

Construct Validation of Experimental Manipulations in Social Psychology: Current Practices and Recommendations for the Future

David S. Chester^{ID} and Emily N. Lasko

Department of Psychology, Virginia Commonwealth University

Perspectives on Psychological Science
1–19

© The Author(s) 2020

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/1745691620950684

www.psychologicalscience.org/PPS



Abstract

Experimental manipulations in social psychology must exhibit construct validity by influencing their intended psychological constructs. Yet how do experimenters in social psychology attempt to establish the construct validity of their manipulations? Following a preregistered plan, we coded 348 experimental manipulations from the 2017 issues of the *Journal of Personality and Social Psychology*. Representing a reliance on “on-the-fly” experimentation, the vast majority of these manipulations were created ad hoc for a given study and were not previously validated before implementation. A minority of manipulations had their construct validity evaluated by pilot testing before implementation or via a manipulation check. Of the manipulation checks administered, most were face valid, single-item self-reports, and only a few met criteria for “true” validation. In aggregate, roughly two fifths of manipulations relied solely on face validity. To the extent that they are representative of the field, these results suggest that best practices for validating manipulations are not commonplace—a potential contributor to replicability issues. These issues can be remedied by validating manipulations before implementation using validated manipulation checks, standardizing manipulation protocols, estimating the size and duration of manipulations’ effects, and estimating each manipulation’s effects on multiple constructs within the target nomological network.

Keywords

construct validation, experimental manipulation, social psychology, manipulation check, meta-analysis

Social psychology emphasizes the power of the situation (Lewin, 1939). To examine the causal effects of situational variables, social-psychological studies often use experimental manipulations of such factors and examine their impact on human thoughts, feelings, and behaviors (Campbell, 1957; Cook & Campbell, 1979). However, experimental manipulations are useful only to the extent that they exhibit construct validity (i.e., that they meaningfully affect the psychological processes that they are intended to affect; Brewer, 2000; Garner, Hake, & Eriksen, 1956; Wilson, Aronson, & Carlsmith, 2010). Yet few recent studies have systematically documented the approaches that social-psychological experiments use to estimate and establish the construct validity of their manipulations. In an effort to address this limitation in our understanding, we meta-analyzed the frequency with which various manipulation validation practices were adopted (or not adopted) by a

representative sample of studies from what is widely perceived as the flagship publication for experimental social psychology: the *Journal of Personality and Social Psychology* (*JPSP*).

Validity in Experimental Manipulations of Psychological Processes

Experimental social psychologists often focus on *internal validity* and *external validity* (Haslam & McGarty, 2004). Internal validity is present when experimenters (a) eliminate extraneous variables that might incidentally influence the outcome of interest and (b) maximize

Corresponding Author:

David S. Chester, Department of Psychology, Virginia Commonwealth University

E-mail: dschester@vcu.edu

features of the experimental manipulation that ensure a precise, causal conduit from manipulation to outcome (Brewer, 2000). Experimenters establish internal validity via practices such as removing sources of experimenter bias and demand characteristics and by cultivating *experimental realism* to maximize the chances that the manipulation is the source of experimental effects and not some unwanted artifact of design (Cook & Campbell, 1979; Wilson et al., 2010). Other efforts are directed toward maximizing external validity to ensure that the experiment captures effects that exist in the real world and that findings of the experiment can generalize to other settings, populations, time periods, and cultures (Highhouse, 2009; c.f. Berkowitz & Donnerstein, 1982; Mook, 1983). Integral to both internal and external validity is a concept most often invoked in the context of clinical assessments and personality questionnaires—*construct validity*.

Psychological Constructs and the Nomological Network

Psychological scientists often seek to measure and manipulate psychological constructs, which are psychological entities constructed by people and not objective realities (Cronbach & Meehl, 1955). Constructs are considered latent because they are readily imperceptible compared with the associated manifestations that are designed to capture (e.g., psychological questionnaires) or influence (e.g., experimental manipulations) them. Latent constructs exist in a nomological (i.e., lawful) network, which is a prescribed array of relationships (or lack thereof) to other constructs (Cronbach & Meehl, 1955). In a nomological network, constructs exist in varying degrees of proximity to one another, with closer proximities reflecting stronger patterns of association. Each construct has its own idiographic network that includes construct-specific arrays of associated constructs and construct-specific patterns of associations with those constructs. The constellations of constructs within each nomological network are articulated by psychological theory (Gray, 2017). Nomological networks, when distilled accurately from strong theory, are the basis of construct validity (Messick, 1995).

Construct Validity of Psychological Measures

Construct validity is a methodological and philosophical property that largely reflects how accurately a given manifestation of a study has mapped onto a construct's latent nomological network (Borsboom, Mellenbergh, & van Heerden, 2004; Embretson, 1983; Strauss & Smith, 2009). Conventionally, construct validity has been

largely invoked in the context of psychological measurement, assessment, and tests. In this context, construct validity is present when a manifest psychological measure (a) accurately quantifies its intended latent psychological construct, (b) shares theoretically appropriate associations with other latent variables in that construct's nomological network, and (c) does not capture confounding extraneous latent constructs (Cronbach & Meehl, 1955; Messick, 1995; Fig. 1). According to modern standards in psychology, construct validity does not pertain to a property of a given measure or the scores derived from it but instead to the uses and interpretations of the scores that are derived from the measure (American Educational Research Association, American Psychological Association, & National Council on Measurement in Education, 2014).

As depicted in Figure 1, a measure of a given construct (e.g., a scale that measures feelings of rejection) should exhibit a pattern of associations with theoretically linked variables (e.g., positive correlations with pain and shame, negative correlations with happiness) and null associations with variables outside of the nomological network (e.g., awe).

Estimating the Construct Validity of Psychological Measures

The process of testing the construct validity of measures is well defined (for an overview, see Flake, Pek, & Hehman, 2017). First, investigators should conduct a comprehensive literature review to define the properties of the construct, prominent theories of the construct, and its associated nomological network (Simms, 2008). This substantive portion of construct validation and research design more broadly is perhaps the most crucial (and often neglected) aspect. Rigorous theoretical work is needed before constructing a measure to ensure that the manifestation of the measure accurately captures the full range of the construct, distinguishes it from related constructs, and includes measures of other constructs to test the construct's nomological network (Benson, 1998; Loevinger, 1957; Zumbo & Chan, 2014).

Second, researchers apply their theoretical understanding to design the content of the measure to capture the breadth and depth of the construct (i.e., content validity; Haynes, Richard, & Kubany, 1995), often in consultation with experts outside of the study team. Third, this preliminary measure is administered, and empirical analyses (e.g., item-response theory, exploratory and confirmatory factor analyses) are used on the resulting data to (a) ensure that the measure's data structure exhibits the expected form, (b) select content with good empirical qualities, and (c) ensure the measure is invariant across the groups it should be invariant

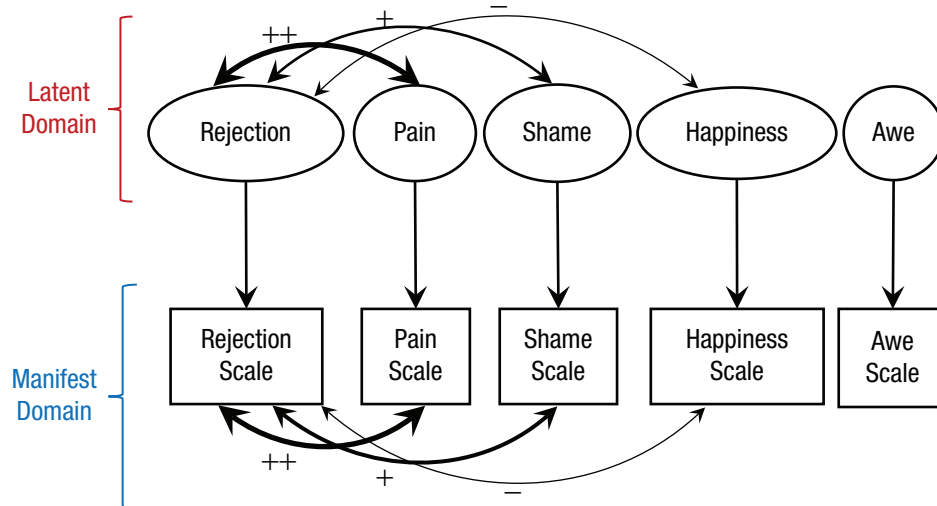


Fig. 1. Schematic depiction of a hypothetical nomological network surrounding the construct of rejection. Plus signs depict positive associations; minus signs depict negative associations. Greater numbers of plus signs and thicker arrows depict stronger associations and effects.

across (Clark & Watson, 2019). Fourth, a refined version of the measure is administered alongside other measures to ensure that it (a) positively corresponds to measures of the same or similar constructs (i.e., convergent validity), (b) negatively or weakly corresponds to measures of different or dissimilar constructs (i.e., discriminant validity), (c) is linked to theoretically appropriate real-world outcomes (i.e., criterion validity), and (d) differs across groups as it should (G. T. Smith, 2005). Measures that meet these stringent psychometric criteria can be said to exhibit construct validity (i.e., they measure the construct they are intended to measure and do not capture problematically large amounts of unintended constructs). Yet how do these concepts and practices translate to experimental manipulations of psychological processes?

