Making inferential leaps: manipulation checks and the road towards strong inference

Stefan. L.K. Gruijters

Open University of the Netherlands

** Note that this is a pre-print and not the Version of Record. A definitive version of this manuscript has been published Open Access in the Journal of Experimental Social Psychology. Please access (for free!) the published version of this manuscript before citing at:

<https://doi.org/10.1016/j.jesp.2021.104251>**

The paper can be cited as:

“Gruijters, S.L.K. (2021). Making inferential leaps: manipulation checks and the road towards strong inference. *Journal of Experimental Social Psychology*. In Press.
<https://doi.org/10.1016/j.jesp.2021.104251>”

Abstract

Experiments in psychology often target hypothetical constructs to test some causal hypothesis or theory. In light of this goal, it is pertinent to use a manipulation that actually changes the focal hypothetical construct, and only that construct. In assessing whether such manipulation “success” can be assumed, researchers often include manipulation validity checks in their designs—a measure of the focal construct which should be responsive to the manipulation. One interpretation of a positive manipulation check is that it lends credence to a particular causal interpretation of a phenomenon. Scrutinizing the results of such manipulation checks supposedly enables a more stringent test of a causal hypothesis. This paper submits that manipulation checks do not improve our inferences to causal explanations, but may in practice result in weaker hypothesis tests. Rather than being useful, manipulation checks are at best uninformative, but more likely compromise the appraisal of a causal hypothesis. The second half of this paper advocates four methodological heuristics, offered as alternatives to manipulation validity checks, to more severely test causal hypotheses. The heuristics call for a burgeoning focus on (a) assessing the specificity of manipulations, (b) evaluating theoretical risk, (c) attempts to cast doubt on alternatives, and (d) appraising the relative merits of explanations. I conclude that rather than relying on manipulation checks as a ‘Band-Aid’ method to alleviate validity concerns, inferential rigor can be improved by virtue of these heuristics.

Keywords: manipulation check, validity, causal explanation, inference to the best explanation, falsificationism, hypothetico-deductive method.

“I have devised seven separate explanations, each of which would cover the facts as far as we know them. But which of these is correct can only be determined by the fresh information which we shall no doubt find waiting for us.” (Sherlock Holmes—Conan Doyle, 1892/1992, p. 258).

1. Introduction

Experiments are ubiquitous in psychology and useful for making inferences about causes of phenomena. Causal inference is in turn important, because linking causes to consequences provides an important inroad towards scientific explanation. Within experimental (social) psychology, manipulations commonly target inferred (mental) states and processes—a consequence of the cognitive revolution which steered psychology away from mere behavioral analysis (e.g., Miller, 2003). One practical consequence of such progress is that it has not made the job easier for psychological experimenters. In particular, when dealing with unobserved variables, a critical question to tackle is whether a manipulation influenced the construct of interest. Manipulation checks are argued to be a helpful (even necessary) analytical tool for experimental psychologists dealing with this issue (see Fiedler et al., 2021; Simons & Holcombe, 2014; Stroebe & Strack, 2014; Wilson et al., 2010).

It is worth noting at the outset that the concept “manipulation check” is variously and somewhat inconsistently used in the literature (see also Ejelöv & Luke, 2020). “Manipulation check” is used to denote a check on whether participants noticed a manipulation (e.g., Sigall & Mills, 1998), understood the manipulation (Aronow et al., 2019), or read and attended the experimenters’ instructions (Hauser & Schwarz, 2015; Kung et al., 2018; Oppenheimer et al., 2009). Additionally, ‘manipulation check’ is used to describe checks on whether the construct of interest was successfully manipulated, by including a measurement of the focal construct and assessing its responsiveness to manipulation (Chester & Lasko, 2020; Ejelöv & Luke, 2020; Fayant et al., 2017; Sigall & Mills, 1998). The above listed usages, though all assessing boundary

conditions for successful manipulation, have different intents and purposes. In this paper, only the measurement interpretation of ‘manipulation check’ is of interest—the Litmus test for such checks being the measure’s responsiveness to a manipulation.

Data suggests that a majority of social psychology researchers consider manipulation checks a hallmark of properly designed experiments (Fayant et al., 2017; see also Hauser et al., 2018; Sigall & Mills, 1998). As a general principle, it is of course worthwhile to obtain some level of confidence that a manipulation manipulates what it is supposed to manipulate. However, in this paper I argue that including a manipulation check in an experiment does not pave the way towards strong inferences about causal hypotheses. Even though manipulation check are part of the methodological staple in experimental psychology, this paper specifically critiques how manipulation checks are used to substantiate causal inferences (see also Fayant et al., 2017; Sigall & Mills, 1998; Trafimow & Rice, 2009).

