“Rarely safe to assume”: Evaluating the use and interpretation of manipulation checks in
experimental social psychology

Emma Ejelöv
Timothy J. Luke

University of Gothenburg

This manuscript is published in *Journal of Experimental Social Psychology*
(<https://doi.org/10.1016/j.jesp.2019.103937>).

The text of the preprint may differ slightly from the version of record.

Accepted: 2019-11-24

Minor text corrections: 2019-12-02

Version of record published online: 2019-12-10

Corrections: 2020-01-17

Authors' note: The authors each contributed equally. Order of authors was determined by a badminton match, played to two sets (21-19 to EE, 21-16 to EE). For their comments on earlier drafts, we would like to thank Karl Ask, Will Crozier, and Patrik Michaelsen, who are listed alphabetically by surname, rather than by athletic performance.

Abstract

Although the use of manipulation checks is widespread in social psychology, several researchers have raised methodological concerns about their use and interpretations. However, knowledge of how they are actually being used has been lacking. Extracting data from published reports of 207 recent experiments, we provide an empirical review of current practices concerning manipulation checks in social psychology. Our review suggests that there are serious deficiencies in the manner in which manipulation checks are used and interpreted. For example, published reports tend to contain highly limited quantitative reasoning about the effectiveness of manipulations, and researchers report little or nothing to address the possibility that manipulation checks might cause undesirable reactivity among participants. However, we argue that manipulation checks can be highly beneficial components of experiments when used properly, and they have untapped potential for the quantitative assessment of the strength of manipulations relative to the effect on the dependent variable (i.e., causal efficacy). To assist with such assessments, we provide empirical benchmarks for causal efficacy in social psychology. Additionally, we provide several recommendations for researchers and reviewers for improving the use and reporting of manipulation checks.

Keywords: manipulation check, construct validity, experimental methods, causal efficacy

“Rarely safe to assume”: Evaluating the use and interpretation of manipulation checks in experimental social psychology

“It is rarely safe to assume beforehand that the operations used to manipulate variables will be successful and will tie in directly with the concept the experimenter has in mind.” (Festinger, 1953, p.145)

In their chapter on experimental research in the *Handbook of Social Psychology*, Wilson, Aronson, and K. Carlsmith (2010) write, “An experiment cannot test a hypothesis unless the independent variable manipulates what it is supposed to manipulate” (p.71). It is difficult to disagree. But what is actually the optimal method of assessing the effectiveness of experimental manipulations? The typical approach is, of course, a *manipulation check*. Manipulation checks are a staple of modern experimental, especially social, psychology (Hauser, Ellsworth, & Gonzalez, 2018), and are considered a fundamental technique. Researchers openly assert their importance (see, e.g., Fayant, Sigall, Lemonnier, Retsin, & Alexopoulos, 2017), and they are even addressed in undergraduate-level research methods textbooks (e.g., Haslam & McGarty, 2019). Manipulation checks are purported to provide evidence of construct validity (Sigall & Mills, 1998). However, only a small amount of data currently exists concerning *how* manipulation checks are actually used, making it difficult to ascertain the extent to which manipulation checks provide meaningful evidence for the effectiveness of a manipulation.

In this review, we present an empirical description and evaluation of current practices concerning manipulation checks in social psychology. We begin by briefly reviewing the practice of construct validation. We discuss how this process relates to the way manipulation checks have been and are currently being conceptualized and issues regarding the various uses and interpretations of manipulation checks. Using data collected from 207 recent social psychology experiments, we then examine how manipulation checks are used in practice, and we end with recommendations for research practice and offer some statistical tools for capitalizing on the untapped potential that we believe manipulation checks possess (specifically, benchmarks for causal efficacy). We find that many criticisms of manipulations checks have merit, but like many past (e.g., Festinger, 1953) and contemporary researchers (e.g., Fabrigar & Wegener, 2016; Flake, Pek, Hehman, 2017; Foschi, 2014), we argue that manipulation checks can be highly beneficial components of experiments when used properly. Indeed, researchers ought to make more extensive efforts to check the validity of manipulations.

Construct validation

Researchers and methodologists in the social sciences have written extensively about how to establish construct validity, particularly in the context of measuring psychological variables. In a widely-influential monograph on psychological testing, Loevinger (1957)

distinguished between *substantive*, *structural*, and *external* components of construct validity¹. The substantive component concerns the extent to which the scope and content an operationalization adequately represents the intended construct. The structural component concerns statistical properties of an operationalization such as its reliability, dimensionality, and factor invariance (e.g., Reise, Widaman, & Pugh, 1993). The external component concerns generalizability and distinctiveness, such as the extent to which the measure predicts other variables with which it is hypothesized to correlate and not correlate (Cronbach & Meehl, 1955). Whereas the substantive component is largely theoretical, the structural and external components require careful quantitative specification and empirical testing.

These components of construct validity have been considered primarily in the context of psychological measurement (AERA, APA, & NCME, 2014; Flake, Pek, & Hehman, 2017). However, this description of construct validation has been extended to experimental manipulations as well. In an early extension of construct validity theory to manipulations, Cook and Campbell (1979) argued that before deciding how to manipulate a construct, researchers ought to carefully consider the construct's definition and identify constructs from which it must be distinguished (substantive validation) and that the validity of manipulations should be evaluated empirically, with pre-experiment pilot testing that assesses the extent to which the manipulation affects measures of the intended construct and does not affect measures of unintended constructs (structural and external validation; see also Shadish, Cook, & Campbell 2002). In short, essentially the same methodological and statistical considerations that apply to the validation of measurements also apply to the validation of manipulations. However, although there is a long tradition of thorough validation research and theory for measurements (Cronbach & Meehl, 1955; Loevinger, 1957; Nunnally, 1970; Raykov & Marcoulides, 2011), the empirical validation of experimental treatments in social psychology has been dominated by simpler techniques collectively known as *manipulation checks* (Sigall & Mills, 1998; Hauser et al., 2018).

Manipulation checks

To contextualize current practice and understanding, it can be useful to examine the origin and development of the concept of checking manipulations. Both the methods and definition of manipulation checks have to some extent changed from the term's origin in the 1950's to the present day. In what may have been the earliest published discussion of manipulation checks (without using the specific term) as a methodological tool, in the sentence immediately following the quotation that opened this paper, Leon Festinger (1953) commented that, "It is a worthwhile precaution to check on the success of the experimental manipulations" (p.145). He also provided examples of studies in which this had been, in his view, successfully performed (e.g., Kelley, 1951). Concern for the effectiveness of manipulations seems to have been a virtue Festinger vigorously instilled in his students (Aronson, 2010), many of whom

¹ A thorough review of construct validation is beyond our present purposes. For a brief but informative review (from which we have drawn substantially), see Flake, Pek, and Hehman (2017). For a more complete review, see, e.g., Raykov and Marcoulides (2011).

became influential figures in social psychology. These students addressed “manipulation checks” in their writing on methods: Aronson and J. M. Carlsmith (1968) used the term in a chapter on experimentation in the *Handbook of Social Psychology* and Mills (1969) included the term in the index of his textbook *Experimental Social Psychology*.

There is ample evidence that manipulation checks are in widespread use in social psychology. For example, Sigall and Mills (1998) noted that the reported use of manipulation checks had increased from 27% in 1965 to 63% in 1995, for experiments published in the *Journal of Personality and Social Psychology* and the *Journal of Experimental Social Psychology*. Similarly, in a more recent review, Hauser, Ellsworth, and Gonzelez (2018) found that between 41% and 63% of empirical articles in top-ranking social psychology journals featured manipulation checks.

Despite the prevalence of manipulation checks, substantial variation has existed (and still exists) regarding their perceived benefits and drawbacks. Manipulation checks have been subjected to criticism for almost as long as they have been proposed (e.g., Kidd, 1976; Mills, 1969; Sigall & Mills, 1998). Indeed, some critics go so far as to argue that manipulation checks are not only unnecessary but harmful to science (Fayant et al., 2017). Many of these critiques (as we will discuss below) often depend on the manner in which a check is used and interpreted.

Concept and terminology

In many of the early mentions of manipulation checks (e.g., Aronson & Carlsmith, 1968; Mills, 1969), they were described as “a separate measure designed to find out whether the experimental treatment succeeded in producing the internal state that constituted the conceptual variable” (J. M. Carlsmith, Ellsworth, & Aronson, 1976, p.141). Thus, the original definition of a manipulation check was a measure of the construct targeted by an experimental manipulation. While this is still consistent with the modern descriptions of the concept occurring in textbooks (e.g., Haslam & McGarty, 2019; Wilson, Aronson, & K. Carlsmith, 2010), Sigall and Mills (1998) noted that in practice, the term is sometimes used to refer to “measures taken to check on the assumption that the differences between the experimental conditions were noticed” (p.221). Similarly, in their recent discussion of manipulation checks, Hauser, Ellsworth, and Gonzalez (2018) additionally included checks of attention and measures of mediating variables under the term. This indicates that the definition of what constitutes a manipulation check may have changed during the last couple of decades, but to what extent researchers in general include additional measures under the term and whether different measures are evaluated appropriately is not known.

The measurement of manipulation checks is largely concerned with the issue of assessing construct validity – that is, the extent to which the operationalizations adequately reflect the constructs they are meant to represent (Shadish, Cook & Campbell, 2002). Some critics have expressed concerns that researchers may misunderstand exactly how manipulation checks contribute to construct validity (Fayant et al., 2017; Sigall & Mills, 1998). For instance, Sigall and Mills (1998) reasoned that measures of the construct of interest do not add any additional information on construct validity when there are no other plausible explanations for the observed

effect on the dependent variable, and thus are not necessary. When alternative explanations do exist, they argue that the inclusion of such manipulation checks may contribute to construct validity but that they do not safeguard against construct confounding (viz. manipulating more than the intended construct). Still others highlight that manipulation checks used in concert with measures of other possible explanations could give researchers a fair estimation of whether a manipulation has affected only the intended construct or other unintended constructs (Foschi, 2014).