Construct Validity of Psychological Manipulations

Construct validity is not confined to psychometrics and is a crucial element in experimental psychology (Cook & Campbell, 1979). Translated to an experimental setting, construct validity is present when a manifest psychological manipulation (a) accurately and causally affects its intended latent psychological construct in the intended direction, (b) exerts theoretically appropriate effects on other latent variables in that construct's nomological network, and (c) does not affect or weakly affects confounding extraneous latent constructs (Campbell, 1957; Shadish, Cook, & Campbell, 2002). This desired pattern of effects is illustrated in a phenomenon we deem the nomological shockwave.

The nomological shockwave

In a nomological shockwave, a psychological manipulation (e.g., a social-rejection manipulation; Chester, DeWall, & Pond, 2016) exerts its initial and strongest causal effects on the target latent construct in the intended direction (e.g., greatly increased feelings of rejection; Fig. 2). This change in the target construct then ripples out through that construct's latent nomological network—causally affecting related constructs in ways that reflect the degree and strength of their latent associations with the target construct. More specifically, the shockwave exerts stronger effects on constructs that are closer to the manipulation's point of impact (e.g., moderately increased pain). Conversely, the shockwave's effects get progressively weaker as the theoretical distance from the target construct increases (e.g., modestly increased shame, modestly reduced happiness). The shockwave will not reach constructs that lie beyond the target construct's nomological network (e.g., no effect on awe). Back in the manifest domain, these latent shockwave effects are then captured with a manipulation check and the various discriminant validity checks that are causally affected by the latent nomological shockwave.

Internal versus construct validity

Construct validity differs from another type of validity that is critical for experimental manipulations—internal validity. Internal validity reflects the extent to which the intended aspects of the manifest experimental manipulation—and not some artifact(s) of the research methodology—exerted a causal effect on an outcome

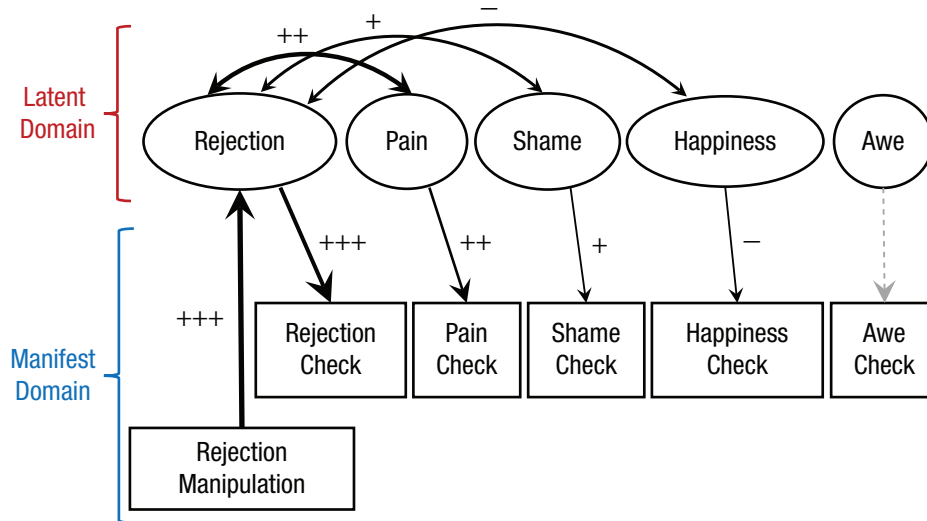


Fig. 2. Schematic depiction of a hypothetical nomological shockwave elicited by a construct-valid social-rejection manipulation. Plus signs depict positive effects; minus signs depict negative effects. Greater numbers of plus signs and thicker arrows depict stronger associations and effects.

(Campbell, 1957; Shadish et al., 2002; Wilson et al., 2010). Threats to internal validity include unintended differences between the participants in the experimental conditions, participant attrition and fatigue over the course of the experiment, environmental and experimenter effects that undermine the manipulation, measures that are not valid or reliable, and participant awareness (of the experiment's hypotheses, of deceptive elements of the study, or that they are being studied; Shadish et al., 2002; Wilson et al., 2010). Each of these issues can elicit spurious effects that are not due to the intended aspects of the experimental manipulation.

Although construct validity requires that the causal chain of events from manipulation to outcome effect was intact (i.e., that the manipulation possessed internal validity), its focus is on the ability of the manipulation to affect the intended constructs in the intended manner (Shadish et al., 2002). In other words, internal validity ensures that the manipulation's effect was causal, whereas construct validity ensures that the manipulation's effect was accurate. Threats to a manipulation's construct validity are *instrumental incidentals*—or confounding aspects of the manipulation that elicited the intended cause in the targeted constructs but were not the aspects of the manipulation that were intended to elicit that effect (Campbell, 1969). For instance, imagine that an experimental condition (e.g., writing an essay that recalls an experience of rejection) was compared with an inappropriate control condition (e.g., writing an essay that tells a story of a brave and adorable otter). This manipulation design would cause an intended increase in rejection, but this effect would be due to

both the intended aspect of the manipulation (i.e., the rejection-related content of the essay) and unintended, confounding aspects as well (e.g., positive attitudes toward brave and adorable otters, ease of writing about a fictional character). Another threat to construct validity is a lack specificity, in which a manipulation exerts a similarly sized impact on a broad array of constructs instead of isolating the target construct (e.g., a rejection manipulation that also increases sadness and anger to the same extent as it does feelings of rejection). An experimental manipulation with construct validity will exert its intended, targeted effects on the intended, specific constructs only through theoretically appropriate aspects of the manipulation (Reichardt, 2006).

Whereas internal validity can be established before testing the construct validity of a manipulation, construct validity first requires that a manipulation exhibit internal validity. Indeed, if an experimental artifact caused by some other aspect of the experiment (e.g., participant selection bias caused by a lack of random assignment) was the actual and unintended source of an observed experimental effect, then it is impossible to claim that the manipulation is what affected the target construct (Cook & Campbell, 1979). This is akin to how psychological questionnaires can have internal consistency among their items without exhibiting construct validity, yet the construct validity of this measure requires the presence of internal consistency. The process through which measures are validated can be instructive for determining how to establish the construct validity of experimental manipulations.

Current Construct Validity Practices for Psychological Manipulations

A survey of the literature on experimental manipulation in social psychology revealed three primary approaches to establishing that a given manipulation has construct validity. These approaches do not map neatly onto the process through which psychological measures are validated, an issue we return to in the Discussion.

Use of previously validated manipulations

The simplest means of establishing the validity of a manipulation is to replicate one that has been already validated in previous research. Many experimental paradigms are frequently reused in other investigations and modified for other purposes. For instance, the seminal article that introduced the Cyberball social-rejection paradigm has been cited more than 1,900 times (Williams, Cheung, & Choi, 2000). However, the value of using previous manipulations is predicated on the extent to which they were adequately validated in such preexisting work. Previously used manipulations, whether they have been validated or not, are often modified before implementation (e.g., the identities of the Cyberball partners are varied; Gonsalkorale & Williams, 2007) or are conceptually replicated by implementing the manipulation through an entirely different paradigm (e.g., being left out of an online chatroom instead of a ball-tossing game; Donate et al., 2017). These conceptual replications are important means for establishing the ability of the manipulated construct's ability to exert its effects irrespective of the manifest characteristics of the manipulation. However, conceptual replication cannot alone establish construct validity.

Pilot validity studies

Whether a manipulation is newly created or acquired from a prior publication, authors often test them before implementation in hypothesis testing. This practice entails conducting at least one separate pilot study of the manipulation outside of the context of the full study procedure (Ellsworth & Gonzalez, 2003). Such pilot studies are used to examine various aspects of the manipulation, from its feasibility to participant comprehension of the instructions to various forms of validity. Of particular interest to the current research, pilot validity studies (a subset of the broader pilot-study category) estimate the manipulation's effect on the target construct (i.e., they test the manipulation's construct validity). In this way, pilot validity studies are a hybrid of experimental pilot studies and the validation studies

used by clinical and personality psychologists who examine the psychometric properties of new measures using the steps we previously outlined.

Pilot validity testing of a new manipulation is an essential step in ensuring that the manipulation has the intended effect on a target manipulation check and to rule out confounding processes (Wilson et al., 2010). Pilot validity testing can also estimate the magnitude and duration of the intended effect. If the effect is so small or transient that it is nearly impossible to detect or if the effect is so strong or long-lasting that it produces ceiling effects or excessive distress among participants, then the manipulation can be altered to address these issues and repiloted. If deception is used, suspicion probes can be included in a pilot study to estimate whether the deception was perceived by the participants (Blackhart, Brown, Clark, Pierce, & Shell, 2012). Even if the manipulation has been acquired from previous work, pilot validity testing is a crucial way of ensuring that the protocol has been accurately re-created and that the validity of the manipulation has been replicated (Ellsworth & Gonzalez, 2003). Because all of these factors have an immense impact on whether a given manipulation will affect its target construct, pilot validity studies are an important means of ensuring the construct validity of a manipulation.