The present paper is bi-purposed. First, I discuss current practice surrounding manipulation checks, including the purported role they play in providing corroborating evidence for a causal hypothesis. Through a discussion of four scenarios, I posit that manipulation checks do not meaningfully improve the causal interpretation of an observed phenomenon—aligning with other published critiques (Fayant et al., 2017; Sigall & Mills, 1998; Trafimow & Rice, 2009). Secondly, having discussed their limited inferential force, I discuss alternative tactics to appraise the merits of a causal hypothesis. More specifically, through a discussion of four methodological heuristics, this work aims to contribute in moving the field towards stronger inferences on causal hypotheses.

2. Manipulations, measures, and hypothetical constructs

This section provides a brief discussion of key terms and concepts forming the conceptual backbone of the ensuing discussions—as some conceptual clarifications are

foundational to the discussion that follows. As an illustrative backdrop (and selected for no other reason), this paper relies on research connecting mood states to selective attentional processes as a running example (e.g., Becker & Leininger, 2011; Richards et al., 1992; Tamir & Robinson, 2007). For instance, Tamir and Robinson (2007) report several experiments testing *mood-congruent attentional biases* (e.g., a positive mood shifts attention towards rewarding stimuli, and the obverse). To induce particular mood states, participants were required to recall ‘sad’, ‘neutral’, and ‘positive’ events, and write about these events (see Tamir & Robinson, 2007, Study 2). A dot-probe task (a reaction-time task designed to measure attentional bias) was used as a measure of mood congruent attention shifts. Other details, when needed for illustration, are discussed in the relevant sections.

2.1. Distinguishing between operational and causal hypotheses

In experimental (social) psychology, causal hypotheses frequently contain reference to *hypothetical constructs*—such as “mood” and “attention”. Hypothetical constructs tend to designate unobserved cognitive states, processes, or dispositions, which may provide explanations for an observable cognitive or behavioral phenomenon (Lovasz & Slaney, 2013; MacCorquodale & Meehl, 1948). As Cronbach & Meehl (1955) phrased it in the context of measurement, a “construct is some postulated attribute of people, assumed to be reflected in test performance” (p. 284)¹. It is critical to conceptually tease apart the hypothetical constructs from the instruments that are used to measure (e.g., a reaction-times task) and manipulate the constructs (e.g., a mood induction). Notably, a manipulation such as “having subjects recall a negative event” is not synonymous with, or merely an operationalization of the hypothetical construct ‘mood’, but a way to causally influence the construct. Manipulations are thus more properly seen as instruments (or vehicles) to bring about change in a hypothetical construct (see Eronen, 2020; Woodward, 2003). Conversely, measurement instruments, such as the dot-

probe task, can be seen as designed to *pick up* variations in hypothetical constructs (Borsboom et al., 2004).

The traditional approach to experimentation in social psychology follows a hypothetico-deductive recipe (see Haig, 2005, 2009, 2014). Researchers using this method take a theory or causal hypothesis to propose a tentative explanation for a target phenomenon. Once a causal hypothesis is formulated (the hypothetico part), a way to manipulate presumed causal factors (often hypothetical constructs) can be designed, from which a specific prediction can be deduced (the deductive part). The deduced prediction is subsequently rephrased in operational terms and usually put to a statistical test. Importantly, causal hypotheses are only indirectly assessed through an inferential leap from the operational hypothesis to the causal hypothesis (e.g., Mayo, 2018; Meehl, 1990a). The distinction is captured by the following two claims (see also Figure 1):

- (I) Recalling negative versus positive events is statistically related to performance on the dot-probe task (operational hypothesis)
- (II) Mood is a causal factor in explaining selective attention processes (causal hypothesis)

Maintaining a sharp distinction between type I and II claims is critical to fend off an “operationalist mode” of inference (cf. Borsboom, 2006)—which would involve an illicit inferential leap in which a statistical test of an operational hypothesis is equated to a test of a substantive (causal) hypothesis (see also Mayo, 2018; Meehl, 1967, 1990a). Examining the type I statement roughly corresponds to what Shadish, Cook & Campbell (2002) describe as *causal description*: identifying the effect of varying levels of a manipulation on an outcome². The veracity of causal description hinges mainly on the internal validity of the experiment (e.g., Campbell, 1957), which requires that a manipulation, exclusively, caused variation on the focal outcome variable (e.g., Shadish et al., 2002). Internal validity is sensitive to issues such as the

randomization procedure (e.g., Gruijters, 2016), selection bias, non-random drop-out (e.g., Groenwold et al., 2012), demand characteristics, and so forth (see Whitley Jr & Kite, 2012).