In summary, construct validity may be strengthened by assessing variables that are hypothesized to be directly influenced by the manipulation, as well as variables that are expected to be uninfluenced by the manipulation. To date however, we do not know to what extent researchers take measures to safeguard against alternative explanations. It is also unknown to what extent researchers use direct measures of their manipulations, or if they rely on indirect measures that are rather hypothesized to correlate with (rather than directly measure) the construct of interest. Different types of measures may have different implications for construct validity and offer different limitations. To clearly distinguish between different types of measures relating to the manipulation and research design may therefore be important to convey that study's particular strengths and weaknesses. The extent to which this is actually done is, however, not clear.

Use and interpretation

Methodological use. Manipulation checks have not only apparently drifted from their original conception but also in their methodological approach. Such checks have not always referred to the kinds of measures we would today recognize as typical manipulation checks – that is, quantitative measures, usually self-reports, taken after the implementation of a treatment. Rather, Aronson and J. M. Carlsmith (1968) advocated extensive pretesting of manipulations in which participants were interviewed at length following the experience of an experimental intervention. This advice is reiterated (almost verbatim) by Aronson, Ellsworth, J. M. Carlsmith, and Gonzales (1990) in their textbook on research methods in social psychology. In contrast to these recommendations, contemporary psychology appears to depend predominantly on verbal measures (e.g., self-report) taken during an experiment (Hauser, Ellsworth, & Gonzalez, 2018). In an early critique of the typical manipulation check, Kidd (1976) used a self-report response on a Likert scale as an example. Mills (1969) also provided an ordinal self-report measure taken during an experiment (rather than during a pilot test) as an example of a manipulation check. To the extent that these examples were representative of methods at the time, it seems that it was not so long after Festinger's (1953) and Aronson and J. M. Carlsmith's (1968) recommendations that researchers began to more heavily rely on verbal measures as manipulation checks, rather than extensive pilot testing (to the extent that researchers relied on piloting to begin with).

The apparently frequent use of self-report measures as manipulation checks in experiments has raised concerns among many critics. One prominent critique concerns the consequences of measuring the manipulation check in the main study, as opposed to in a pilot. If a manipulation check is taken between the manipulation and dependent measure, there is a

possible risk of manipulation checks either revealing the purpose of the study to participants or prompting participants to introspect about the manipulated construct, possibly affecting participants' responses to the manipulation check itself or the dependent variable, respectively referred to demand effects and reactivity (Kidd, 1976; Mills, 1969; Parro, & Hertel, 1999). This critique is still cited as a reason to avoid manipulation checks (Fayant et al., 2017; Hauser, Ellsworth, & Gonzalez, 2018). The order in which a manipulation check is placed could potentially have different effects, such that placing it after the dependent variable could result in weakened effects on the manipulation check, due to either a short-lived effect of the manipulation or the response to the dependent variable allowing for processing of the manipulation. Placing it before the dependent variable could instead result in the manipulation check having an adverse effect on the dependent variable (Fayant et al., 2017). There is evidence that at least some kinds of checks can cause reactive effects (e.g., attention checks can encourage systematic thinking; Hauser & Schwarz, 2015). While previous research has shown that the majority of experiments do not manipulate the order of manipulation checks (Hauser, Ellsworth, & Gonzalez, 2018), this problem is not necessarily specific to manipulation checks, as all measures (e.g., additional dependent measures or proposed mediators) may have unintentional effects on any subsequent measure (Schwarz, & Strack, 1991). One often cited solution to the problem of order effects is to spend some time on pretesting manipulations, thereby avoiding the risk of affecting either the manipulation check or the dependent variable (Hauser, Ellsworth, & Gonzalez, 2018; Kidd, 1976). However, the extent to which manipulation checks are used in pilots as opposed to main studies is not currently known².

Another issue regarding the interpretation of manipulation checks concerns measurement error. A failed manipulation check could potentially be due to an inadequate measure of the independent variable and not a failed manipulation (Fayant et al., 2017; Hauser, Ellsworth, & Gonzalez, 2018). Similarly, if the measurement of the dependent variable is not as reliable as the measure of the manipulation, the apparent success of the manipulation combined with an absence of effect on the dependent variable could lead researchers to incorrectly discard their hypothesis (Kidd, 1976). Possibly unreliable measures are a common issue in social psychology and should be given serious consideration (see Flake, Pek, & Hehman, 2017). The exclusion of a manipulation check, however, will not in itself improve the reliability of dependent measures.

Statistical use.

Internal analysis. An early identified use of manipulation checks is the possibility of using the manipulation check, instead of the experimental assignment, as the independent variable in a statistical analysis, to ascertain whether an unsupported hypothesis test might be due to a failed manipulation or faulty theory (see, e.g., J. M. Carlsmith, Ellsworth, & Aronson, 1976; Festinger, 1953; Wilson, Aronson, & K. Carlsmith, 2010). That is, if the manipulation check measures the construct represented by the independent variables, researchers may still

² There is some inconsistency in the usage of the terms “pilot” and “pretest” (Westlund & Stuart, 2017). We use both terms to refer to any assessments of the effectiveness of a manipulation conducted separately from a hypothesis test (see, e.g., Wilson, Aronson, & K. Carlsmith, 2010). As we use the term, “pilot” can also refer to a study aimed at validating a manipulation, even if it is not intended to prepare for a specific planned experiment.

observe the predicted relationship between the manipulation check measure and the dependent measure, even if the manipulation was ineffective. There have been many critiques of internal analysis. Some have argued that using manipulation checks in this way unduly capitalizes on chance in finding significant effects (Kidd, 1976). Others have argued that the interpretation of internal analysis is ambiguous, because the relationship between the manipulation check and the dependent variable is merely correlational (Hauser, Ellsworth & Gonzalez, 2018). The use of manipulation checks as mediators has similarly been criticized on the grounds that they do not rule out alternative explanations (variables that may correlate with the manipulation check), cannot provide a basis for causal claims, and are subjected to the same problem with order effects as discussed previously (Fiedler, Schott, & Meiser, 2011; Hauser, Ellsworth, & Gonzalez, 2018). Whether or not researchers actually use manipulation checks for these types of (potentially problematic) statistical analyses is currently not well documented.

Manipulation strength and causal efficacy. Another potential use of manipulation checks is to quantitatively assess the strength of a manipulation, which is of prime importance for many purposes. Many researchers have argued that, especially for practical applications, small causes that lead to large effects are more useful than large causes that lead to small effects (Prentice & Miller, 1992). Moreover, the strength of a manipulation can directly bear on the size of the effect observed in the dependent variable (see, e.g., Festinger, 1953; Mills, 1969). That is, an effect of a given size could be the result of either a strong manipulation on a weakly related dependent variable or the effect of a weak manipulation on a strongly related dependent variable. Additionally, variance in manipulation strength can partly account for the heterogeneity in effect sizes in an experimental literature (Tryon, 2001).

Considering the importance of manipulation strength for the interpretation of effect size, Abelson (1995) offered a suggestion for a quantitative measure of *causal efficacy* – the ratio of effect size to cause size – to help assess the magnitude and meaningfulness of an experimental effect. Abelson's (1995) recommended calculation for causal efficacy was simply dividing the unstandardized effect size (effect on the dependent variable) by the cause size (effect on the manipulation check). For example, if a manipulation produces a mean difference of 4.50 scale points on the dependent measure and a mean difference of 2.10 points on the manipulation check, this would be a causal efficacy of 2.14 – the effect size is 2.14 times larger than the cause size. Researchers and data analysts have remarked about the potential usefulness of the causal efficacy ratio but have also remarked that researchers almost never report interpretation of cause size in empirical articles (e.g., Breugh, 2003; Tryon, 2001).

The present review

Critics of manipulation checks have presented persuasive arguments for their concerns. However, these criticisms depend on researchers generally misunderstanding (or failing to present their reasoning about) the strengths and limitations of their measures, but at present, we lack systematic data on researchers' reasoning, argumentation and quantitative practices regarding manipulation checks. Without knowing *how* manipulation checks are used, it is not possible to assess the extent to which criticisms apply to practice. Moreover, the process of

construct validation demands well-specified quantitative reasoning (e.g., predictions about what variables will and will not be influenced; interpretations of the strength of manipulations). To evaluate how well current practices contribute to the validation of manipulations, we would benefit from knowing the extent to which researchers are reporting their reasoning and presenting statistical evidence of construct validity.

Thus, we undertake an empirical review of the manner in which manipulation checks are used in social psychology. This review has two main goals: (1) to provide a description and evaluation of present practice concerning manipulation checks, and (2) to provide recommendations to improve their use and interpretation. The data we collected allow us to assess the validity of many past criticisms and also to consider how manipulation checks might be beneficial in ways in which researchers have not fully capitalized.

Method

Sampling procedure

For this review, we examined empirical articles from the *Journal of Experimental and Social Psychology (JESP)*, *Journal of Personality and Social Psychology (JPSP)*, and *Personality and Social Psychology Bulletin (PSPB)*. These were selected as they were the three journals with the highest impact factor of non-specialty social psychology journals at the time of the coding. To capture the current use of manipulation checks, we decided to include issues from the year 2015 to 2018. To achieve a reasonably representative sample of articles from these journals and years, we randomly sampled two issues from *JPSP* and *PSPB* for each of the four specified years, and one volume from *JESP* (roughly equal to the length of two issues from *JPSP* and *PSPB*) per year. This resulted in 8 issues from *JPSP*, 8 issues from *PSPB*, and 4 volumes from *JESP*.