Manipulation checks

A diverse array of measurements fall under the umbrella term *manipulation check*. The overarching theme of such measures is to ensure that a given manipulation had its intended effect (Hauser, Ellsworth, & Gonzalez, 2018). We adopt a more narrow definition to conform to the topic of construct validity; that is, manipulation checks are measures of the construct that the manipulation is intended to affect. This definition excludes attention checks, comprehension checks, and other forms of instructional manipulation checks (Oppenheimer, Meyvis, & Davidenko, 2009), as they do not explicitly quantify the target construct. These instructional manipulation checks are useful tools, especially because they can identify construct-irrelevant variance that is caused by the manipulation. However, our current focus on construct validity entails that we apply the term manipulation check to measures of a manipulation's target construct. We refer to measures of different constructs that are used to ensure that a given manipulation did not exert similarly robust effects onto other, nontarget constructs as *discriminant validity checks*. Discriminant validity checks are specific to each investigation and should include theoretically related constructs to the target construct so that the manipulation's specificity and nomological shockwave can be estimated.

Many articles have debated the utility and validity of manipulation checks, with some scholars arguing for their exclusion (Fayant, Sigall, Lemonnier, Retsin, & Alexopoulos, 2017; Sigall & Mills, 1998). Indeed, manipulation checks can have unintended consequences (e.g., drawing participants' attention to deceptive elements of the experiment, interrupting naturally unfolding psychological processes). Minimally intrusive validation assessments are thus preferable to overt self-report scales (Hauser et al., 2018). Although many such challenges remain with the use of manipulation checks, they are a necessary source of construct validity data that an empirical science cannot forego. Without manipulation checks, the validity of experimental manipulations would be asserted by weaker forms of validity (e.g., face validity) that are deeply flawed when used as the sole basis for construct validity (Grand, Ryan, Schmitt, & Hmurovic, 2010). In an ideal world, such manipulation checks would be validated according to best psychometric practices (see Flake et al., 2017). Without validated manipulation checks, it is uncertain what construct the given check is capturing. An apparently "successful" manipulation check could thus be an artifact of another construct entirely.

The Current Research

The current research was purposed with a central descriptive aim related to construct validation practices for experimental manipulations in social psychology: document the frequency with which manipulations were (a) acquired from previous research or newly created, (b) paired with a pilot validity study, and/or (c) paired with a manipulation check. It was impractical to estimate whether each manipulation that was acquired from previous research was adequately validated by that prior work, so we gave authors the benefit of the doubt and assumed that the research that they cited alongside their manipulations presented sufficient evidence of the manipulation's construct validity. It is likely, given the findings from the current research, that many of these cited articles did not report sufficient evidence for the manipulation's construct validity. Therefore, this is a relatively liberal criterion that probably overestimates the extent to which manipulations have been truly validated.

We focused on social psychology because of its heavy reliance on experimental manipulations, our membership in this field, and this field's ongoing reckoning with replication issues that may result, in part, from experimental practices. We hope that other experimentally focused fields such as cognitive and developmental psychology, economics, management, marketing, and neuroscience may glean insights into their own

manipulation validation practices and standards from this investigation. Further, clinical and counseling psychologists might learn approaches to improving the construct validity of clinical trials, which are similar to experiments in many ways.

In addition to these descriptive analyses, we also empirically examined several important qualities of pilot validity studies and manipulation checks. There is only a sparse amount of research on these topics, so we aimed to fill this gap in our understanding. Given the widespread evidence for publication bias in the field of psychology (Head, Holman, Lanfear, Kahn, & Jennions, 2015), our primary goal in these analyses was to estimate the extent to which pilot and manipulation-check effects are affected by such biases. First, we tested the evidentiary value of these effects via *p*-curve analyses to estimate the extent to which pilot validity studies and manipulation checks capture "true" underlying effects and are not merely the result of publication bias and questionable research practices (Simonsohn, Nelson, & Simmons, 2014). Second, *p*-curve analyses estimated the statistical power of these reported pilot validity and check effects to examine whether long-standing claims that pilot validity studies in social psychology are underpowered (Albers & Lakens, 2018; Kraemer, Mintz, Noda, Tinklenberg, & Yesavage, 2006). Third, we used conventional meta-analyses to estimate the average size and heterogeneity of pilot validity study and manipulation-check effects, useful information for future power analyses. Fourth, these meta-analyses also estimated the presence of publication bias to establish the extent to which pilot validity studies and manipulation checks are selectively reported on the basis of the favorability of their results.

Finally, we returned to our descriptive approach to examine the presence of suspicion probes in the literature. Given the crucial role of suspicion probes in many social-psychological experiments (Blackhart et al., 2012; Nichols & Edlund, 2015), we examined whether manipulations were associated with a suspicion probe and whether "suspicious" participants (i.e., those who had suspicions about the purpose of the study) were retained or excluded from analyses.

Method

Literature search strategy

We conducted our literature search within *JPSP*, a journal that is often reputed to be the flagship journal of experimental social psychology. We limited our search to a single year of publication (as in Flake et al., 2017), selecting the year 2017 because it was recent enough to reflect current practices in the field. Our preregistration plan stated that we would examine Volume 113 of

JPSP, limiting our coding procedures to the two experimentally focused sections: “Attitudes and Social Cognition” and “Interpersonal Relations and Group Processes.” We excluded the “Personality Processes and Individual Differences” section of *JPSP* because of its focus on measurement rather than manipulation. However, we deviated from our preregistration plan by also including Volume 112 in our analysis to increase our sample size and therefore our confidence in our findings.

Inclusion criteria

We sought first to identify every experimental manipulation within the articles that fell within our search. In our initial preregistration plan, we defined experimental manipulations as any systematic alteration of a study’s procedure meant to change a specific psychological construct. However, this definition did not always provide clear guidance in many instances in which a systematically altered aspect of a given study might or might not constitute an experimental manipulation. The ambiguity around many of these early decisions caused us to rapidly deem it impossible to implement this definition in any rigorous or objective manner. Instead, we revised our preregistration plan to follow two simple heuristics. First, we decided that a study aspect would be deemed an experimental manipulation if it had been described by the authors as a manipulation. This approach lifted the burden of determining whether a given aspect of a study was a true manipulation from the coders and instead allowed a given article’s authors, their peer reviewers, and editor to determine whether something could be accurately described as an experimental manipulation. Second, if participants were randomly assigned to different treatments or conditions, this aspect of the study procedure would be considered an experimental manipulation, as random assignment is the core aspect of experimental manipulation (Wilson et al., 2010). We deviated from our preregistration plans by deciding to exclude studies from our analyses that were not presented as part of the main sequence of hypothesis-testing studies in each article (e.g., pilot studies). This deviation was motivated by the realization that pilot validity studies were often provided as the very sources of purported validity evidence we sought to identify for each article’s main experiments and therefore should be examined separately.

Coding strategy

We coded every experimental manipulation for several criteria that either provided descriptive detail or spoke to the evidence put forward for the construct validity of the manipulation.

Coding process. All manipulations were coded independently because both authors possess considerable expertise and training in experimental social psychology, research methodology, and construct validation. We met frequently throughout the coding process to identify coding discrepancies. Such discrepancies were reviewed until we both agreed on one coding outcome (as in Flake et al., 2017). Before such discrepancy reviews and meetings, we each created 459 codes of the nine key coded variables of our meta-analysis (e.g., whether a given study included a manipulation, how many manipulations were included in each study, whether a manipulation was paired with a manipulation check) from the first 11 articles in our literature review. In an exploratory fashion, we examined the interrater agreement in these initial codes (459 codes per rater \times 2 raters = 918 codes; 102 codes per coded variable), which were uncontaminated because we had yet to meet and conduct a discrepancy review. These initial codes exhibited substantial interrater agreement across all coded variables ($\kappa = .89$). Interrater agreement estimates for each of the uncontaminated coded variables are presented below.

Condition number and type. Each manipulation was coded for the number of conditions it contained ($\kappa = .94$) and whether it was administered in a between- or within-participants fashion ($\kappa = .92$). Deviating from our preregistration plan, we also coded whether each of the between-participants manipulations were described as randomly assigning participants to each condition of the manipulation ($\kappa = .63$).

Use in prior research. We coded each manipulation for whether the manipulation was paired with a citation that indicated the manipulation was acquired from previously published research ($\kappa = .84$). If this was not the case, we assumed that the manipulation was uniquely created for the given study. Manipulations that were acquired from prior publications were then coded for whether the authors stated that the manipulations were modified from the referenced version of the manipulation ($\kappa = .75$). It is important to note that we did not code for or select manipulations on the basis of whether that manipulation had been previously validated by the cited work. We refrained from doing so because (a) each cited manipulation could have required a laborious search through a trail of citations to find evidence of validation and (b) simply citing an article in which the manipulation was previously used is likely an implicit argument that the manipulation has been validated by that work.

Pilot validity studies. As a deviation from our preregistration plans, we also coded each manipulation for whether the manipulation’s construct validity was pilot-tested. More

specifically, we coded whether each manipulation was paired with any pilot validity studies that empirically tested the effect of the manipulation on the intended construct (i.e., tested the manipulation's construct validity; $\kappa = .91$).