Claim 2 moves (a step) towards a *causal explanation* (Shadish et al., 2002; Woodward, 2003)³.

To corroborate a causal explanation, the operational hypothesis (type I claim) needs to connect with the causal hypothesis (type II claim).

Figure 1. The distinction between the operational and causal hypotheses (causal description versus explanation). Note that the figures use structural equation modelling convention to denote observed (rectangles) and unobserved variables (ovals). Single-headed arrows denote causal direction, whereas double-headed arrows suggest a statistical association.

2.2. Connecting type I and II claims with validity assumptions

A leap from causal description towards explanation requires (at least) two validity assumptions. One assumption pertains to the manipulation, the other to the measure used to assess the outcome (See Figure 2). The validity of measures and manipulations is often depicted as a question of whether a manipulation or measure adequately *represents* or *instantiates* a hypothetical construct (e.g., Whitley Jr & Kite, 2012; Wilson et al., 2010). But such definitions are unclear about the relationship between hypothetical constructs and their

measures and manipulations. For instance, Whitley Jr and Kite (2012) describe the issue of validity as: “Whenever a manipulation is intended to operationally define a hypothetical construct, one must be sure that the manipulation accurately represents the construct.” (p. 247). Wilson and colleagues (2010) describe the validity of manipulations as concerning whether a manipulation “is a good instantiation of the conceptual variable” (p. 56). However, validity can also be conceived of as a simple causal notion (Borsboom et al., 2004, 2009; see also Grujters & Fleuren, 2018; Hood, 2009). In this view, a measure is valid when it captures variations in—is causally sensitive to—the processes/states alluded to by a hypothetical construct (Borsboom et al. 2004, 2009). Similarly, the causal approach to validity readily extends to manipulations: a manipulation can be considered valid when it causes variation in the focal hypothetical construct. In terms of the running example, valid manipulation and measurement requires that the mood induction procedure influences the hypothetical construct “mood”, and that the dot-probe task is causally sensitive to variation in “attentional processes.”

The chain of causation involved in an experiment using hypothetical constructs (see Figure 2), such as between mood and attention, thus begins from a manipulation → hypothetical construct (antecedent) → hypothetical construct (consequent) → outcome measure. The first hypothetical construct being the causal antecedent, modified by manipulation, whereas the latter construct is hypothesized to change as a result—a change which subsequently should be reflected on the outcome measure. The causal connection between the two hypothetical constructs (type II claim) is inferred (rather than observed) from the evidence for the operational hypothesis (type I claim). In summary, the move from causal description to explanation thus minimally requires a) a valid manipulation, b) a valid outcome measure, and c) internal validity.

Thinking of validity as an additional research question about causation (between manipulation and the hypothetical construct), invites additional quantitative questions about manipulations. Relevant considerations (see also Chester & Lasko, 2020; Ejelöv & Luke, 2020) are the effect size of a given manipulation on the hypothetical construct, but also the nature of the relationship between the manipulation and the hypothetical construct (i.e., dose-response relationships). Notably, an ideal manipulation is not only causally related to the hypothetical construct, but it also generates *sufficient* between-group variance on the construct, so that a statistical basis is put in place enabling exploration of its causal role. These are all questions about manipulations that as of yet receive little attention in the literature (see Chester & Lasko, 2020; Ejelöv & Luke, 2020).

Figure 2. Connecting a causal description to a causal explanation depends on validity assumptions. Note that the depiction loosely draws from Sigall & Mills (1998).

2.3. The role of in-experiment manipulation checks for causal inference

Manipulation checks are typically included in experimental studies to get insight into—or “check”—the validity of manipulations (Ejelöv & Luke, 2020; Sigall & Mills, 1998; Whitley Jr & Kite, 2012; Wilson et al., 2010). Manipulation checks are a way to find evidence

that a manipulation causes a change on a measure of the hypothetical construct and usually involve a single non-validated ‘on the fly’ measurement (Flake et al., 2017) of the hypothetical construct (Chester & Lasko, 2020). Researchers are, of course, not precluded from using various measures to assess validity—but this is not a particularly common practice (see reviews by Chester & Lasko, 2020; Ejelöv & Luke, 2020).

The current standard method involves analyzing a causal fork (see Figure 3), which usually involves the following protocol. First an assessment of the operational hypothesis is conducted by testing for a manipulation effect on the outcome variable. The second step involves an additional test of the manipulation check, by inspecting whether a measure of the focal hypothetical construct was also responsive to the manipulation. Data suggests that the majority of social psychology researchers use significance tests (with H_0 rejected when the inferred effect size > 0) to analyze manipulation check results (Ejelöv & Luke, 2020). Significance tests, of course (see Cohen, 1994), merely allow researchers to reach a categorical verdict (as opposed to quantitative description) as to whether a focal construct was likely influenced by the manipulation. The rationale behind the causal fork tactic is that, if one finds evidence that a manipulation affects the outcome variable (Figure 3, observation *a*), and also that a manipulation affects the manipulation check (Figure 3, observation *b*), the combined observations (*a+b*) provide cumulative evidence for the causal hypothesis.