We searched for relevant articles using Google Scholar, limiting the search to the pre-specified journals, years, issues/volumes and search terms. Our search terms included: “manipulation check”, “attention check”, “pilot”, “pretest”, “pre-test”, “prestudy” and “pre-study”. A pretest of this procedure³ indicated that it adequately identified articles, compared to manually searching for articles. Google Scholar also outperformed other search tools. For example, using ProQuest, SCOPUS and PsycNET did not adequately identify the search terms in the full text, and numerous articles were missed.

Note that our approach cannot (and was not intended to) estimate the prevalence of manipulation checks. Rather, it was aimed at getting a representative view of how manipulation checks are used. For estimates of prevalence, see Sigall & Mills (1998) and Hauser, Ellsworth, and Gonzalez (2018). We also do not offer a direct assessment of the usefulness of any given check (e.g., whether authors have adequately measured the construct or have sufficiently measured potential confounds), since such evaluations would likely require substantial subject-matter expertise in the areas of each experiment we examined.

³ We registered our procedures on the Open Science Framework (osf.io/zw6jg/). We made some alterations to our approach after unexpected difficulties, so we registered after each change in the plan.

The articles resulting from this search strategy were manually examined and included based on whether the article included some form of assessment of the validity or effectiveness of an experimental manipulation. We also included checks of attention to procedures and experimental stimuli (see coding description below), since these measures are sometimes called manipulation checks (see, e.g., Hauser, Ellsworth, & Gonzalez, 2018). The sampling procedure rendered 87 articles, from which we coded 411 manipulation checks from 207 experiments.

Coding procedure

Complete documentation of our coding procedures, a list of references for coded articles, and all coded data are available on the Open Science Framework (osf.io/zw6jg/). In the interest of space, we present a streamlined summary of the procedures here.

As we noted previously, the term “manipulation check” can refer to several kinds of measures. We coded the manipulation checks firstly based on what they measured. Manipulation checks were labeled *independent variable measures* when the measure was a variable intended to be directly affected by the manipulation (e.g., manipulating the sound of rain and measuring the perceived sound of rain). We made a distinction between *correlates* and independent variable measures. We draw this distinction as variables that correlate with the construct of interest can be used to assess convergent validity, rather than directly assessing construct validity. We coded a measure as a correlate when the measure was a variable not intended to be directly affected by the manipulation but was predicted to correlate with the construct of interest (e.g., manipulating social class belonging and measuring perceived income). We also made a distinction between manipulation checks that were used to assess discriminant validity. Measures that were predicted to be uninfluenced by the manipulation were coded as *discriminant variables* (e.g., manipulating a person's emotional expression and measuring attractiveness). We additionally made a distinction between different types of attention checks. Attention checks that measured whether participants were engaged in the task at hand were coded as *general attention checks* (e.g., measuring whether participants had read a specific question), while measures of whether participants had sufficiently encoded the stimulus material were coded as *stimulus attention checks* (e.g., manipulating group membership and asking participants what group membership they had been exposed to), and measures requesting participants to comply with specific instructions to indicate that they are alert and attentive were coded as *instructional manipulation checks* (e.g., asking participants to leave a particular item response blank; Oppenheimer et al., 2009).

We also coded the statistical use of manipulation checks, specifically whether the manipulation check was used as an independent variable in a statistical analysis, used in a mediation analysis, and used as a covariate. To assess authors' explicit quantitative reasoning, we recorded whether authors analyzed manipulation check data using significance testing, whether they reported power calculations for the manipulation check, whether they reported frequentist equivalence tests, whether they made explicit assessments about the size of the effect on the manipulation check (e.g., making comparisons to other relevant effects or using general benchmarks), and whether they made claims of equivalence based on non-significant *p*-values.

We also recorded basic information about the study designs, including the sample size and number of experimental groups.

Afterward, all manipulation checks that were coded as independent variable measures were revisited, and we extracted or calculated the size of the effect of the manipulation on the manipulation check (the “cause size”) and the dependent variable of primary interest in the experiment (the “effect size”). This resulted in 109 pairs of effect sizes and cause sizes, which were used in an analysis of causal efficacy described below.

Inter-rater reliability. To test the inter-reliability of our procedures, both authors coded all relevant articles in two journal issues, *JPSP*, 114 (6) and *PSPB*, 42 (5). Inter-rater reliability based on Gwet’s AC showed adequate reliability or better for all variables (minimum .73), except one variable intended to assess whether there was a meaningful control group in the design. We decided to drop that variable before moving on to the main coding. The discrepancies between the two coders were resolved through discussion. A table of reliability statistics for all variables is available on OSF (osf.io/zw6jg/).

Results and discussion

We divide the results into three major sections: (1) *concept and terminology*; (2) *use and interpretation*, with three subsections, (a) methodological approach, (b) statistical use, and (c) causal efficacy; and lastly (3) *format and style of reports*. For clarity of presentation, in each section, alongside the results, we provide interpretations of the data as well as recommendations for improved use of manipulation checks.

Concept and terminology

How do researchers label their checks? For each coded manipulation check, we noted the terms that authors used to describe them. In our data, the most common term used to describe any kind of check was “manipulation check” (57.25%, 235/411). However, of these, less than half (44.7%, 105/235) were used to refer to measures of the construct of interest (i.e., independent variable measure). The term “manipulation check” was also widely applied to other kinds of measures. It was the most commonly used term for stimulus checks (53/71), correlates (51/75), and discriminant variables (21/69). It was the second most common term used to describe instructional manipulation checks, 5 out of 28 (“attention check” was the most common, 14/28), and never for general attention checks (“attention check” was the most common, 15/19).

As can be seen, the term “manipulation check” is applied to measures that differ quite substantially, and sometimes its use does not necessarily make clear the expected relationship between the manipulation and the check (e.g., using the term “manipulation check” for variables that are assumed to be unaffected by the manipulation). Inconsistency in usage and possible inappropriateness of some terms is concerning, as it may reflect a lack of consideration of the usefulness and limitations of a given measure. Given that explicit hypotheses are not typically provided for manipulation checks, inconsistent terminology makes it difficult for readers to interpret the results of a reported check. Sometimes, it is simple enough to guess what should be expected, but readers probably should not be required to guess about researchers’ predictions. Consider, for example, that both correlates (expected to be associated with the manipulation) and

discriminant variables (*not* expected to be associated with the manipulation) are commonly referred to simply as “manipulation checks.”

Ultimately, the evidence a check provides is more important than its labeling. However, even if not strictly incorrect, inconsistencies in labeling may impede understanding. For these reasons, it would be useful for researchers to establish a more standardized set of terminology. Although we do not think we should necessarily have the last word on this issue, we offer some recommendations for terms (largely based on the distinctions drawn in our coding system).

Recommendations.

Independent variable measures: (Direct) manipulation checks. In our view, the term “manipulation check” (used without any additional qualifiers) should generally be reserved specifically for measures of the construct of interest, as such measures are the most direct assessment of the effectiveness of a manipulation. In situations in which additional clarity is desirable, we suggest the term *direct manipulation check* for this type of measure.

Correlates: Indirect manipulation checks. Sometimes, however, it is not possible to directly measure the construct of interest, or there may be serious concerns that a direct measure could render a manipulation ineffective if it alerts participants to the purpose of the manipulation (Hauser, Ellsworth, & Gonzalez, 2018; Mills, 1969). It can also be useful to measure related constructs in order to provide convergent validity evidence (Cronbach & Meehl, 1955). In such cases, it may be desirable to measure a close correlate of the independent variable. We suggest the term *indirect manipulation check* to refer to such measures, as they can provide important information about the influence on the targeted construct, but that information is only indirect.

Discriminant variables. If a variable is measured to rule out alternative explanations or potential confounds, this should be made clear – and, in our opinion, should not be referred to as a manipulation check without clarifying qualification, to avoid confusion with measures of the targeted construct. Consistent with our coding system, we suggest the term *discriminant variable*, *discriminant measure*, or *discriminant manipulation check* for these types of measures.

The fact that we discourage the use of the unqualified term “manipulation check” for these variables should not be taken as a suggestion that they are unimportant. On the contrary, measuring discriminant variables in addition to manipulation checks could be quite useful, especially if a manipulation is conceptually distal from the independent variable. If, for example, you are using different imagery to manipulate mood, the manipulation can plausibly affect other constructs in addition to mood. Some of these constructs may be correlated with your independent variable while others are measures you do not wish to influence, as they may have separate effects on your dependent variable.

Attention and stimulus checks. We find it somewhat concerning that the term “manipulation check” is frequently used for checks of attention to or processing of the stimulus material. Although this usage is understandable, whether participants have perceived or understood the manipulation is not necessarily a test of whether the manipulation has had the intended effect. In our opinion, a clearer distinction should be made between measures of the targeted construct and measures of whether the participant paid attention to the experimental

stimuli. Again similar to our coding system, we suggest the term *stimulus check* to refer to measures specifically designed to assess attention or memory for the content of a manipulation. Sigall and Mills (1998) suggested the similar term *treatment check* for this kind of measure.