Manipulation checks. Each manipulation was coded for whether a manipulation check was used ($\kappa = .88$). If such a check were used, we coded the form of the manipulation check (e.g., self-report measure) and whether it was validated in previously published research or was created uniquely for the given study and not validated. We did not rely on authors to make this determination; that is, we did not deem a measure a manipulation check simply because the authors of an article referred to it as such, and we did not exclude a measure from consideration as a manipulation check simply because the authors did not refer to it as such. Instead, we defined a manipulation check as any measure of the construct that the given manipulation was intended to influence (Hauser et al., 2018; Lench, Taylor, & Bench, 2014) and included any measure that met this criterion. This process therefore excluded instructional manipulation checks and other measures that authors deemed manipulation checks but did not actually assess the construct that the manipulation was designed to alter (as in Lench et al., 2014). For each manipulation check we identified, we then coded the form that it took (e.g., self-report questionnaire) and the number of measurements that composed it (e.g., the number of items in the questionnaire).

Suspicion probes. We also coded for whether investigators assessed for participant suspicion of their manipulation ($\kappa = .92$). If such a suspicion probe were used, we coded the form that it took and whether participants who were deemed suspicious were excluded from analyses ($\kappa = .92$).

Results

The "Attitudes and Social Cognition" and "Interpersonal Relations and Group Processes" sections of Volumes 112 and 113 of *JPSP* contained 58 articles. Four of these articles were excluded because they were meta-analyses or nonempirical, leaving 54 articles that summarized 355 independent studies. Of these studies, 244 (68.73%) presented at least one experimental manipulation for a total of 348 experimental manipulations acquired from 49 articles.

Manipulations per study

The majority of studies that contained experimental manipulations reported one (66.80%) or two (25.00%) manipulations, although there was considerable variability in the number of manipulations per study ($M = 1.43$, $SD = 0.68$, mode = 1, range = 1–4).

Conditions per manipulation

The majority of studies reported two (82.18%) or three (12.64%) conditions for each manipulation, although we observed wide variation in the number of conditions per manipulation ($M = 2.30$, $SD = 0.98$, mode = 2, range = 2–13).

Between- versus within-participants designs

The overwhelming majority of manipulations were conducted in a between-participants (94.54%) rather than within-participants (5.46%) manner. Variability in the number of conditions was observed in both within- and between-participants manipulations. These frequencies are depicted in Figure 3, which is an alluvial plot created with SankeyMATIC (<https://github.com/nowthis/sankeymatic>). Alluvial plots visually mimic the flow of rivers into an alluvial fan of smaller tributaries. These figures depict how frequency distributions fall from left to right into a hierarchy of categories. In each plot, a full distribution originates on the left side that then "flows" to the right into different categories whose width is based on the proportion assigned to that initial category. These streams then flow into even more specific subcategories on the basis of their proportions in an additional category.

Manipulation validation practices

Only a modest majority of the manipulations ($n = 202$; 58.04%) were accompanied by at least one of the following sources of purported validity evidence: a citation indicating that the manipulation was used in prior research, a pilot validity study, or a manipulation check (for a breakdown of these statistics, see Table 1 and Fig. 4). Pilot validity study analyses were not preregistered and therefore exploratory.

Citations from previous publications. Of all manipulations, 67 (19.25%) were paired with a citation that indicated the manipulation was used in previously published research. Of these cited manipulations, 16 (23.88%) were described as being modified in some way from their original version. The majority of the remaining 51 cited manipulations were not described in a way in which it was clear whether they had been modified from the original citation. Therefore, the number of modified manipulations provided here may be an underestimation of their presence in the larger body of research.

Manipulation checks. Across all manipulations, 127 (36.49%) were accompanied by a manipulation-check measure. These 127 manipulation checks took the form

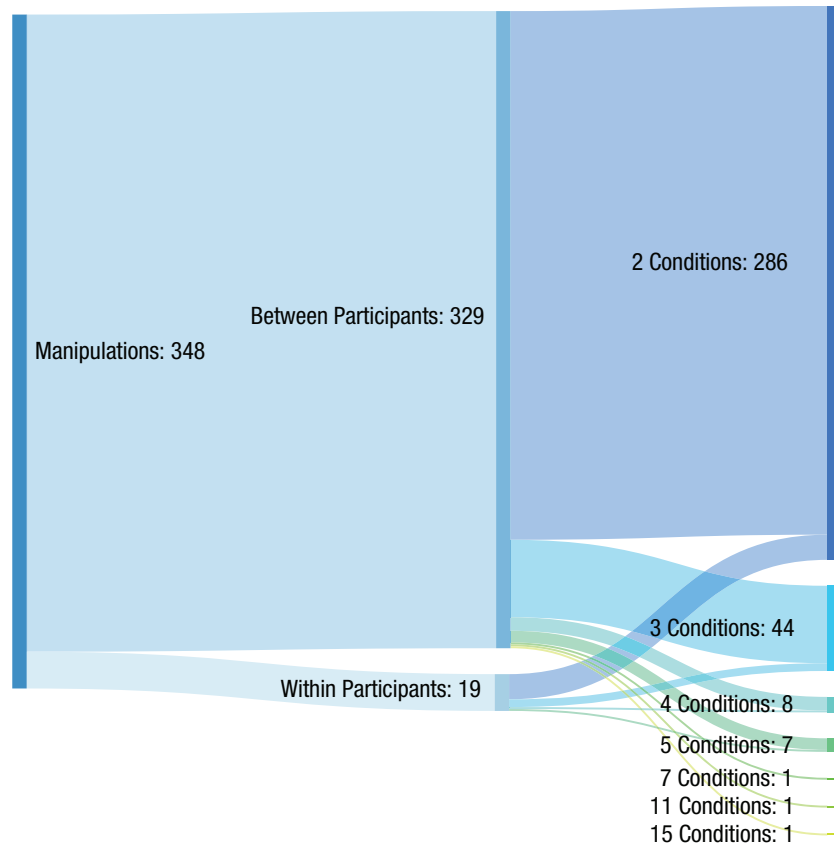


Fig. 3. Alluvial plot of condition frequencies by condition type.

of self-report questionnaires ($n = 105$; 82.68%), coded behavior ($n = 3$; 2.36%), behavioral-task performance ($n = 9$; 7.09%), or an unspecified format ($n = 10$; 7.87%; Fig. 5). Of the 105 self-report manipulation-check questionnaires, 68 (64.76%) consisted of only a single item; the rest included a range of items, $M = 1.68$, $SD = 1.27$, range = 1–10 (Fig. 5).

Suspicion probes

Of all manipulations, only 31 (8.90%) were accompanied by a suspicion probe. Probing procedures were invariably described in vague terms (e.g., “a funnel interview”), and no experimenter scripts or sample

materials were provided that gave any further detail. Of these probed manipulations, only five (16.10%) from two articles reported that they excluded suspicious participants from analyses. The exact criteria for what determined whether a participant was suspicious were not provided in any of these cases, and the impact of excluding these participants was not estimated.

Exploratory analyses

Random assignment. We found that 205 (62.31%) of between-participants manipulations declared that participants were randomly assigned to conditions. No articles described the method they used to randomly assign participants.

Table 1. Frequencies and Percentages of the Number of Manipulations That Were Presented Alongside Each Type of Purported Validity Evidence

| | Without citation | | With citation | |
|---------------|------------------|-------------|---------------|-----------|
| | Not piloted | Piloted | Not Piloted | Piloted |
| Without check | 146 (41.96%) | 35 (10.06%) | 36 (10.34%) | 4 (1.15%) |
| With check | 63 (18.10%) | 37 (10.63%) | 26 (7.47%) | 1 (0.29%) |

Note: The types of validity evidence were a citation indicating that the manipulation had been used in prior research, a pilot validity study, and/or a manipulation check.

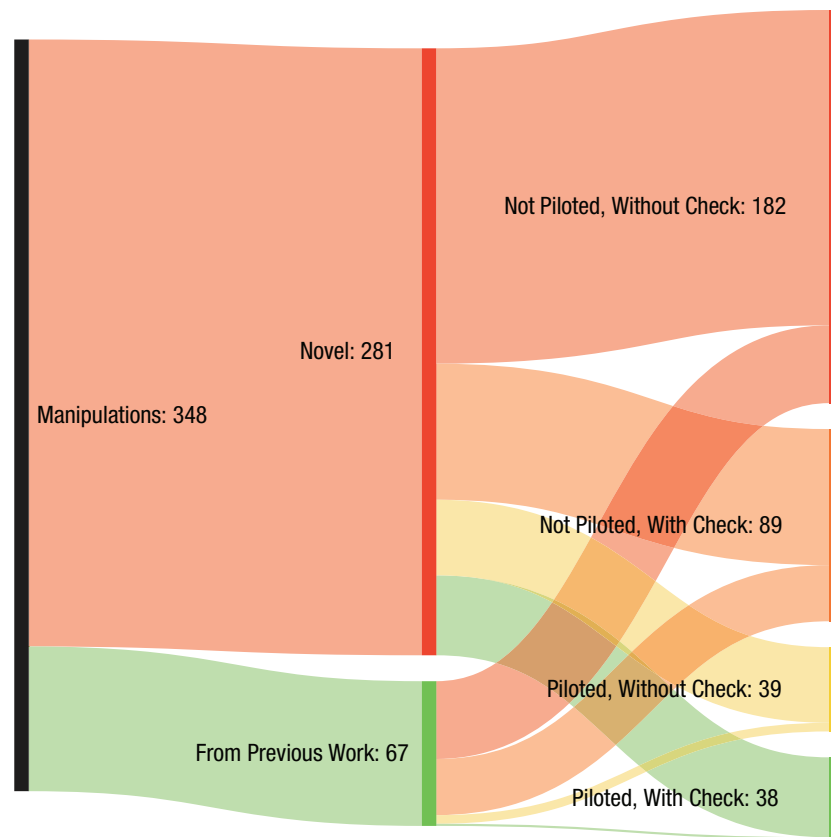


Fig. 4. Alluvial plot depicting distributions of the types of purported validity evidence reported for each manipulation.