Figure 3. Diagram depicting the role of a manipulation check. A common method is to statistically assess both (a) and (b) to provide convergent evidence for a causal explanation.

Academic legend has it that by analyzing a manipulation check—a measure of the hypothetical construct—one is therefore justified in moving from an operational hypothesis to a claim about the causal hypothesis. But what role do manipulation checks have in shaping inferences about causal hypotheses? There are four scenarios which may occur when studying a manipulation’s effect on the outcome measure and an additional manipulation check (see also Fayant et al., 2017; Sigall & Mills, 1998).

Scenario (1): A confirmed operational hypothesis combined with an effect on the manipulation check.

Scenario (2): A confirmed operational hypothesis combined with no effect on the manipulation check.

Scenario (3): A disconfirmed operational hypothesis combined with no effect on the manipulation check.

Scenario (4): A disconfirmed operational hypothesis combined with an effect on the manipulation check.

To evaluate the contribution of the manipulation check results on causal inference, it is helpful to consider a counterfactual: Would the causal inference in this scenario have been different had the manipulation check not been considered? The running example of associations between mood states and selective attention will be used once again for illustrative purposes. Notably, in mood research, a common procedure to check whether a mood induction succeeds in affecting mood is by examining self-reports on how ‘happy’, ‘sad’, and so forth, participants feel after exposure to a mood induction (e.g., Tamir & Robinson, 2007).

3. Possible scenarios when analyzing manipulation checks in experiments.

3.1. Scenario 1

The first scenario occurs when researchers find supporting evidence for an operational hypothesis, aligning with the causal hypothesis underlying the prediction. This finding is conjoined with a positive manipulation check result, such that participants who received a positive mood induction indicate significantly higher ratings on reports of “happiness”, “joy”, and so forth. So can one conclude based on the manipulation check results that the causal hypothesis is strongly corroborated? Summarily, the critical reason why a positive manipulation check provides limited information (beyond the test of the focal prediction) is because there is no reason to suppose a manipulation has not impacted other hypothetical constructs as well (see also Chester & Lasko, 2020; Fayant et al., 2017; Sigall & Mills, 1998; Trafimow & Rice, 2009). So much the worse, given what is known about psychological measures and manipulations, it is for two related reasons likely that multiple explanations are always in play.

First, manipulations in psychology are often “fat-handed” (Eronen, 2020)—manipulations are unlikely to impact only a single construct, but are more akin to “buckshots” impacting several constructs (see also Chester & Lasko, 2020). Secondly, psychological variables are often either directly or indirectly related to each other, that is, there are vast networks of causal connections between various psychological attributes (the Crud-factor; Meehl, 1990a, 1990b; Orben & Lakens, 2020)⁴. Manipulations are thus likely to cause what Chester & Lasko (2020) label a “nomological shockwave”, and thus cannot be assumed to only influence the focal construct. Overall, this state of affairs permits the possibility that a manipulation influenced the outcome through some other indirect pathway, rather than through the hypothesized focal construct.

In the example of the mood-attention hypothetical, one conceivable alternative may be that a broader cognitive phenomenon (e.g., conceptual priming) rather than mood primarily drove the effect (see Figure 4). Although I do not claim this to necessarily constitute a strong alternative, it illustrates that most operational hypotheses can be connected to a plurality of causal hypotheses. There are various conceivable causal pathways that may explain the findings—including the manipulation affecting an alternative process (e.g., conceptual priming), which as a side effect leads to differences on self-reported mood. Or perhaps both processes were affected by the manipulation, implying that the mood hypothesis only explains part of the findings. Contrasting conditions on a measure of the presumed causal construct does not make it more or less likely that other unmeasured constructs were manipulated, which is why such a finding adds little to the interpretation of the experiment.

Figure 4. Alternative explanations for a positive “mood” manipulation check.