Moreover, we believe it is important to delineate general *attention checks* from stimulus checks. Whereas stimulus checks provides assurance that participants have processed the content of the manipulation, the general attention checks and instructional manipulation checks do not necessarily provide direct assurance that the manipulation content has been processed appropriately. Likewise, a stimulus check might not provide any assurance that a participant was paying attention to anything else in the procedure. In order to avoid confusion about the purpose of the check, we recommend against the term “instructional manipulation check” (cf. Oppenheimer et al., 2009) and instead suggest the term *instructional attention check* (a phrase that sometimes appeared in the articles we coded) for measures that request that a participant comply with instructions to demonstrate their attention. (Note that for the remainder of this paper, we retain the term “instructional manipulation check” or “IMC” to be consistent with our coding procedures.)

Use and interpretation

Methodological approach. What are the methodological characteristics of manipulation checks? Despite the fact that researchers have long advocated the practice of pretesting and pilot testing manipulations (Aronson & J. M. Carlsmith, 1968; Festinger, 1953; Hauser, Ellsworth, & Gonzalez, 2018; Westlund & Stuart, 2017), only 13.5% of coded checks (49/364; excluding attention checks and IMCs) were reported in pilot tests. However, we do not know whether this low rate reflects a low frequency of pilot testing or a failure to report pilots that were conducted. One of the major criticisms of manipulation checks is the potential effect they might have on participants’ mental states and behavior (and as a consequence, the dependent variable), when measured before the dependent variable (e.g., Fayant et al., 2017; Kidd, 1976). The low number of pilots, coupled with previous research showing that order effects are rarely accounted for (Hauser, Ellsworth, & Gonzalez, 2018), suggests that reactivity is a plausible problem in research, but we do not know the magnitude of this possible effect.

Sample sizes were generally larger in main experiments (i.e., those in which the dependent variable was measured) compared to pilot tests. Descriptive statistics for sample sizes are presented in Table 1. There were a total of 24 pilot tests in our data (reporting 49 checks). Of these, one quarter (25.0%, 6/24) did not report a sample size, and just over half (54.2%, 13/24) did not report sufficient information to determine whether the pilot participants were drawn from the same population as the participants in the main experiment. No coded cases reported power calculations for a pilot or specifically for the manipulation check.

Table 1.

Sample sizes in main experiments and pilots

Type	Mean <i>N</i> (SD)	Median <i>N</i>	Maximum <i>N</i>	Minimum <i>N</i>
Main experiment	201.63 (164.89)	155	1,168	25
Pilot	78.67 (46.35)	72.5	199	17

Note: Data were available for 187 main experiments and 18 pilots. All *N*s represent sample sizes after exclusions.

Table 2.

Frequencies of manipulation check usages, by type of check

	Frequency	Type of check
Used for exclusion	2.2% (3/134)	Direct
	2.7% (2/75)	Indirect
	0% (0/69)	Discriminant
	94.7% (18/19)	Attention
	50.7% (36/71)	Stimulus
	64.3% (18/28)	IMC
Used as independent variable		
	0% (0/134)	Direct
	1.3% (1/75)	Indirect
	0% (0/69)	Discriminant
	0% (0/19)	Attention
	0% (0/71)	Stimulus
	0% (0/28)	IMC
Used as covariate		
	0% (0/134)	Direct
	1.3% (1/75)	Indirect
	5.8% (4/69)	Discriminant

	0% (0/19)	Attention
	0% (0/28)	Stimulus
	0% (0/28)	IMC
Used as mediator		
	3.0% (4/134)	Direct
	1.3% (1/75)	Indirect
	31.9% (22/69)	Discriminant
	0% (0/19)	Attention
	0% (0/71)	Stimulus
	0% (0/28)	IMC
Use of frequentist statistics		
	91.8% (123/134)	Direct
	86.7% (65/75)	Indirect
	92.8% (64/69)	Discriminant
	5.3% (1/19)	Attention
	23.9% (17/71)	Stimulus
	0% (0/28)	IMC
Unwarranted claims of equivalence		
	17.2% (23/134)	Direct
	14.7% (11/75)	Indirect
	60.9% (42/69)	Discriminant
	0% (0/19)	Attention
	5.9% (1/71)	Stimulus
	0% (0/28)	IMC

Reasoning about effect sizes

0.7% (1/134)	Direct
1.3% (1/75)	Indirect
0% (0/69)	Discriminant
0% (0/19)	Attention
0% (0/71)	Stimulus
0% (0/28)	IMC

Claims of manipulation effectiveness

65.7% (88/134)	Direct
73.3% (55/75)	Indirect
10.2% (7/69)	Discriminant
5.3% (1/19)	Attention
22.5% (16/71)	Stimulus
0% (0/28)	IMC

Statistical use. In what ways do researchers use the quantitative data from their manipulation checks? Table 2 provides frequencies for statistical use of each type of coded check. One out of every 5 checks (20.0%, 82/411) was used as a criterion to exclude data from participants. These exclusions were almost always made on the basis of a general attention check, stimulus check, or IMC. While data exclusion based on a direct or indirect manipulation check was rare (2.4%, 5/209), this practice is potentially problematic. Data exclusions on the basis of stimulus checks is more difficult to interpret. If the stimulus check in addition to measuring attention to the stimulus material, also captures whether the manipulation has been effective, exclusions based on the stimulus check can create biased groups (e.g., by selecting only participants for whom the manipulation works). This may be more likely in certain contexts, such as when the construct of interest is related to comprehension or attention. When designing stimulus check, this potential limitation should be given serious consideration, and ideally, researchers should offer arguments for why their stimulus checks do not additionally measure the construct of interest.

In our sample, manipulation checks (of any type) were rarely used for analytic purposes other than data exclusion. They were almost never used as independent variables in analyses (e.g., for internal analysis). They were used as covariates in only 1.2% (5/411) of cases. Manipulation checks were used as mediators in 6.6% (27/411) of cases, and these cases occurred in six experiments reported in four articles (4.6% of the sample of 87 articles). Of these 27 mediation analyses, 22 of them (81.5%) involved discriminant variables, for the apparent purpose of ruling out alternative explanations for an observed effect.

Using attention checks and IMCs as exclusion criteria is, in our view, often justified (but see Abbey & Meloy, 2017). However, as with any data exclusion, in order to maximize transparency, authors should clearly report (1) how many cases they have excluded, (2) the basis for those exclusions, and (3) when the decision to exclude was made (see, e.g., Simmons, Nelson, & Simonsohn, 2011). The former two points were often included in the experiments we coded, but the third was rarely, if ever, provided. Ideally, decisions about whether and how to use checks to exclude data should be made prior to data collection, to avoid inflating the risk of false positives with flexible data analysis procedures.

Given the low rate at which direct and indirect manipulation checks are used in mediation analyses, some of the concerns raised by Hauser, Ellsworth, and Gonzalez (2018) may be moot. Although we agree with these critiques in principle, these methods may not be extensively problematic in practice, given that researchers tend not to actually use manipulation checks as mediators. Similarly, both Kidd (1976) and Hauser, Ellsworth, and Gonzalez (2018) point out potential problems with internal analysis, which involves using a manipulation check as an independent variable or predictor in a statistical analysis (see, e.g., Wilson, Aronson, & K. Carlsmith, 2010), something which was used only once in the experiments we examined. Although rarely performed, any use of internal analysis should be clearly labeled and interpreted as a correlational analysis, to avoid giving the impression that the analysis support any causal claim.

Frequentist statistics were practically ubiquitous, with 68.1% (280/411) of coded checks reporting some significance testing. However, these tests were often misused. For instance, authors frequently misinterpreted nonsignificant statistical tests as evidence in favor of the null hypothesis. Approximately 1 out of every 5 cases (21.2%, 77/364; excluding attention checks and IMCs) we coded involved an unwarranted claim of statistical equivalence based on a nonsignificant p -value (e.g., “There was no difference between Groups A and B.” “There was no effect of the manipulation on...”)⁴. About half (54.5%, 42/77) of these unwarranted inferences of equivalence were related to discriminant variables, and 29.9% (23/77) were related to measures of the independent variable⁵. Out of 207 experiments, 128 reported only a single check (61.8%).

⁴ We coded statements that excluded statistical qualifications as making a claim of equivalence. That is, we did not code the presence of a claim of equivalence if an author stated, “There was no *significant* difference between the groups.” However, if the word “significant” were excluded from such a statement, we did code the presence of a claim of equivalence.

⁵ For these measures of the independent variable, claims of equivalence often referred to nonsignificant interactions between two or more manipulations.

Of these 128, 73 (57.0%) used significance testing (20 of these 128 were attention checks and IMCs; 15.6%). In short, most experiments reported only one check, and these checks were typically assessed using significance testing.

Beyond the use of significance testing, there was strikingly little quantitative reasoning in the reports we examined. We identified only two cases for which authors provided explicit reasoning about the size of the effect on the manipulation check (one for a main effect and one for an interaction). No cases assessed manipulation check effect sizes using general benchmarks (e.g., Cohen, 1988) or empirical benchmarks (e.g., Richard, Bond, Stokes-Zoota, 2003). Although there were very few cases in which authors provided explicit quantitative assessments of manipulation strength, a random-effects meta-analysis of the size of the effects on the manipulation checks indicates those effects tend to be relatively large, with a weighted mean effect of $d = 1.58$, 95% CI [1.31, 1.85]. As we could only calculate this effect for 95 out of 134 coded manipulation checks that measured the independent variable, combined with the fact that we do not know the extent to which more ineffective manipulation checks are not being reported, this effect may be an overestimation of the general manipulation strength in social psychology.