Pilot validity study meta-analyses. Pilot validity studies were reported as purported validity evidence for 77 (22.13%) of all manipulations. However, the majority of these studies did not report inferential statistics, described the results too vaguely to identify the target effect, or were drawn from overlapping samples of participants. The results of pilot validity studies were often summarized in a qualitative fashion without accompanying inferential statistics or methodological details (e.g., “pilot testing suggested that the effect . . . tended to be large”; Gill & Cerce, 2017, p. 364). *P*-curve analyses based on the 15 pilot-validity-study effects that we could extract revealed that pilot validity studies exhibited remarkable evidentiary value and were statistically powered at 99% (Fig. 6).

Exploratory random-effects meta-analyses on 14 of the Fisher’s *Z*-transformed pilot validity effects (one effect could not be translated into an effect-size estimate) revealed an overall medium-to-large effect size, $r = .46$, 95% confidence interval (CI) = [.34, .59], $SE = .06$, $Z = 7.28$, $p < .001$, with significant underlying interstudy heterogeneity, $Q(13) = 136.70$, $p < .001$. The average sample size of these studies was 186.47, which explains the high statistical power we observed for such

relatively strong effects. Minimal evidence was found for publication bias in pilot validity studies (see the Supplemental Material available online).

Manipulation-check meta-analyses. Of the 127 manipulations with manipulation checks, six did not report the results of the manipulation check, and 14 others reported incomplete inferential statistics (e.g., a range of *p* values, no test statistics), making it difficult to verify the veracity of the claims. From these manipulation checks, 82 independent manipulation-check effects were extracted and submitted to exploratory *p*-curve analyses, which revealed that manipulation checks exhibited remarkable evidentiary value and were statistically powered at 99% (Fig. 7).

Exploratory random-effects meta-analyses on these Fisher’s *Z*-transformed manipulation-check effects revealed an overall medium-to-large effect size, $r = .55$, 95% CI = [.48, .62], $SE = .03$, $Z = 16.31$, $p < .001$, with significant underlying interstudy heterogeneity, $Q(81) = 2,167.90$, $p < .001$. The average sample size of these studies was 304.79, which explains the high statistical power we observed for such relatively strong effects. No evidence was found for publication bias (see the Supplemental Material).

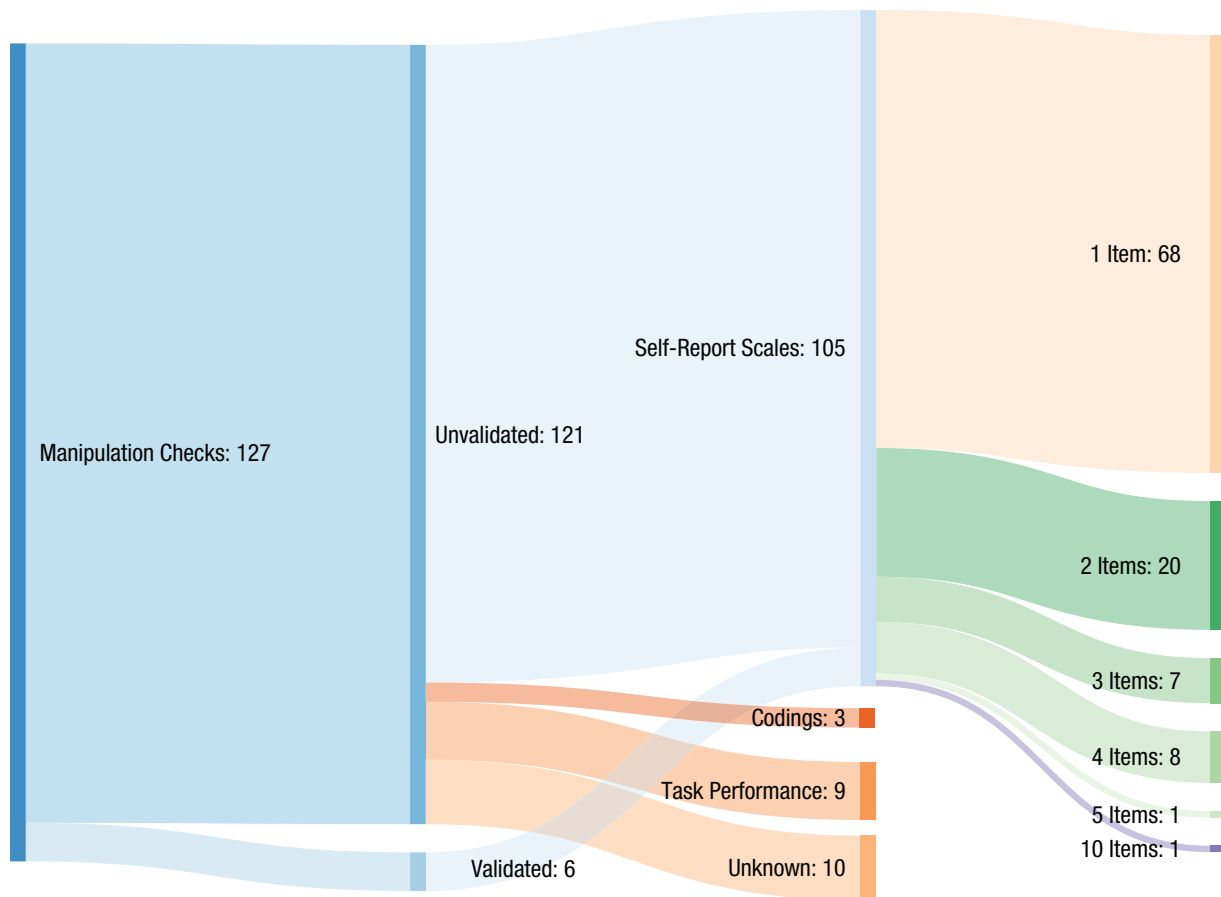


Fig. 5. Alluvial plot depicting distributions of the types of manipulation-check measures reported for each manipulation and numbers of self-report items.

Internal consistency of manipulation checks. Among the 37 manipulation checks that took the form of multiple-item self-report scales, exact Cronbach's alphas were provided for 18 (48.65%) of them, and these estimates mostly exhibited sufficient internal consistency ($M = .83$, $SD = .12$, range = .49–.98).


Validity of manipulation checks. Only eight of all of the manipulation checks (6.30%) were accompanied by a citation indicating that the check was acquired from previous research. After reading the cited validity evidence for each case, only six (4.27%) manipulation checks actually met the criteria for established validation, taking the forms of the Need-Threat Scale (Williams, 2009) and the Positive and Negative Affect Schedule (Watson, Clark, & Tellegen, 1988).

Discussion

Construct-valid measures in psychology can accurately capture the target construct while excluding extraneous variables (Borsboom et al., 2004; Cronbach & Meehl, 1955; Embretson, 1983; Strauss & Smith, 2009). Such construct validity is not limited to psychometrics but

applies equally to experimental manipulations of psychological processes. Indeed, construct-valid manipulations must affect their intended construct in the intended way and not exert their effect via confounding variables (Cook & Campbell, 1979). To better understand the current practices through which experimental social psychologists provide evidence that their manipulations possess construct validity, we examined published articles from *JPSP*.

Chief among our findings was that approximately 42% of experimental manipulations were paired with no evidence beyond the face validity of their underlying construct validity—no citations, no pilot validity testing, and no manipulation checks. Indeed, the most common approach was to present no construct validity evidence whatsoever. To the extent that this estimate generalizes across the field, this suggests that social psychology's experimental foundations rest on considerably unknown ground instead of empirical adamant. In what follows, we highlight other key findings from each domain of our meta-analysis while providing recommendations for future practice in the hope of improving the state of experimental psychological science.



Note: The observed p -curve includes 15 statistically significant ($p < .05$) results, of which 15 are $p < .025$. There were no non-significant results entered.

Fig. 6. Results of the p -curve analysis on pilot-validity-study effects.

Prevalence and complexity of experimental manipulations

At first glance, we find that experimental manipulation is alive and well in social psychology. A little more than two thirds of the studies we reviewed had at least one experimental manipulation. Suggesting a preference for simplicity, more than 90% of studies with manipulations used only one or two manipulations, and a similar number of manipulations contained only two or three conditions. This prevalence of relatively simple experimental designs is promising because exceedingly complex designs (e.g., a $2 \times 3 \times 2$ factorial design) undermine statistical power and inflate Type I and II error rates (R. A. Smith, Levine, Lachlan, & Fediuk, 2002).