3.2. Scenario 2

In this scenario, the manipulation is found to influence the outcome, but no difference was found on the manipulation check. So something has influenced selective attention, but have we lost confidence that a mood was responsible? Note that the absence of an effect on the manipulation check does not necessarily falsify the hypothesis that mood was manipulated. One reason for this is simply that measurement error (e.g., Loken & Gelman, 2017; Schmidt & Hunter, 1999) or test invalidity (e.g., Hussey & Hughes, 2020) could have made the manipulation check insensitive to the manipulation. Another possibility is that the size of the manipulation effect, what some label *cause size* (Ejelöv & Luke, 2020), is not large enough to cause a detectable group-level change on the manipulation check. It may very well be that the mood induction effect is detected on the response-times used by the dot-probe task, but that the self-report of mood may not be as sensitive to the manipulation. Given that manipulations checks are often implemented using verbal measures—as opposed to behavioral measures

more commonly used for primary outcomes (Hauser et al., 2018)—it is not in the least inconceivable that the lower sensitivity of such measures explains the incongruent findings. Without an alternative explanation of what was manipulated, and until such alternatives are examined further, the findings corroborate the causal hypothesis.

3.3. Scenario 3

In this scenario, we did not find support for a mood induction effect on selective attention, and this finding is accompanied by a null-effect on the manipulation check. There are several explanations for the findings: one or more measures used in the study may be invalid, the manipulation may not have affected the focal construct, or the causal hypothesis is just false. Alternatively, without further examination of the manipulation, it may also be that the manipulation was not causally potent enough to engender a substantive change on the focal construct—which may explain the absence of any effect on outcomes measures. Unfortunately, it is not possible to determine which interpretation provides the best explanation without running more experiments. The only inference that can plausibly be made in light of such findings is that the causal hypothesis has not been corroborated. Similar to Scenario 1 and 2, the results of the manipulation check has not substantially affected or altered this inference.

3.4. Scenario 4

The final scenario involves a situation where a researcher does not find support for a core prediction, but does find an effect on the manipulation check. This scenario does not substantially change things for our priming researcher compared to Scenario 3. If we fail to provide evidence for an effect of the mood induction on a measure of selective attention, then we fail to provide evidence for the causal hypothesis. Sigall & Mills (1998), however, consider

manipulation checks to be useful (though not necessary) to bolster the interpretation of null-findings. In particular, the idea is that when results are accompanied by evidence that a manipulation did influence a measure of the focal hypothetical construct, one may therefore be *more* justified in inferring the causal hypothesis as “falsified”. A Bayes-minded researcher would say that one has “stacked up” more evidence against the causal hypothesis.

A special interest in the current scenario also comes from authors discussing the value of “failed” (direct) replication studies (e.g., Ejelöv & Luke, 2020; Fiedler et al., 2021). For example, suppose that researchers attempt to replicate the mood research in a different cultural context, using the same procedures (so an exact replication) and run into scenario 4—they fail to replicate the operational hypothesis, but observe an effect on the manipulation check. With such a scenario in mind, Fiedler and colleagues recently wrote:

“It is so essential to distinguish between replication failures reflecting invalidity of a hypothesis $\Delta x \rightarrow \Delta y$ and the banal case that an ineffective manipulation failed to establish the premise, Δx ” (Fiedler et al., 2021, p. 8).

In other words, failing to replicate evidence for a causal hypothesis, that changing construct X results in changes on outcome Y, is much more interesting when X has actually been changed. At first blush, there seems little to object to the logic behind such arguments. However, the reasoning assumes that *any* change of X should be sufficient to test the causal hypothesis. More concretely, there is little quantitative reasoning involved when one considers any (perhaps statistically significant?) effect on a manipulation check as grounds for pinning the blame on the falsity of the causal hypothesis. Such a strong interpretation of the Scenario 4 results would be to rush to conclusions, because the possibility that the manipulation failed to induce a sufficiently sized change on X has not been ruled out (see also Ejelöv & Luke, 2020).

Rather than discarding the causal hypothesis, one could be inclined to first examine whether the treatment *insufficiently* changed the focal construct. Indeed, this illustrates that exact replication is arguably not the best tool for appraising causal hypotheses, but is rather useful to build up the empirical evidence-base for a specific operational hypothesis—a particular *finding* (see also Crandall & Sherman, 2016; Earp & Trafimow, 2015).

In short, even though a positive manipulation check appears to strengthen the falsification of the causal hypothesis, it could nonetheless be premature to discard the causal hypothesis without first attempting to modify the ‘dose’ delivered by the manipulation. To provide proper footing for a leap from the operational hypothesis to an inference about the falsity of a causal hypothesis, one needs information beyond manipulation check results.