For 41.3% (170/411) manipulation check reports, authors explicitly claimed support for the effectiveness of the manipulation. Here, we are referring to instances where authors verbally state that a manipulation is effective, as opposed to authors merely reporting a statistical test with no accompanying interpretation. Although direct and indirect manipulation checks and discriminant variables are perhaps best suited for evaluating the effectiveness of a manipulation, 22.5% of the coded stimulus checks were used to make claims of the effectiveness of a manipulation. Authors stated that a manipulation was ineffective in only two cases. Thus, in approximately 3 out of every 5 cases (58.2%), authors offered no explicit appraisal of the effectiveness of the manipulation. This rate decreases to 46.1% when only considering claims of effectiveness based on direct and indirect manipulation checks and discriminant variables. Bear in mind that our data speak only about experiments in which manipulation checks were present, so this may in fact be an overestimate of the extent to which authors report assessments of manipulations.

This review unfortunately cannot speak to the potential problem of measurement error in relation to manipulation checks (see e.g., Fayant et al., 2017; Hauser, Ellsworth & Gonzalez, 2018). While we noted that manipulation checks were measured using both single items and scales, the low extent to which researchers engaged in any kind of quantitative reasoning about their manipulation checks suggest that potential measurement issues may be neglected in a similar way (see Flake, Pek, & Hehman, 2017).

Summary and recommendations for methodological approach and statistical use. In short, our review suggests general deficits in the reported quantitative reasoning about manipulation checks. In the majority of cases, there is simply no argumentation about the evidence for the effectiveness of a manipulation, and there seem to be few safeguards against the problem of reactivity, considering the overwhelming use of self-report measures and the low frequency of pilot testing. Measuring a manipulation check in a pilot can alleviate the problem of

reactivity, as different people are responding to the manipulation check and dependent variable. Implementing manipulation checks in pilot studies may also facilitate fine tuning of a manipulation, as discriminant variables can be assessed and potentially corrected for, prior to the main experiment.

Pilot testing is enormously useful, as it provides opportunities to test the validity of manipulations in an environment in which there is no risk that the dependent measure could be biased by measuring the constructed targeted by the manipulation. In our view, pilot testing is typically more useful than using manipulation checks in main experiments (see, e.g., Fayant et al., 2017). However, we hasten to add that there are some considerable limitations to the pilots that were conducted and reported in our sample. First, many reports of pilot tests did not include important information for assessing their usefulness and generalizability (e.g., whether the pilot used a sample drawn from the same population as the one used in the main experiment; Fabrigar & Wegener, 2016). Second, the sample sizes in pilots were often quite small, limiting their power and precision. The median sample size among pilots was $N = 72.5$. For a two group comparison, this median sample size provides more than adequate power for very large effects (.99 power to detect $d = 1.00$), but it provides alarmingly low power for more subtle effects (.56 power to detect $d = .50$). We believe that pilots have the potential of providing important evidence about the validity of manipulations, but like any other study, they require adequate sampling to provide the necessary statistical power and precision. Moreover, they ought to be reported transparently, to facilitate scrutiny and criticism and also to make clear what evidence exists to support a manipulation.

Given that successfully manipulating independent variables is the *sine qua non* of experimental methodology, it is highly important that researchers take seriously the task of vetting their manipulations. As we briefly discussed above, construct validation requires substantive theorizing and thorough quantitative specification (e.g., predictions about convergent and discriminant variables). Quantitative specification was conspicuously absent in the sample of experiments we examined. Our recommendation here is straightforward (but potentially difficult in its application): Authors should present clear arguments about the extent to which a manipulation is effective, and they should be clear about the empirical (or theoretical) basis for their argument. It should not be exclusively left to the reader to assess the effectiveness of a manipulation.

We have seen that significance testing is the most common way of assessing a quantitative manipulation check. In our view, there is nothing inherently wrong with using frequentist statistics with manipulation checks. However, the general issues with the use of frequentist statistics apply here as well. A significant p -value alone does not inherently provide evidence for the effectiveness of a manipulation (see, e.g., Cohen, 1990, 1994). In our sample, authors often claimed that their manipulations were effective, apparently on the sole basis of significant p -values – often from just a single measure. Significance testing can be informative, but it should not form the entire basis of the validation of a manipulation.

Other qualities of the data, such as the size of the effect or the distribution of responses, can provide important information about the performance of a manipulation. Interpretation of the size of the effect on the manipulation check can be particularly useful. Such interpretation can (and should) go beyond simple use of benchmarks for effect sizes (e.g., Cohen, 1988) and ideally should be grounded in substantive reasoning about the meaningfulness of effects (see, e.g., Anvari & Lakens, 2019). Consider, for example, a manipulation designed to increase anger. A significant p -value could indicate that a manipulation has increased anger in a quantity that is not zero, but alone it gives no information about whether, on average, the manipulation has made people palpably angry or mildly irritated. Testing a given hypothesis may require a particular level of change in an independent variable and, thus, require reasoning about the strength of a manipulation. However, in our data, researchers almost never offered interpretations of the strength of manipulations.

Additionally, it is inappropriate to interpret nonsignificant p -values as evidence of equivalence (i.e., the absence of an effect; see Lakens, 2017; Welleck, 2010). The high frequency of inappropriate inferences of equivalence is alarming, particularly given that these errors are especially common for discriminant variables, for which researchers typically desire evidence for equivalence. That is, researchers regularly make unwarranted claims in favor of their manipulations on the basis of nonsignificant effects on discriminant variables.

Interpreting null results and failed replications. Well-implemented manipulation checks can greatly assist in the interpretation of experimental results. Hypotheses can be falsified effectively when there is good evidence for the effectiveness of a manipulation and good evidence for the absence of an effect on a dependent variable (or an effect too small to be of interest). Conversely, we might not be surprised if a manipulation has little or no effect on a dependent variable, even if the hypothesized effect is highly plausible, if the manipulation had only a small effect on a measure of the construct it was intended to manipulate. In such cases, one might argue that a result is plausibly a false negative and that a stronger manipulation is needed. Similarly, failure to influence a dependent variable when there is clear evidence for a large effect of the manipulation on the independent variable can help build the case that a hypothesis is false. Another possibility is that there is no observable effect on a manipulation check, despite there being an apparent effect on the dependent variable. Such a situation might be explained by the presence of substantial error in the manipulation check, a spurious result on the dependent variable, or the possibility that the manipulation in fact influences an unmeasured causal variable that explains the effect on the dependent variable. These possibilities would ideally be addressed empirically. But none of these arguments can be made or issues addressed if one does not look (or takes only a cursory look) at the data relevant to evaluating a manipulation. These arguments also cannot be made if no data from pilots of manipulation checks exist.

The ability of manipulation checks to inform interpretation of null effects is of course of great interest not only for primary studies, but also for the ability of replication studies to assert whether a failed replication is due to unsuccessful manipulations or a false hypothesis. The ability of both direct and conceptual replications to provide information about whether a given

effect is real depends, in part, on what Fabrigar and Wegener (2016) call “psychometric invariance” – the extent to which manipulations engage the same psychological processes when used with different people, in different settings, and at different times. If manipulations vary in their effectiveness across contexts, failure to replicate original effects may not speak to the reliability of the original study but rather be due to operationalizations not generalizing to a specific context or alternative operationalizations not capturing the intended construct equally well. As Cook and Campbell (1979; and see Shadish, Cook, & Campbell, 2002) argue, establishing generalizability requires the accumulation of empirical evidence. Thus, pilot testing a manipulation in new contexts is a key method for addressing concerns about psychometric invariance.

Additionally, when conducting replication studies where the original study did not include a manipulation check or any other assessment of construct validity, it may further be prudent to pilot test whether the original manipulation in fact manipulated the intended construct. If the original manipulation does not seem to manipulate the intended construct, there is reason to question the generality of the original findings. Depending on whether there is reason to think that the original manipulation manipulated something other than intended, or that the original manipulation does not replicate for the current sample, different actions may be considered. In the latter case, a conceptual replication study might be warranted, while in the former case, further pre-testing might give an indication as to what was originally manipulated.

Causal efficacy. None of the experiments we coded reported a quantitative assessment of causal efficacy. Researchers who have recommended the calculation and interpretation of causal efficacy have similarly noted that it is rarely reported (Breaugh, 2003; Tryon, 2001). There are many reasons causal efficacy might not have caught on following Abelson’s (1995) suggestion more than 20 years ago. One reason may be that researchers have not had benchmarks for what should be considered high or low levels of efficacy, in the same way we have, for example, Cohen’s (1988) general benchmarks for small, medium, and large effect sizes or more empirically-based benchmarks provided by Richard, Bond, and Stokes-Zoota (2003), Schäfer and Schwarz (2019), and others. In short, perhaps no one uses causal efficacy because no one knows how to put it in context. To help provide researchers with data to guide their expectations, we extracted the necessary statistics and calculated causal efficacy ratios from the experiments we coded.

Establishing causal efficacy benchmarks. As we mentioned above, for each manipulation check that measured the independent variable of interest, we extracted or calculated the size of the effects of the manipulation on the manipulation check and on the dependent variable, for a total of 112 cases. Using these “effect size” and “cause size” pairs, we calculated causal efficacy. We made two modifications to Abelson’s original calculation to facilitate combination and comparison of results across experiments. First, we used standardized measures of effect/cause size. Second, we took the natural logarithm of the ratio.

Using standardized measures is directly contrary to Abelson’s (1995) original recommendations. He specifically advocated unstandardized effect sizes and cause sizes, in order

to keep the original scale units and thus facilitate more concrete interpretation (e.g., “A change in 2 scale points on the manipulation check corresponds to a change of 5 points on the dependent variable.”). Unstandardized measures are indeed desirable for interpreting the results of an individual experiment, but statistics that are independent of measurement scale are necessary for comparison and combination⁶.

