how
970 you define the familywise error rate. *Review of General Psychology*, 21(3), 269-275.
- Rubin, M. (2020). The costs of HARKing. *The British Journal for the Philosophy of Science*.
- Schäfer, T., & Schwarz, M. A. (2019). The meaningfulness of effect sizes in psychological research: Differences between sub-disciplines and the impact of potential biases. *Frontiers in Psychology*, 10, 813.
- 975 Shieh, G. (2009). Detection of interactions between a dichotomous moderator and a continuous predictor in moderated multiple regression with heterogeneous error variance. *Behavior Research Methods*, 41(1), 61-74.
- Shieh, G. (2010). On the misconception of multicollinearity in detection of moderating effects: Multicollinearity is not always detrimental. *Multivariate Behavioral Research*,
980 45(3), 483-507.
- Simonsohn, U. (2014). No-way interaction. <http://datacolada.org/17%20Simonsohn>.
- Sisk, V. F., Burgoyne, A. P., Sun, J., Butler, J. L., & Macnamara, B. N. (2018). To what extent and under which circumstances are growth mind-sets important to academic achievement? Two meta-analyses. *Psychological Science*, 29(4), 549-571.
- 985 Stone-Romero, E. F., Alliger, G. M., & Aguinis, H. (1994). Type II error problems in the use of moderated multiple regression for the detection of moderating effects of dichotomous variables. *Journal of Management*, 20(1), 167-178.
- Tepe, B., & Byrne, R. M. J. (2022). Cognitive processes in imaginative moral shifts: How judgments of morally unacceptable actions change. *Memory & Cognition*, 50(5),
990 1103-1123.

- Van't Veer, A. E., & Giner-Sorolla, R. (2016). Pre-registration in social psychology—A discussion and suggested template. *Journal of Experimental Social Psychology*, 67, 2-12.
- 995 Wagenmakers, E.-J., Wetzels, R., Borsboom, D., van der Maas, H. L., & Kievit, R. A. (2012). An agenda for purely confirmatory research. *Perspectives on Psychological Science*, 7(6), 632-638.
- Wang, Y., & Xie, X. (2021). Halfway to my request is not halfway to my heart: Underestimating appreciation for partial help. *Personality and Social Psychology Bulletin*, 47(10), 1466-1479.
- 1000 Wang, Y. A., & Rhemtulla, M. (2021). Power analysis for parameter estimation in structural equation modeling: A discussion and tutorial. *Advances in Methods and Practices in Psychological Science*, 4(1), 2515245920918253.
- White II, M. H. (2018). *On calculating power for interactions in 2 x 2 factorial designs*. <https://www.markhw.com/blog/power-twoway>.
- 1005 Wiley, J. F., & Pace, L. A. (2015). Probability distributions. In *Beginning R: An Introduction to Statistical Programming* (pp. 53-65). Apress. https://doi.org/10.1007/978-1-4842-0373-6_6