4. Towards stronger forms of hypothesis testing

Essentially, what the scenarios demonstrated is that—with or without manipulation checks—the traditional hypothetico-deductive method is limited in providing a robust test of a causal hypothesis (Haig, 2005, 2009; Rozeboom, 1997). As previously described, the hypothetico-deductive method consists of using theory or a causal hypothesis to deduce a prediction. What follows is an attempt to provide empirical evidence for that prediction. If supporting evidence for the prediction is found, this is considered to confirm the theory or causal hypothesis from which the prediction was deduced. However, whether a prediction pertaining to some operationalized hypothesis panned out is not sufficient to infer a particular causal hypothesis (e.g., Haig, 2005). To justify such an inferential leap, various validity assumptions need to be made, and alternative explanations need to be excluded. Neither of these issues are firmly resolved by manipulation checks. So how can we move towards stronger forms of causal inference in (basic) experimental psychology?

Methodological heuristics to generate stronger causal inferences can chiefly be distilled from Popper's (1959, 1972) methodology of science. Popperian-inspired heuristics to generate stronger tests include taking *risk* with predictions and *ruling-out* ways in which an incorrect causal inference could have been reached (Mayo, 2018; Meehl, 1967, 1978, 1990a, 1990b). Additionally, a more comprehensive appraisal of hypotheses or theory, by focusing on explanatory virtues, provides another useful methodological approach (Haig, 2005, 2009; Thagard, 1978, 1989). Borrowing from the above-cited works, this final section reviews several methodological heuristics which may serve as alternatives to manipulation checks. The four heuristics can be considered when evaluating an inference to a causal explanation.

4.1. Assessing the specificity of manipulations

Valid measurement and manipulation is a prerequisite for leaping from causal description to explanation, but this does not imply that checks for validity should be included in research testing a substantive theory or hypothesis. There are various procedural limitations and problems with using in-experiment manipulations checks that give reason for avoiding this approach (e.g., Bless & Burger, 2016; Fayant et al., 2017; Hauser et al., 2018; Hauser & Schwarz, 2015; Sigall & Mills, 1998). Issues of validity, be it of a measure or manipulation, are not simple matters to be checked "on the side" or "on the fly" (Chester & Lasko, 2020; Flake et al., 2017). Rather, which, and how, hypothetical constructs relate to a particular manipulation (or measure) involves appraising a theory of how a manipulation (or measurement) is thought to operate (see Smith, 2005; Strauss & Smith, 2009). As a consequence, such independent, auxiliary theories are deserving of equal scientific scrutiny to the primary research question. Indeed, recommended best practice for the validation of measures and manipulations is to

examine such questions in suitably representative, independent studies (see Chester & Lasko, 2020; Ejelöv & Luke, 2020; Flake et al., 2017).

Chester & Lasko (2021) recently reviewed the current state of validation practices and found that most manipulations are either unvalidated or are merely assessed by convergent measures. Given the previous discussion, on the problem of fat-handed manipulations (Eronen, 2020), it is not inconceivable that (in the typical psychological study) one or more non-focal constructs are manipulated. This is problematic however, as having a “thin-fingered” manipulation is critical for making a claim about the causal relevance of a specific variable (see Eronen, 2020; Woodward, 2003). For this reason, it is recommended that manipulations are extensively evaluated on *specificity*, before being sent off to the frontlines to tackle some substantive research question. It is particularly important that researchers and reviewers inquire which constructs close in the nomological net to the focal construct may be affected by the manipulation as well (see also Chester & Lasko, 2020)—especially when one could reasonably expect these constructs to relate to the outcome variable as well. Identification of such constructs is optimally accompanied by attempts to *empirically* assess the impact of a manipulation on those constructs. Notably, assessing the specificity of manipulations thus minimally requires examination of *both* convergent and discriminant measures (see also Chester & Lasko, 2020; Fiedler et al., 2021; Flake et al., 2017). Such examinations ideally provide a *quantitative* analysis of the relationship between manipulation and focal constructs (e.g., using effect size indices), providing valuable information over and above significance tests (Ejelöv & Luke, 2020).

4.2. Assessing theoretical risk

Once the specificity of a manipulation has been sufficiently established, it is worth considering the connection between a prediction (i.e., the results of the operational hypothesis test) and the causal hypothesis. In various contributions, Meehl (1978, 1990a, 1990b) elaborated on Popper's (1959, 1972) notion of bold, *risky predictions*. Concisely, a sufficiently risky prediction implies that—absent the focal theory or causal hypothesis—one has no reason to expect to find statistical evidence for the operational hypothesis (Meehl, 1990a). If without the focal theory or causal hypothesis at hand (e.g., by merely invoking common sense, background knowledge or other theories), one could predict the same effect, this would imply the presence of low theoretical risk. In terms of the running example, a sufficiently risky prediction requires that 'having people recall negative versus positive events' can *only* be expected to change attention processes *given that* the mood theory is correct. By contrast, if the same prediction pertaining to the effect of the mood induction could reasonably be made based on some other considerations (e.g., a priming phenomenon) this suggests the presence of low-risk.