Log ratios are useful because raw ratios are problematic to summarize and compare, since they behave asymmetrically on either side of 1: They extend infinitely in a positive direction (above 1) and asymptotically approach zero (below 1). Although raw ratios can lend themselves to more intuitive interpretation when examining single cases, log ratios are more mathematically useful, as they behave symmetrically about 0: Positive values indicate a larger numerator, and negative values indicate a larger denominator. Thus, in the case of log causal efficacy, positive values indicate a larger effect on the dependent variable relative to the manipulation check, and negative values indicate a larger effect on the manipulation check relative to the dependent variable. The log transformation allows us to calculate meaningful summary statistics, such as the mean and median, which would be distorted if calculated for the raw ratios.

Across the 112 cases we coded, the mean log causal efficacy was -1.34 ($SD = 1.17$, $median = -1.17$). This mean equates to a raw ratio of .26. In other words, on average, effect sizes are less than one-third of their respective cause size (as measured by a manipulation check). Figure 1 illustrates the frequency distribution of log causal efficacy values. To establish benchmarks for low, medium, and high efficacy, we can examine 33.33% quantiles. These quantiles are presented in Table 3.

⁶ In our data, for consistency, we converted all effect and cause sizes to standardized mean differences (d). However, causal efficacy can be calculated using any effect size measure that does not have upper or lower bounds (e.g., r is bounded at -1 and 1; odds ratios and explained variance measures are bounded at 0) and behaves symmetrically about 0. Causal efficacy also requires that the effect and cause sizes use the same measure (e.g., both d).

Figure 1. *Log causal efficacy in experimental social psychology*

Note: Dotted lines represent 33.33% quantile breaks. A heavy dashed line is drawn at 0, to the right of which effect sizes are larger than cause sizes and to the left of which cause sizes are larger than effect sizes.

Table 3.

Benchmarks for causal efficacy

Quantile	Lower bound	Median	Upper bound
Low efficacy (0% to 33.33%)	-5.23 (0.005)	-2.29 (0.10)	-1.65 (0.19)
Medium efficacy (33.33% to 66.66%)	-1.65 (0.19)	-1.18 (0.31)	-0.76 (0.47)
High efficacy (66.66% to 100%)	-0.76 (0.47)	-0.20 (0.82)	0.65 (1.91)

Note: Values are log causal efficacy, with raw ratios in parentheses.

Practical applications. Considering causal efficacy has a straightforward pragmatic benefit: It forces researchers to think about the strength of their manipulations (Abelson, 1995). Currently, researchers report almost no explicit reasoning on this issue – like arguing about the effectiveness of a medication while saying nothing about dosage. Consider a hypothetical case in which causal efficacy is very low (cause size $d = 4.00$, effect size $d = .50$). In such a case, researchers must consider the plausibility that real-world instantiations of the independent variable are stronger or weaker than the manipulation. If they are weaker, perhaps an effect size

that might otherwise be regarded as large enough to be practically meaningful ($d = .50$) should be interpreted with caution, since the effect outside the laboratory may be much smaller.

Beyond encouraging better reasoning, these benchmarks for causal efficacy may be useful for planning experiments, especially when researchers want to conduct *a priori* power calculations. Causal efficacy considerations can make pilot testing even more useful. For example, if you perform a pilot test of a manipulation and take a direct measure of the construct of interest (prior to conducting an experiment to test a substantive hypothesis), you can obtain an estimate of the cause size. With that estimate and benchmarks for causal efficacy, it is possible to take a reasonable guess about the plausible size of the effect on the dependent variable. Going beyond general benchmarks like those provided here, it is also possible to examine cause sizes (and causal efficacy) for similar manipulations reported in previous experiments, in order to obtain more specific estimates applicable to your particular research question. These estimates can further guide expectations about effect sizes and thus inform sample size planning and power calculations. For example, if you are using a manipulation that is weaker than those used in previous experiments and base your planned sample size to achieve adequate power for effect sizes observed in past research, you may inadvertently underpower the experiment.

The middle quantile of log efficacy ratios in our data, where scores are the densest, suggests that what we might call medium efficacy effects range from about one-half to one-fifth the size of their causes. It is relatively rare that effect sizes exceed cause sizes, occurring in only about 1 out of 10 cases (10.7%, 12/112). Thus, researchers should probably generally expect effects on their dependent variable to be substantially smaller than effects on their manipulation check.

Limitations. Unfortunately, at present, we know little or nothing about what predicts causal efficacy. It is possible, for instance, that causal efficacy is influenced by the properties of the dependent measure. However, we have no evidence of such differences in our data. As we extracted the statistics from articles, we coded whether the dependent variable was behavioral (e.g., frequency of a specific behavior, reaction time; 32 variables) or perceptual (e.g., self-reported attitudes; 77 variables). Behavioral and perceptual measures had similar mean log causal efficacy (-1.51 and -1.22, respectively), but we have insufficient data to make a claim that these means are statistically equivalent. Absent specific evidence, it is difficult to argue compellingly that a given manipulation should be expected to be of low, medium, or high efficacy. If you must resort to a heuristic to guess about an effect size, however, it is probably not unreasonable to suppose that it will be about one-fifth to one-half of the estimated cause size.

To what extent can we generalize from these estimates of causal efficacy? There are several reasons to believe that they should be applicable to contemporary social psychology experiments. First, we sampled from top-ranking social psychology journals, which should be broadly representative of the currently dominant topics and methods. Second, conducting a random effects meta-analysis on the effect sizes (on the dependent variables) for our 112 cases, we find a mean of $d = .49$, 95% CI [0.42, 0.57]. This is quite similar (but slightly higher) than the overall meta-meta-analytic effect size in social psychology estimated by Richard, Bond, and

Stokes-Zoota (2003; $d = .42$). In fact, when splitting the sample of effect sizes by the journals in which they were published, we find that *JESP* and *PSPB* have mean effects of $d = .42$, 95% CI [0.17, 0.67], and $d = .43$, 95% CI [0.23, 0.63], respectively – nearly exactly the meta-meta-analytic average. *JPSP* stands out with a higher mean effect of $d = .57$, 95% CI [0.39, 0.75]. Quantitatively, our data are fairly representative of social psychology.

However, we cannot be certain the extent to which these results are applicable to other domains of psychology (e.g., cognitive psychology). It may be a worthwhile future endeavor to review other subfields in order to establish benchmarks for those research areas. Additionally, it is possible that our sample contains many overestimates of effect sizes and cause sizes. If manipulation checks or dependent variable were selectively reported, the estimates may be inflated (Rosenthal, 1979). Moreover, relatively few of the experiments we reviewed were preregistered, and there is reason to believe that unregistered studies tend to find larger effects than registered studies (Schäfer & Schwarz, 2019). But these are problems that have less to do with our data in particular and more to do with reporting practices in psychological science at large (Bakker et al., 2012). To investigate the potential issue of inflated effect sizes, it would be useful for future research to assess causal efficacy using data from replication projects and registered reports.

Finally, measures of causal efficacy rest on several assumptions. First, both the manipulation check and the dependent variable must be reliably measured. A causal efficacy ratio greater than 1 (or a log ratio greater than 0) could indicate that we are observing a small cause leading to a large effect, or it could indicate that the manipulation check was measured with considerably more error than the dependent variable. Alternatively, if one does not account for construct confounding, the observed effect on the dependent variable could be due to variables not captured by the cause size estimate, thus inflating causal efficacy. Second, the estimates of cause size and effect size are assumed to be unbiased. If a manipulation check evoked demand characteristics, this may lead to an underestimate of causal efficacy (by inflating the cause size), while reactivity to the manipulation check may instead produce an overestimate of causal efficacy (by inflating the effect size). Here, our data are taken primarily from studies in which the manipulation check was included in the main study.

Taken together, these issues are reason to err on the side of being relatively conservative in the use of causal efficacy measures. If more cause size data were available from pilots and pretests (which safeguard against reactivity effects on the dependent measure), it would be possible to more directly assess the severity of these potential problems.

Format and style of reports

How do researchers report their manipulation checks? One striking feature of manipulation check reports is their brevity. They are typically reported in only a brief subsection of the results or even in a footnote. As we noted above, it is common for authors to offer no interpretation of the meaning of the outcome of the manipulation check. There are other signs that little space and attention is devoted to manipulation checks.

Out of our 411 coded checks, only three appeared in short reports rather than regular articles. Although we do not have the data to draw strong conclusions about the reason for this, it may reflect a constraint imposed by the short report format. That is, perhaps the format does not allow sufficient space for transparent and thorough reporting of manipulation checks or pilot testing and instead prioritize streamlined presentations of methods and results, at the expense of nuance (see Giner-Sorolla, 2012; McGuire, 2013).

Perhaps authors and editors in social psychology hesitate to clutter articles with too much information and uninteresting detail about the evidence supporting a manipulation. If so, we might expect to find information about manipulation checks in supplementary material. However, only 7 of our coded manipulation checks were found in supplementary material. We cannot determine, however, whether there were few cases in supplementary material because authors are not taking advantage of the opportunity to archive existing but unreported information or because there is simply no information to store (viz. manipulation checks are reported when they are used but tend to be reported very briefly). Neither possibility is desirable.

Conclusions

Current methodological practices intended to assess the validity of experimental manipulations leave much to be desired. Manipulation checks pervade social psychology (Sigall & Mills, 1998; Hauser, Ellsworth, & Gonzalez, 2018), but they are typically accompanied by little argumentation and reasoning, sometimes statistically misinterpreted, and allocated little space and attention. That being said, we believe many of the complaints lodged at manipulation checks have been somewhat misaimed. In some cases, criticisms apply to practices that rarely or never actually occur (at least in our data), such as using manipulation checks for mediation or internal analysis. In other cases, critiques have missed some of the potential benefits of manipulation checks. For example, Sigall and Mills (1998) argued that in cases in which there is no plausible alternative explanation for an effect, manipulation checks are unnecessary because they add no further evidence that the construct of interest was manipulated. But even under these circumstances, quantitative checks may still be informative about the strength of the manipulation and can potentially be used for assessments of causal efficacy.