Between- versus within-participants designs

More than 90% of manipulations were conducted in a between-participants manner, demonstrating a neglect of within-participants experimental designs. Within-participants designs can better maximize statistical power compared with between-participants designs (Aberson, 2019). The overreliance we observed on between-participants designs may thus undermine the overall power of the findings from experimental social

psychology. However, many manipulations may simply be impossible to present in a repeated measures fashion without undermining the internal validity thereof.

Random assignment and the lack of detail in descriptions of manipulations

Of the between-participants manipulations, a considerable number (approximately two fifths) failed to mention whether participants were randomly assigned to their experimental conditions. Given that random assignment is a necessary condition for a true experimental manipulation (Cook & Campbell, 1979; Wilson et al., 2010), explicit statements of what assignment procedure was used to place participants in their given condition should be included in every report of experimental results. Furthermore, none of the manipulations that mentioned random assignment to a condition described precisely what procedure was used to randomize the assignment process. Without this information, it is impossible to know whether the condition assignment was truly randomized or the randomization procedure could have introduced a systematic bias of some kind. Relatedly, we did not determine whether or how within-participants manipulations randomized the order of the conditions across participants. Future



Note: The observed p -curve includes 96 statistically significant ($p < .05$) results, of which 91 are $p < .025$. There were 2 additional results entered but excluded from p -curve because they were $p > .05$.

Fig. 7. Results of the p -curve analysis of manipulation-check effects.

research would benefit from examining the prevalence of these practices and their impact on the construct validity of within-participants manipulations.

This lack of information about random assignment reflected a much more general lack of basic information that authors provided about their manipulations. It was often the case that manuscripts did not even mention the validity information we sought. Pilot validity studies and manipulation checks were frequently described in a cursory fashion and failed to provide the necessary methodological detail and inferential statistics. More transparency is needed to evaluate each manipulation's validity and for researchers to replicate the procedure in their own labs. Toward this end, we have created a checklist of information that we hope peer reviewers will apply to new research to ensure that each manipulation, manipulation check, and pilot validity study is described in sufficient detail (see the appendix). We further encourage experimenters to use this checklist to adequately detail these important aspects of their experimental methodology.

Previously used versus on-the-fly manipulations

Approximately 80% of manipulations were not acquired from previous research and were instead created ad

hoc for a given study. This suggests that researchers rely heavily on “on-the fly” manipulations (term adapted from Flake et al., 2017), in which ad hoc manipulations are routinely created from scratch to fit the parameters of a given study. The prevalence of this on-the-fly manipulation is almost twice that of on-the-fly measurements in social and personality psychology (~46%; Flake et al., 2017). This prevalence rate may be inflated by a tendency for authors to simply fail to provide such citations for manipulations that have, in fact, been implemented in prior publications. We encourage experimenters to cite publications that empirically examine the validity of their manipulations whenever they exist. These ad hoc procedures appear to acutely afflict experimental designs, and future work is needed to determine the reasons underlying this disproportionate practice.

The field's reliance on creating manipulations de novo is concerning. This practice means that much time and resources are spent on creating new manipulations instead of implementing and improving on existing, validated manipulations. This tendency toward on-the-fly manipulation may reflect psychological science's bias toward novelty and away from replicating past research (Neuliep & Crandall, 1993), which has known adverse consequences (Open Science Collaboration, 2015). We therefore recommend that experimenters

avoid on-the-fly manipulation and instead use existing, previously validated manipulations whenever possible (Recommendation 1), although we concede that not many such manipulations are likely available.

Of the relatively small number of manipulations that were acquired from previous research, roughly one fourth were modified from their original form. This is likely an underestimation of modification rates, as none of the articles we coded explicitly stated that their manipulation was not modified in any way. Modification rates may thus be considerably higher. This practice can have consequences, as modifying a manipulation undermines the established validity of that manipulation, just as modifying a questionnaire often requires it to be revalidated (Flake et al., 2017). This practice of unvalidated modification compounds these issues when the original manipulation that has been modified was never validated itself. We therefore recommend that experimenters avoid modifying previously validated manipulations whenever possible (Recommendation 2A). When modification is unavoidable, we recommend that investigators revalidate the modified manipulation before implementation (Recommendation 2B).

We realize that Recommendations 1 and 2 are likely to be difficult to adhere to given the pessimistic nature of our findings. Indeed, it is difficult to avoid on-the-fly manipulation development and modification when there are no validated versions of a given manipulation already in existence. However, we are optimistic that if experimenters begin to improve their validation practices this will not be an issue for long. These recommendations are given with that bright future in mind.

Pilot validity testing

Approximately one in five manipulations were associated with a pilot validity study before implementation in hypothesis testing. This low adoption rate of pilot validity studies suggests that the practice of pilot validity testing is somewhat rare, which is problematic because such testing is a critical means of establishing the construct validity of a manipulation (Ellsworth & Gonzalez, 2003; Wilson et al., 2010). Pilot validity testing has several advantages over simply including manipulation checks during hypothesis testing. First, pilot validity testing prevents unwanted effects of a manipulation check from intruding on other aspects of the study (Hauser et al., 2018). Second, pilot validity studies allow for changes to be made to the manipulation to optimize its effects before it is implemented. Pilot validity testing would further ensure that time and resources are not wasted on testing hypotheses with manipulations of unknown construct validity. We therefore recommend that experimenters conduct well-powered pilot validity studies

for each manipulation before implementation in hypothesis testing (Recommendation 3A).

These relatively rare reports of pilot validity studies may have been artificially suppressed by the practice of not publishing pilot validity evidence (Westlund & Stuart, 2017). However, all pilot validity evidence should be published alongside the later studies it was used to develop to transparently communicate the evidence for and against the validity of the given manipulation (Asendorpf et al., 2013). Keeping pilot validity studies behind a veil may also reflect a broader culture that undervalues this crucial phase of the manipulation validation process. Pilot validity studies should not be viewed as mere “dress rehearsals” for the main event (i.e., hypothesis testing) but should be granted the same importance, resources, and time as the studies in which they are subsequently used. Robust training, investment, and transparency in pilot validity testing will produce more valid manipulations and therefore more valid experimental findings. We therefore recommend that the results of pilot validity studies should be published as validation articles (Recommendation 3B) and that these validation articles should be accompanied by detailed protocols and stimuli needed to replicate the manipulation (Recommendation 3C).

On an optimistic note, meta-analyses revealed that pilot validity studies exhibited substantial evidentiary value and a robust meta-analytic effect size. These findings imply that researchers are conducting pilot validity tests that capture real and meaningful effects and are not just capitalizing on sources of flexibility or variability. Little evidence of *p*-hacking (Simonsohn et al., 2014) or publication bias was observed, suggesting that researchers are neither simply selectively reporting their pilot validity data to artificially evince an underlying effect nor merely submitting unsuccessful pilot validity studies to the “file drawer” and cherry-picking those that obtain effects. These meta-analyses also revealed that these studies were statistically powered to a maximal degree, thus refuting characterizations of pilot validity studies as underpowered (Albers & Lakens, 2018; Kraemer et al., 2006).

Manipulation checks

Approximately one third of manipulations were paired with a manipulation-check measure. This estimate is much lower than those from other meta-analyses. Hauser and colleagues (2018) reported that 63% of articles in the “Attitudes and Social Cognition” section of the 2016 *JPSP* included at least one manipulation check. Sigall and Mills (1998) reported that 68% of *JPSP* articles in 1998 reported an experimental manipulation. The differences in our estimates likely resulted from

our focus on the manipulation level rather than the article level. We focused on the former because articles present multiple studies with multiple manipulations, and article-level analyses obscure these statistics. We also applied a strict definition of a manipulation check, whereas the authors of these other investigations may have counted any measure that the authors referred to as a manipulation check. It is also possible that manipulation-check prevalence rates have actually decreased in recent years as a result of published critiques of manipulation checks (e.g., Fayant et al., 2017; Sigall & Mills, 1998).

A central issue with manipulation checks is that they intrude on the experiment, calling participants' attention and suspicion to the manipulation and subsequently to the construct under study (Hauser et al., 2018). For instance, asking participants how rejected they felt may raise suspicions about the ball-tossing task from which they were just excluded. Such effects can be manifold and insidious, causing participants to guess at the experimenters' hypotheses, heighten their suspicion, change their thoughts or feelings by reflecting on them, or change the nature of the manipulation itself (Hauser et al., 2018). However, the concerns raised by these critiques are obviated if the manipulation check is administered during the pilot validation of the manipulation and excluded during implementation of the manipulation in hypothesis testing. We therefore recommend that experimenters administer manipulation checks during the pilot validity testing of each manipulation (Recommendation 4A), and postpilot manipulation checks should be administered only if they do not negatively affect other aspects of the study (Recommendation 4B).