As another example, suppose one uses cognitive dissonance theory (Festinger, 1957) to deduce the prediction that persons driving electric cars have an above average concern about climate change. This is not a particularly risky prediction, as one could have concocted the prediction without ever having heard of cognitive dissonance theory. The implication of low-risk predictions is that examining such predictions cannot provide a stringent test of a theory or causal hypothesis. No matter how robust the evidence in support of a prediction (e.g., that climate change concerns relate to the electric car usage) turns out to be, there is little confirmation or disconfirmation it can bestow on cognitive dissonance theory. For this reason, observing a relationship between electric car use and climate change concerns only provides

a *weak test* of cognitive dissonance theory (see Meehl, 1967, 1978, 1990a). The concrete consequence of such a circumstance is that additional tactics are needed to further appraise the theory or causal hypothesis.

One way to increase the riskiness of a prediction, according to Meehl (1978, 1990a), is to deduce more *precise* predictions from a theory or hypothesis (e.g., a prediction on the point effect size, or a pattern of results). But such precise predictions, absent formalized theory, are not easy to come by in psychology (Broers, 2021; Gruijters & Peters, 2020). Indeed, most predictions in experimental psychology are not of the high-risk sort, because they merely predict the presence or absence of a phenomenon (e.g., Meehl, 1967, 1978, 1990a). Though theoretical risk is perhaps not easily quantitatively evaluated, a qualitative evaluation can still be useful. It would perhaps be opportune to include a theoretical risk assessment in discussion sections (instead of emphatic denials that correlation implies causation), to aid readers' assessment whether a theory or causal hypothesis has been subjected to a stringent test. A risk assessment could accordingly discuss whether a particular finding, an effect of a particular manipulation on an outcome measure, could have been predicted without the specific hypothesis the researchers happened to be interested in. Such an evaluation specifically requires inquiring whether a predicted effect (e.g., an association between recalling negative events and the dot-probe task scores) is *only* to be expected in light of the focal hypothesis (e.g., the mood theory), or whether there are alternative hypotheses from which one could have deduced the same prediction. In the latter case, the heuristic discussed in the following section might be considered as a guide to follow-up research.

4.3. Ruling-out rather than Ruling-In.

The likely common circumstance of low theoretical risk warrants actively examining potential alternative explanations for a corroborated operational hypothesis. Within the broader hypothetico-deductive method, an approach that is well-known and stronger is to attempt to cast doubt on it (Mayo, 2018; Popper, 1959). Such a falsificationist approach (Popper, 1959) rests on the realization that accumulating evidence in support of a particular explanation cannot lead to any certainty about its validity. However, the converse tactic, gathering evidence against a particular explanation of a relationship can help to rule it out as a plausible candidate. To use this tactic, it is pertinent to theorize beyond confirming predictions (e.g., if X is causally relevant to behavior Y , then we expect to observe P), to falsifying predictions (e.g., if X is causally relevant to behavior Y , then it should not be the case that we observe R). Though offering an viable alternative to work that merely aims to strengthen the proposed explanation with supporting evidence, the approach is seldom explicitly used to gauge causal hypotheses (see also Earp & Trafimow, 2015; Meehl, 1978, 1990a; Trafimow, 2009).

As an example of this tactic, consider again the work by Tamir & Robinson (2009), whose Study 5 actively strove to cast doubt on an alternative explanation for the observed mood-induction effect. In particular, the authors considered the possibility that the manipulation caused a semantic priming effect rather than a mood change. To cast doubt on the alternative, the researchers used music as a mood-induction, which generated a similar effect to the verbal recall induction method (see also Siemer, 2005). This study exemplifies a more stringent test of the mood hypothesis, not because it aligned with the results of the previous studies, but because it exposed the mood hypothesis to risk of an incongruent

finding. Of course, there may be other viable alternatives that would require ruling-out, but the study illustrates the general tactic. In sum, the ruling-out heuristic (see Mayo, 2018; Meehl, 1978, 1990a)—exploring the ways a causal hypothesis may be incorrectly inferred—puts a causal hypothesis to a more stringent test compared to a manipulation check.

4.4. Contrasting explanations using explanatory virtues

The not unlikely result of the previous three steps is the identification of multiple plausible explanations for an observed effect. Optimally such competing explanations are pitted against each other in an experiment designed to decide between them (e.g., Platt, 1964), but such crucial empirical tests are unfortunately not always feasible. The hypothetico-deductive method has been critiqued in a number of published works for not offering a stringent framework to appraise the relative merits of theories and hypotheses (Haig, 2005, 2009, 2014; Platt, 1964; Rozeboom, 1997; Thagard, 1978). Some argue that a more comprehensive appraisal of the merits of an hypothesis or theory is offered by an approach called ‘inference to the best explanation’ (Haig, 2009; Harman, 1965; Thagard, 1978), which has for instance an important role in Haig’s (2005b) abductive theory of method.