However, recall that very few manipulation checks were reported in pilot testing and that there is a general lack of reported reasoning about manipulation checks. Taken together, these tendencies in reporting practices suggest that building arguments for the effectiveness of manipulations is not being given nearly enough recognition as a vital part of the experimental process. It is plausible that the low degree of attention given to the quality of manipulations is due in large part to (1) incentives to be concise in scientific reporting and (2) a *lack* of incentives for performing and reporting rigorous vetting of manipulations (see Smaldino & McElreath, 2016).

In the past, high-profile researchers have advised against the reporting of preliminary work, despite its importance (e.g., Bem, 2003). For example, Wilson, Aronson, and K. Carlsmith (2010) state

[Pretesting manipulations] is where some of the most important work in experimental design occurs, but it is rarely reported in published articles, because it would not be very informative or interesting to begin the methods section by saying, “We will first tell you about all the ways of manipulating the independent variable that didn’t work...” ...It is another matter, however, if the manipulation works as intended but does not influence the dependent variable in the predicted manner... This *is* informative. (p.71, emphasis in the original)

Prioritizing the reporting of effects on the dependent measure over (un)successful manipulations is understandable, insofar as one of the main objectives of science is to test hypotheses and generate new knowledge. But critical voices in psychological science have recently emphasized the importance showing the blemishes of a rigorous scientific process, as opposed to presenting a good, streamlined narrative (Bakker et al., 2012; Giner-Sorolla, 2012). Our perspective is decidedly more in line with the latter view than the former.

In addition to explicit statements about priorities in publication, current practices may send implicit signals to researchers that preliminary work is simply not valued, since it is rarely showcased in published articles and thus not rewarded. To borrow some rhetoric from John Tukey (1969), current practices might send the message that assessing manipulations is “something done hastily [and] ...done to show that one’s work is acceptable, rather than something carefully done” (p.89)⁷. Why would a researcher go out of her or his way to extensively pilot test a manipulation and carefully implement manipulation checks, knowing that no one would value the labor and knowing that she or he could get away with doing much less? One obvious incentive for validation is avoiding the use of faulty manipulations and wasting resources on flawed ideas. This incentive is, however, already in operation, and it does not appear to have been sufficient to prompt thorough validation of manipulations. Therefore, to improve practice, we may need not only to provide researchers with the possibility of thoroughly reporting preliminary work to support the effectiveness of a manipulation but also strongly encourage them to do so. For researchers to provide this kind of detailed descriptions of their manipulation checks and pilot testing, editors and reviewers may need to specifically request this material and appropriately offer criticism when such work has not been conducted.

There would be several benefits to such a shift in priorities when evaluating experiments. Complete information about pilots can also be useful for evaluating the generality of conclusions and construct validity. For example, if researchers severely struggled to find a manipulation that “worked,” perhaps the phenomenon under examination only occurs under select circumstances. By reporting all pilots and preliminary work, it would also allow other researchers to avoid wasting resources repeating unsuccessful manipulations. Moreover, such information would also paint a more realistic image of how much work is necessary to conduct rigorous experimental research. Along similar lines, McGuire (2013) bemoaned the fact that the preliminary work that

⁷ Tukey was not talking about manipulation checks here. Rather, he was arguing against treating data analysis as a perfunctory ritual rather than careful detective work (see, e.g., Gigerenzer, 2004). In our view, the same critiques apply to the vetting of experimental manipulations.

precedes many experimental manipulations may contain valuable information about the boundary conditions of the manipulation (as well as the theory being studied), and this work is rarely reported. To the extent that the careful conduct, documentation, and reporting of preliminary work is treated as a desirable component of scientific labor, there will be incentives for producing more rigorous and more informative publications (Fraley & Vazire, 2014; Kraus, 2013; Vazire, 2018)⁸.

Many journals now allow for virtually unlimited space in online supplementary sections, and authors can also use external repositories (e.g., Open Science Framework, GitHub) to store their supplemental information. To the extent that researchers have additional data that bears on the effectiveness of their manipulations but may not be reporting it, supplementary materials are a largely untapped resource. Editorial policy might understandably prioritize allocating space only to information that most research consumers would find interesting, and this could constrain space for thorough reporting of painstaking preliminary work. If this is the case, researchers can take advantage of online repositories to provide detailed descriptions of all checks of the manipulation and all the pilot tests and preliminary studies that were conducted in relation to the tests of their hypotheses.

Individual initiatives, such as the use of supplemental material repositories, may be useful, but journals and editors also have the capacity to provide substantial incentives for work on construct validation for manipulations. Similar incentives already exist for the validation of psychological measures; indeed, some journals are entirely devoted to the subject. Developing a similar tradition for the validation and scrutiny of experimental manipulations could dramatically improve research practice by giving researchers incentives, in the form of citable publications, to validate manipulations.

At least two obstacles stand in the way of such an initiative. First, journals in social psychology do not seem to encourage manipulation validation. Second, to the best of our knowledge, no specific guidelines exist for the conduct and reporting of such validation. We who are not editors of journals can do little about the former problem. However, to begin to address the latter problem, we have developed a template for the preregistration of manipulation validation studies⁹. The template can be accessed online here: <https://osf.io/p6szz/>. The registration template serves two key objectives: (1) providing researchers with a guide to planning pilot tests to validate manipulations and (2) limiting researchers' flexibility in data analysis and reporting of validation studies, thus increasing the credibility of claims (see, e.g., Nelson, Simmons, & Simonsohn, 2018; Simmons, Nelson, & Simonsohn, 2011). It is based, in part, on popular preregistration templates but is specifically tailored to emphasize

⁸ Though above we criticize the position articulated by Wilson, Aronson, and K. Carlsmith (2010), the analogous chapter by Aronson and J. M. Carlsmith (1968) in the second edition of the *Handbook* expresses ideas quite similar to ours: "An experimenter can decrease the probability of failure if he [sic] pretests each aspect of his procedure, strengthening his independent variable... This renders failure more meaningful and, hence, more conclusive... One cannot overemphasize the importance of pretesting" (p.74).

⁹ We are grateful to Roger Giner-Sorolla for suggesting the creation of guidelines for reporting manipulation validation studies.

methodological issues related to the assessment of manipulations, particularly structural and external validation (Loevinger, 1957). Part of its potential usefulness derives from encouraging researchers to think about the extent to which their proposed methods contribute to the construct validation process and alerts them to common methodological problems related to manipulation checks. For example, because the purpose of manipulation check measures can often be transparent to participants, the document specifically prompts authors to note how, if at all, they are safeguarding against demand effects. Additionally, to promote consideration of how constructs have been conceptualized and of potential limits on the generalizability of manipulations, the template also prompts authors to provide background information about any substantive validation that has preceded the registered test. We hope that this template will serve as a helpful tool for researchers.

The publication of articles focused on manipulation validation could not only encourage researchers to thoroughly vet their manipulations, but could also lead to a more efficient use of resources. Unpublished pilot tests may provide the individual researcher with important information, but published pilot studies will provide important information to a whole community of researchers, who will not have to repeat the same mistakes. Moreover, researchers who spend time and resources on the vetting of manipulations receive due credit.

Summary of recommendations. Throughout our review, we have offered numerous suggestions for how to avoid the hazards and how to reap the benefits of manipulation checks. Table 4 summarizes each of our recommendations for practice, organized according to the steps in an idealized research process in which they are relevant.

Most of our suggestions are directed at researchers planning and conducting experiments, but these recommendations are also important for reviewers and editors, whose task it is to scrutinize reports of experiments and, in our opinion, should request that authors implement better practices. The recommendations provided in Table 4 complement the registration template we have designed for pilot testing of manipulations. That is, authors conducting manipulation validation research and reviewers of such research can use these as a set of guidelines for reporting and critiquing such studies.

We find it unlikely that better practices will be adopted widely until such time that prescriptive norms take hold to demand them (see, e.g., Smaldino & McElreath, 2016). Reviewers and editors are the enforcers of such norms, as the gatekeepers of publication. More rigorous standards are likely to be demanding and may implicitly require researchers to slow down and do (and report) more “legwork” with piloting and pretesting. Given that productivity in the form of publications is highly incentivized in science today, researchers may find new demands unappealing (Nosek, Spies, & Motyl, 2012; Tullett, 2015; Vazire, 2018). Such demands may, however, be necessary and may be made more palatable if journals encourage the publication of validation studies.

Many critics have appropriately called for more attention to be paid to the process of validating our methods (e.g., Fayant et al., 2017; Flake, Pek, & Hehman, 2017; Hauser, Ellsworth, & Gonzalez, 2018), with the hope that researchers will improve their practices. To be

realized, this hope must be given teeth. It is rarely safe to assume that a manipulation will be effective – nor is it safe to assume that methodological weaknesses will be improved without intervention.

Table 4.

Summary of recommendations

Consideration	Recommendation
<hr/> <i>Guidelines for conducting and reporting manipulation checks and validation</i> <hr/>	
<i>Concept & terminology</i>	<p>Consistent use of terms</p> <ul style="list-style-type: none"> ● Select type of manipulation check according to purpose, e.g., <ul style="list-style-type: none"> - Assess impact on construct of interest: measure of independent variable (<i>direct manipulation check</i>) or <i>indirect manipulation check</i> - Assess construct confounding: measure <i>discriminant variables</i> - Assess attention to manipulation: <i>stimulus checks/treatment checks</i> - Assess general attention: <i>general</i> or <i>instructional attention checks</i> ● Clearly identify and report the purpose of checks. ● Label manipulation checks clearly and consistently.