Pilot validity studies may differ substantially from the primary experiments that use the manipulations that they seek to validate. Indeed, the presence of other manipulations, measures, and environmental factors might lead a manipulation that exhibited evidence of possessing construct validity to no longer exert its "established" effect on the target construct. When such differences occur between pilot validity studies and focal experiments, including a manipulation check in the focal experiment could establish whether these changes have affected the manipulation's construct validity. If there are legitimate concerns that including a manipulation check could negatively affect the validity of the manipulation, then experimenters could randomly assign participants to either receive the check or not to estimate the effect that the check has on the manipulation's hypothesized effects (assuming sufficient power to detect such effects).

As with the manipulations themselves, the overwhelming majority of manipulation checks were created

ad hoc for the given manipulation. The purported validity evidence provided for the manipulation checks was often simple face validity and, in some cases, a Cronbach's α . Many were single-item self-report measures. These forms of purported validity evidence are insufficient for establishing the construct validity of a measure (Flake et al., 2017). Not knowing whether the check captured the latent construct of interest, or instead tapped into some other construct(s), renders any inferences drawn on such measures theoretically compromised. We therefore recommend that experimenters validate the instruments they use as manipulation checks before use in pilot validity testing (Recommendation 4C). Requiring that manipulation checks be validated would entail a large-scale shift in the practices of experimental social psychologists, who would now often find themselves having to preempt new experiments with the task of creating and validating a new state measure. This would require a new emphasis on training in psychometrics, resources devoted to the manipulation-check validation process, and rewards given to those who do so.

Meta-analyses revealed that manipulation checks exhibited evidentiary value and a robust meta-analytic effect size. Although these findings are promising indicators that the manipulations used in these studies exerted true effects that these checks were able to capture, they cannot speak to the underlying construct validity of these manipulation effects. Indeed, just because manipulations exert some effect on their manipulation checks does not tell us whether the intended aspect of the manipulation exerted the observed effect or whether the manipulation checks measured the target construct. Manipulation-check effects were also maximally statistically powered, which implies that manipulations are at least well powered enough to influence their intended constructs. As with pilot validity studies, there was no evidence for publication bias.

Suspicion probes

Only approximately one tenth of manipulations assessed the extent to which participants were suspicious of the deceptive elements of the study. Although studies vary in the extent to which they are deceptive, almost all experimental manipulations entail some degree of deception in that participants are being influenced without their explicit awareness of the full nature and intent of the manipulation. The majority of studies thus could not estimate the extent to which participants detected their manipulation procedures. Even fewer studies adequately described how suspicion was assessed, often referring vaguely to an experimenter interview or an open-ended survey question. No specific criteria were given for what delineated suspicious from nonsuspicious participants,

and only five studies excluded participants from the former group. Given that no well-validated, standardized suspicion-assessment procedures exist and that there is little in the way of data on what effect that removing suspicious participants from analyses might have on subsequent results (Blackhart et al., 2012), we do not make any recommendations in this domain. Much work is needed to establish the best practices of suspicion assessment and analysis.

Size and duration of manipulation effects

Although many articles established the size of a manipulation's effect on the manipulation check, no manipulation checks repeatedly assessed any manipulation's effect to estimate the time course of these effects. The effect of a given experimental manipulation wanes over time (e.g., Zadro, Boland, & Richardson, 2006), and its time course is a critical element to determine for several reasons. First, experimenters need to know whether the manipulation's effect is still psychologically active at the time point at which they administer their outcome measures and its strength at that given time point. This would allow experimenters to identify an experimental "sweet spot" when the manipulation's effect is strongest. Second, for ethical reasons it is crucial to ensure that the manipulation's effect has adequately decayed by the time the study has ended and participants have returned to the real world. This is especially important when the manipulated process is distressing or interferes with daily functioning (Miketta & Friese, 2019). We therefore recommend that experimenters estimate the time course of their manipulation's effect whenever possible by repeatedly administering manipulation checks during pilot validity testing (Recommendation 5).

Estimating the nomological shockwave via discriminant validity checks

Across the manipulations we surveyed, construct validity was most often assessed (when it was assessed) by estimating the manipulation's effect on the construct that the manipulation was primarily intended to affect. However, a requisite of construct validity is discriminant validity, such that the given manipulation influences the target construct and not a different, confounding construct (Cronbach & Meehl, 1955). Absent this practice, "successful" manipulation checks may obscure the possibility that although the manipulation influences the desired construct, it also affects a related, nontargeted variable to a confounding degree. In this context, discriminant validity can be established by examining the manipulation's nomological shockwave (i.e., the

manipulation's effect on other constructs that exist within the target construct's nomological network). This can be done by administering discriminant validity checks, which are measures of constructs within the target construct's nomological network. In its simplest form, the nomological shockwave can be empirically established by demonstrating that the manipulation's largest effect is on the target construct and then exerts progressively weaker and nonoverlapping effects on theoretically related constructs as a function of their proximity to the target construct in the nomological network. We therefore recommend that experimenters administer measures of theoretically related constructs in pilot testing (i.e., discriminant validity checks; Recommendation 6A) and that these measures are used to estimate the nomological shockwave of the manipulation (Recommendation 6B).

Estimating the nomological shockwave by simply comparing effect sizes and their confidence intervals is admittedly a crude empirical approach. The shockwave rests inherently on the assumption that the manipulation exerts a causal effect on the target construct; this target construct then exerts a causal effect on the discriminant validity constructs by virtue of their latent associations. Causal models could ideally test this sequence of effects, although such quantitative approaches are often limited in their abilities to do so (Fiedler, Schott, & Meiser, 2011). Future research is needed to understand the accuracy and utility of using causal modeling to estimate nomological shockwaves.

Limitations and future directions

This project examined only articles from the *JSPS* and did not include a wider array of publication outlets in social psychology. It may be that our assessment of validation practices would change if we had cast a wider meta-analytic net. Future work should test whether our findings are replicated by studies reported in other journals and in other subfields of psychology. Other experimentally focused fields such as cognitive, developmental, and biological psychology may also vary in their approaches to the validation of their experimental manipulations. Future research is needed in these areas to see whether this is the case. We also used subjective codes and definitions of the manipulation features that we coded, allowing for our own biases to have influenced our findings. We have made all of our codes publicly available so that interested parties might review them for such biases and modify the codes according to their own sensibilities and examine their effect on our results. Indeed, although we do not see our findings as conclusive, the coded data set we have created will be a resource for other investigators to examine in the future.

Conclusion

Experimental manipulations are the methodological foundation of much of social psychology. Our meta-analytic review suggests that the construct validity of such manipulations rests on practices that could be improved. We have made recommendations for how to make such changes that largely revolve around translating the validation approach taken toward personality questionnaires to experimental manipulations. This new model would entail that validated manipulations are used whenever available and that when new manipulations are created they are validated (i.e., pilot validated) before implementation in hypothesis testing. Validity would then be established by demonstrating that the manipulation has its strongest effect on the target construct and theoretically appropriate effects on the nomological network surrounding it. Adopting this model would mean a dramatic change in practices for most laboratories in experimental social psychology. The costs inherent in doing so should be counteracted by a rise in replicability and veridicality of the field's findings. We hope that our assessment of the field's practices is an important initial step in that direction.

Appendix: Peer Reviewer Manipulation Information Checklist

Below are pieces of information that should be included for research using experimental manipulations in psychology. If any of this information is missing, peer reviewers should consider requesting that the authors ensure that it is explicitly stated in the article.

- The number of manipulations in each study
- The number of conditions in each manipulation
- The definition of the construct that each manipulation was intended to affect
- Whether each manipulation was administered between or within participants
- Whether random assignment (for between-participants designs) or counterbalancing (for within-participants designs) were used in each manipulation
- How random assignment or counterbalancing was conducted in each manipulation
- Whether each manipulation was acquired from previous research or newly created for the study
- The preexisting validity evidence for each manipulation that was acquired from previous research
- Whether each manipulation that was acquired from previous research was modified from the version of the manipulation detailed in the previous research
- The validity evidence for each manipulation that was modified from previous research
- Whether each manipulation was pilot-tested before implementation
- The validity evidence for each measure used in each pilot study
- The pilot validity evidence for each manipulation that was pilot-tested
- The detailed methods and results of each pilot study
- Whether each manipulation was paired with a manipulation check that quantified the manipulation's target construct
- The validity evidence for each manipulation check
- Whether each manipulation was paired with a discriminant validity check that quantified potentially confounding constructs
- The validity evidence for each discriminant validity check
- Whether deception by omission was used for each manipulation (i.e., facts about the manipulation were withheld from participants)
- Whether deception by commission was used for each manipulation (i.e., untrue information about the manipulation was provided to participants)
- Whether each deceptive manipulation was paired with a suspicion probe
- The methodological details of each suspicion probe
- The validity evidence for each suspicion probe
- How each suspicion probe was scored
- How participants were deemed to be suspicious or not for each suspicion probe
- How suspicious participants were handled (e.g., excluded from analysis, suspicion used as a covariate) in each manipulation study

Transparency

Action Editor: Richard Lucas

Editor: Laura A. King

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding

This work was supported by National Institute on Alcohol Abuse and Alcoholism Grant K01-AA026647 (to D. S. Chester).

ORCID iD

David S. Chester  <https://orcid.org/0000-0002-9383-7138>

Acknowledgments

This project was intended to capture an exploratory snapshot of the literature; therefore, no hypotheses were advanced a priori. The preregistration and amended plan for the current research have been made publicly available via <https://osf.io/rbtwj> and <https://osf.io/zvg3a>, respectively. The disclosure table of all included studies and their associated codes have been made publicly available via <https://osf.io/je9xu/files>.

Supplemental Material

Additional supporting information can be found at <http://journals.sagepub.com/doi/suppl/10.1177/1745691620950684>

References

- Aberson, C. L. (2019). *Applied power analysis for the behavioral sciences*. New York, NY: Routledge.
- Albers, C., & Lakens, D. (2018). When power analyses based on pilot data are biased: Inaccurate effect size estimators and follow-up bias. *Journal of Experimental Social Psychology, 74*, 187–195.
- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. Washington, DC: American Educational Research Association.
- Asendorpf, J. B., Conner, M., De Fruyt, F., De Houwer, J., Denissen, J. J., Fiedler, K., . . . Perugini, M. (2013). Recommendations for increasing replicability in psychology. *European Journal of Personality, 27*, 108–119.
- Benson, J. (1998). Developing a strong program of construct validation: A test anxiety example. *Educational Measurement: Issues and Practice, 17*(1), 10–17.
- Berkowitz, L., & Donnerstein, E. (1982). External validity is more than skin deep: Some answers to criticisms of laboratory experiments. *American Psychologist, 37*, 245–257.
- Blackhart, G. C., Brown, K. E., Clark, T., Pierce, D. L., & Shell, K. (2012). Assessing the adequacy of postexperimental inquiries in deception research and the factors that promote participant honesty. *Behavior Research Methods, 44*, 24–40.
- Borsboom, D., Mellenbergh, G. J., & van Heerden, J. (2004). The concept of validity. *Psychological Review, 111*, 1061–1071.
- Brewer, M. B. (2000). Research design and issues of validity. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research: Methods in social and personality psychology* (pp. 3–39). New York, NY: Cambridge University Press.
- Campbell, D. T. (1957). Factors relevant to the validity of experiments in social settings. *Psychological Bulletin, 54*, 297–312.
- Campbell, D. T. (1969). Prospective: Artifact and control. In R. Rosenthal & R. L. Rosnow (Eds.), *Artifact in behavioral research* (pp. 351–382). San Diego, CA: Academic Press.
- Chester, D. S., DeWall, C. N., & Pond, R. S. (2016). The push of social pain: Does rejection's sting motivate subsequent social reconnection? *Cognitive, Affective, & Behavioral Neuroscience, 16*, 541–550.
- Clark, L. A., & Watson, D. (2019). Constructing validity: New developments in creating objective measuring instruments. *Psychological Assessment, 31*, 1412–1427.
- Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design & analysis issues for field settings*. Chicago, IL: Rand McNally.
- Cronbach, L., & Meehl, P. (1955). Construct validity in psychological tests. *Psychological Bulletin, 52*, 281–302.
- Donate, A. P. G., Marques, L. M., Lapenta, O. M., Asthana, M. K., Amodio, D., & Boggio, P. S. (2017). Ostracism via virtual chat room: Effects on basic needs, anger and pain. *PLOS ONE, 12*(9), Article e0184215. doi:10.1371/journal.pone.0184215
- Ellsworth, P. C., & Gonzalez, R. (2003). Questions and comparisons: Methods of research in social psychology. In M. Hogg & J. Cooper (Eds.), *The SAGE handbook of social psychology* (pp. 24–42). Thousand Oaks, CA: SAGE.
- Embretson, S. (1983). Construct validity: Construct representation versus nomothetic span. *Psychological Bulletin, 93*, 179–197.
- Fayant, M. P., Sigall, H., Lemonnier, A., Retsin, E., & Alexopoulos, T. (2017). On the limitations of manipulation checks: An obstacle toward cumulative science. *Psychology, 30*, 125–130.
- Fiedler, K., Schott, M., & Meiser, T. (2011). What mediation analysis can (not) do. *Journal of Experimental Social Psychology, 47*, 1231–1236.
- 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.
- Garner, W., Hake, H., & Eriksen, C. (1956). Operationism and the concept of perception. *Psychological Review, 63*, 149–159.
- Gill, M., & Cerce, S. (2017). He never willed to have the will he has: Historicist narratives, “civilized” blame, and the need to distinguish two notions of free will. *Journal of Personality and Social Psychology, 112*, 361–382.
- Gonsalkorale, K., & Williams, K. D. (2007). The KKK won't let me play: Ostracism even by a despised outgroup hurts. *European Journal of Social Psychology, 37*, 1176–1186.
- Grand, J. A., Ryan, A. M., Schmitt, N., & Hmurovic, J. (2010). How far does stereotype threat reach? The potential detriment of face validity in cognitive ability testing. *Human Performance, 24*, 1–28.
- Gray, K. (2017). How to map theory: Reliable methods are fruitless without rigorous theory. *Perspectives on Psychological Science, 12*, 731–741.
- Haslam, S. A., & McGarty, C. (2004). Experimental design and causality in social psychological research. In C. Sanson, C. C. Morf, & A. T. Panter (Eds.), *Handbook of methods in social psychology* (pp. 235–264). Thousand Oaks, CA: SAGE.
- Hauser, D. J., Ellsworth, P. C., & Gonzalez, R. (2018). Are manipulation checks necessary? *Frontiers in Psychology, 9*, Article 998. doi:10.3389/fpsyg.2018.00998
- Haynes, S. N., Richard, D. C., & Kubany, E. S. (1995). Content validity in psychological assessment: A functional

- approach to concepts and methods. *Psychological Assessment*, 7, 238–247.
- Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., & Jennions, M. D. (2015). The extent and consequences of p-hacking in science. *PLOS Biology*, 13(3), Article e1002106. doi:10.1371/journal.pbio.1002106
- Highhouse, S. (2009). Designing experiments that generalize. *Organizational Research Methods*, 12, 554–566.
- Kraemer, H. C., Mintz, J., Noda, A., Tinklenberg, J., & Yesavage, J. A. (2006). Caution regarding the use of pilot studies to guide power calculations for study proposals. *Archives of General Psychiatry*, 63, 484–489.
- Lench, H. C., Taylor, A. B., & Bench, S. W. (2014). An alternative approach to analysis of mental states in experimental social cognition research. *Behavior Research Methods*, 46, 215–228.
- Lewin, K. (1939). Field theory and experiment in social psychology: Concepts and methods. *American Journal of Sociology*, 44, 868–896.
- Loevinger, J. (1957). Objective tests as instruments of psychological theory. *Psychological Reports*, 3, 635–694.
- Messick, S. (1995). Validity of psychological assessment: Validation of inferences from persons' responses and performances as scientific inquiry into score meaning. *American Psychologist*, 50, 741–749.
- Miketta, S., & Friese, M. (2019). Debriefed but still troubled? About the (in) effectiveness of post experimental debriefings after ego threat. *Journal of Personality and Social Psychology*, 117, 282–309.
- Mook, D. (1983). In defense of external invalidity. *American Psychologist*, 38, 379–387.
- Neuliep, J. W., & Crandall, R. (1993). Reviewer bias against replication research. *Journal of Social Behavior and Personality*, 8(6), 21–29.
- Nichols, A. L., & Edlund, J. E. (2015). Practicing what we preach (and sometimes study): Methodological issues in experimental laboratory research. *Review of General Psychology*, 19, 191–202.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349, 253–267.
- 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.
- Reichardt, C. S. (2006). The principle of parallelism in the design of studies to estimate treatment effects. *Psychological Methods*, 11, 1–18.
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston, MA: Houghton Mifflin.
- 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.
- Simms, L. J. (2008). Classical and modern methods of psychological scale construction. *Social and Personality Psychology Compass*, 2, 414–433.
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143, 534–547.
- Smith, G. T. (2005). On construct validity: Issues of method and measurement. *Psychological Assessment*, 17, 396–408.
- Smith, R. A., Levine, T. R., Lachlan, K. A., & Fediuk, T. A. (2002). The high cost of complexity in experimental design and data analysis: Type I and type II error rates in multiway ANOVA. *Human Communication Research*, 28, 515–530.
- Strauss, M. E., & Smith, G. T. (2009). Construct validity: Advances in theory and methodology. *Annual Review of Clinical Psychology*, 5, 1–25.
- Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54, 1063–1070.
- 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.
- Williams, K., Cheung, C., & Choi, W. (2000). Cyberostracism: Effects of being ignored over the Internet. *Journal of Personality and Social Psychology*, 79, 748–762.
- Williams, K. D. (2009). Ostracism: Effects of being ignored and excluded. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 41, pp. 279–314). San Diego, CA: Academic Press.
- Wilson, T. D., 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, pp. 51–81). Hoboken, NJ: Wiley.
- Zadro, L., Boland, C., & Richardson, R. (2006). How long does it last? The persistence of the effects of ostracism in the socially anxious. *Journal of Experimental Social Psychology*, 42, 692–697.
- Zumbo, B. D., & Chan, E. K. H. (2014). Setting the stage for validity and validation in social, behavioral, and health sciences: Trends in validation practices. In B. D. Zumbo & E. K. H. Chan (Eds.), *Validity and validation in social, behavioral, and health sciences* (pp. 3–8). Cham, Switzerland: Springer.