<i>Use & interpretation</i>	Methodological approach	<ul style="list-style-type: none"> • Consider potential measurement error (e.g., use validated scales or pretest scales). • Consider how potential demand effects could affect results (e.g., use indirect manipulation checks, conceal manipulation check among other measures) • Consider the extent to which the manipulation will function as intended in other contexts (psychometric invariance)
	Statistical use	<ul style="list-style-type: none"> • Perform power calculations for manipulation checks (especially for pilots) • Preregister decisions on statistical use of manipulation checks • Avoid unwarranted statistical inferences • Evaluate discriminant variables using appropriate statistical methods (e.g., equivalence tests). • Do not rely solely on significance testing for interpretation of manipulation checks; interpret the magnitude of effects
	Format and style of report	<ul style="list-style-type: none"> • Report manipulation checks and pilots as detailed as the main results (e.g., number of participants, effect size measures etc) • Present explicit appraisal (theoretical and/or statistical) of the effectiveness of manipulations. • When space is limited, use supplementary material sections and online repositories

Guidelines for planning experiments

- Consider causal efficacy when conducting power calculations for main experiment
 - To avoid reactivity , use pilot tests whenever possible
 - If using a check in the main experiment, account for reactivity (e.g., vary order of measures)
-

References

- Abbey, J. D. & Meloy, M. G. (2017). Attention by design: Using attention checks to detect inattentive respondents and improve data quality. *Journal of Operations Management*, 53, 63-70.
- Abelson, R. P. (1995). *Statistics as Principled Argument*. Hillsdale, NJ: Lawrence Erlbaum.
- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. Washington, DC: Joint Committee on Standards for Educational and Psychological Testing.
- Anvari, F. & Lakens, D. (2019). Using Anchor-Based Methods to Determine the Smallest Effect Size of Interest. <https://doi.org/10.31234/osf.io/syp5a>
- Aronson, E. (2010). *Not by Chance Alone: My Life as a Social Psychologist*. New York, NY: Basic Books.
- Aronson, E. & Carlsmith, J. M. (1968). Experimentation in social psychology. In G. Lindzey & E. Aronson (Eds.), *Handbook of Social Psychology* (2nd ed., pp. 1-79). Reading, MA: Addison-Wesley.
- Aronson, E., Ellsworth, P., Carlsmith, J. M., & Gonzales, M. H. (1990). *Methods of Research in Social Psychology* (2nd ed.). New York: McGraw-Hill.
- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, 7, 543-554.
- Bem, D. (2003). Writing the Empirical Journal Article. In J. M. Darley, M. P. Zanna, and H. L. Roediger III (Eds.), *The Compleat Academic: A Practical Guide for the Beginning Social Scientist* (2nd ed.). Washington, DC: American Psychological Association.
- Breaugh, J. A. (2003). Effect size estimation: Factors to consider and mistakes to avoid. *Journal of Management*, 29, 79-97.
- Carlsmith, J. M., Ellsworth, P., & Aronson, E. (1976). *Methods of Research in Social Psychology* (1st ed.). Reading, MA: Addison-Wesley.
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49, 997-1003.
- Cohen, J. (1990). Things I have learned (so far). *American Psychologist*, 45, 1304-1312.
- Cohen, J. (1988). *Statistical power for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Lawrence Erlbaum.
- Cook, T. D. & Campbell, D. T. (1979). *Quasi-Experimentation: Design and Analysis Issues for Field Settings*. Boston: Houghton Mifflin.
- Cronbach, L., & Meehl, P. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52, 281-302.
- Fabrigar, L. R., & Wegener, D. T. (2016). Conceptualizing and evaluating the replication of research results. *Journal of Experimental Social Psychology*, 66, 68-80.
- Fayant, M. P., Sigall, H., Lemonnier, A., Retsin, E., & Alexopoulos, T. (2017). On the Limitations of Manipulation Checks: An Obstacle Toward Cumulative Science. *International Review of Social Psychology*, 30, 125-130.

- Festinger, L. (1953). Laboratory experiments. In L. Festinger & D. Katz (Eds.), *Research Methods in the Behavioral Sciences* (pp. 136-172). New York: Dryden Press.
- Fiedler, K., Schott, M., & Meiser, T. (2011). What mediation analysis can (not) do. *Journal of Experimental Social Psychology*, 47, 1231-1236.
- Foschi, M. (2014). Hypotheses, operationalizations, and manipulation checks. In: Webster, M. J., & Sell, J. (Eds.), *Laboratory Experiments in the Social Sciences* (2nd ed., pp. 247–268). San Diego, CA: Elsevier Academic Press.
- Flake, J. K., Pek, J., & Hehman, E. (2017). Construct validation in social and personality research: Current practice and recommendations. *Social Psychological and Personality Science*, 8, 370-378.
- 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.
- Gigerenzer, G. (2004). Mindless statistics. *The Journal of Socio-Economics*, 33, 587-606.
- 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, 562-571.
- Haslam, S. A. & McGarty, C. (2019). *Research Methods and Statistics in Psychology* (3rd ed.). Los Angeles: Sage.
- Hauser, D. J., Ellsworth, P. C., & Gonzalez, R. (2018). Are Manipulation Checks Necessary?. *Frontiers in Psychology*, 9, 998.
- Hauser, D. J. & Schwarz, N. (2015). It's a trap! Instructional manipulation checks prompt systematic thinking on "tricky" tasks. *Sage Open*, 5(2), 2158244015584617.
- Kelley, H. H. (1951). Communication in experimentally created hierarchies. *Human Relations*, 4, 39-56.
- Kidd, R. F. (1976). Manipulation checks: advantage or disadvantage? *Representative Research in Social Psychology*, 7, 160-165.
- Kraus, M. (2013). *Quality v. quantity in publication*. Available at <http://psych-your-mind.blogspot.com/2013/06/quality-v-quantity-in-publication.html>. Accessed 2019 June 29.
- Lakens, D. (2017). Equivalence tests: A practical primer for t tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, 8, 355-362.
- Loevinger, J. (1957). Objective tests as instruments of psychological theory. *Psychological reports*, 3, 635-694.
- McGuire, W. J. (2013). An additional future for psychological science. *Perspectives on Psychological Science*, 8, 414-423.
- Mills, J. (1969). *Experimental Social Psychology*. Toronto: MacMillian.
- Nelson, L. D., Simmons, J., & Simonsohn, U. (2018). Psychology's renaissance. *Annual review of psychology*, 69, 511-534.

- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific utopia: II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7, 615–631.
- Nunnally, J.C. (1970). *Introduction to psychological measurement*. New York: McGraw-Hill.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45, 867-872.
- Parro, W. G. & Hertel, P. (1999). Research methods in cognition and emotion. In T. Dalgleish & M. J. Power (Eds.), *Handbook of cognition and emotion* (pp. 61-81). Chichester, England: Wiley.
- Prentice, D. A. & Miller, D. T. (1992). When Small Effects Are Impressive. *Psychological Bulletin*, 112, 160-164.
- Raykov, T. & Marcoulides, G. A. (2011). *Introduction to psychometric theory*. Routledge.
- Reise, S. P., Widaman, K. F., & Pugh, R. H. (1993). Confirmatory factor analysis and item response theory: Two approaches for exploring measurement invariance. *Psychological bulletin*, 114, 552-566.
- 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, 331-363.
- Rosenthal, R. (1979). The “File Drawer Problem” and Tolerance for Null Results. *Psychological Bulletin*, 86, 638-641.
- Schäfer, T. & Schwarz, M. (2019). The meaningfulness of effect sizes in psychological research: Differences between sub-disciplines and the impact of potential biases. *Frontiers in Psychology*, 10, 813.
- Schwarz, N. & Strack, F. (1991). Context effects in attitude surveys: Applying cognitive theory to social research. In W. Stroebe & M. Hewstone (Eds.), *European Review of Social Psychology* (Vol. 2, pp. 31-50). Chichester: Wiley.
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Belmont, CA: Wadsworth Cengage.
- Sigall, H., & Mills, J. (1998). Measures of independent variables and mediators are useful in social psychology experiments: But are they necessary? *Personality and Social Psychology Review*, 2, 218-226.
- 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.
- Smaldino, P. E., & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, 3(9), Article 160384.
- Tryon, W. W. (2001). Evaluating Statistical Difference, Equivalence, and Indeterminacy Using Inferential Confidence Intervals: An Integrated Alternative Method of Conducting Null Hypothesis Statistical Tests. *Psychological Methods*, 6, 371-386.

- Tukey, J. W. (1969). Analyzing data: Sanctification or detective work? *American Psychologist*, 24, 83-91.
- Tullett, A. M. (2015). In search of true things worth knowing: Considerations for a new article prototype. *Social and Personality Psychology Compass*, 9, 188-201.
- Vazire, S. (2018). Implications of the credibility revolution for productivity, creativity, and progress. *Perspectives on Psychological Science*, 13, 411-417.
- Wellek, S. (2010). *Testing statistical hypotheses of equivalence and noninferiority*. Chapman and Hall/CRC.
- Westlund, E., & Stuart, E. A. (2017). The nonuse, misuse, and proper use of pilot studies in experimental evaluation research. *American Journal of Evaluation*, 38, 246-261.
- Wilson, T., Aronson, E., & Carlsmith, K. (2010). The Art of Laboratory Experimentation. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of Social Psychology* (Vol 1., 5th ed., pp. 51-81). Hoboken, NJ: Wiley.