Pivotal to inference to the best explanation is to not only appraise an hypothesis on its predictive success, but also on whether the hypothesis provides an adequate explanation for a phenomenon. It has been argued that considering theoretical virtues—or an hypothesis’ ‘goodness of fit to the empirical evidence’ (Haig, 2009, *p.* 231)—may direct researchers towards the best available explanation (see Haig, 2005, 2009; Thagard, 1978, 1989). Unfortunately, a lacking consensus prohibits a definitive and exhaustive list of virtues to be reviewed here (see Haig, 2009), so a selective listing of criteria is unavoidable. Popper (1972) mentions, in addition to his well-known criteria of falsifiability, *relative satisfactoriness*—which holds that one should

prefer the theory ‘which tells us more’ than its alternative (p. 217). By this Popper (1972) specifically meant that preference should go towards a theory ‘which contains the greater amount of information; which is logically stronger; which has the greater explanatory *and* predictive power; and which can therefore be more severely tested by comparing predicted facts with observations’ (p. 217, my emphasis). Kuhn (1977) also hailed a similar criterion but labeled it empirical *scope* and further famously listed *simplicity*, *fruitfulness*, *consistency*, *and* *accuracy* as valuable criteria (as cited in Dienes, 2008; Haig, 2009). Thagard (1978) also discusses simplicity and scope but adds *analogy* (the availability of similar explanations in related domains)⁵.

At least two virtues—*scope* and *simplicity*—are thus widely agreed upon and expounded by philosophers of science (e.g., Dienes, 2008; Haig, 2009; Kuhn, 1977; Thagard, 1978, 1989). Thagard (1978) defines empirical *scope* not only as a theory’s ability to make sense of the “facts”, but also its ability to explain phenomena in different domains. The *simplicity* criterion urges that theories making fewer assumptions, mentioning fewer hypothetical constructs, should be preferred over those making more such claims. I propose that applying such criteria may help to more systematically compare the merits of rival explanations, once they emerge in the literature (for extensive discussion of these topics, see Haig, 2005, 2009; Thagard, 1978, 1989). Because of the broad consensus towards the scope and simplicity virtues, the following briefly illustrates the application of only these two virtues.

The running example of mood research led to the identification of various plausible explanations for the mood induction effect: the mood hypothesis, a broader semantic priming phenomenon, or some combination of both factors. Other work (Siemer, 2005; Tamir & Robinson, 2007), using variations of mood induction tasks (e.g., music rather than verbal

recall), added relevant findings for contrasting the hypotheses on the basis of scope and simplicity. The limited findings used here for illustrative purposes suggests that the mood hypothesis has a comparatively larger explanatory breadth than the semantic priming alternative. In particular, the mood hypothesis more easily explains *both* the verbal recall and music mood induction effects on attention—whereas the latter semantic priming hypothesis explains only one of the two findings.

Yet another alternative explanation (see Figure 4) is that both mood and semantic priming effects occurred, and that the manipulation exerted its influence through some indirect causal route. From a simplicity perspective, it would notably require more assumptions to argue that an indirect effect through both mood and semantic priming explains the findings (e.g., that both constructs are theoretically relevant, that they can interact, and so on)—rendering the simpler mood alternative preferable given its sufficient scope. This comparison of two hypotheses' ability to explain mood induction phenomena is, of course, undergirded by only few empirical findings. As a general principle, the broader and more robust the available evidence-base, the better such a comparative appraisal of virtues can be done.

5. Conclusion

5.1. The inferential value of manipulation checks

This paper made a case against the ill-considered use of manipulation checks in experimental designs to aid the interpretation of the findings (see also Fayant et al., 2017; Hauser et al., 2018; Sigall & Mills, 1998; Trafimow & Rice, 2009). To be clear: the message of this paper is not that the validity of manipulations is not an important issue. Rather, it is argued that manipulation checks only lend the appearance of benefitting inferential rigor,

whereas their contribution to the appraisal of causal hypotheses or theory is limited. In other words, manipulation checks do not impose a severe test of an explanation, but—by generating a false sense of certainty—may rather result in a weaker test of a causal hypothesis. In particular, when a manipulation check signals that a measure of a construct may be responsive to manipulation, researchers might be disincentivized to consider alternative explanations. Conversely, in some scenarios a ‘negative’ manipulation check result may lead to premature abandonment of a causal hypothesis—even though the limited responsiveness could be due to problems with the measure, or by insufficient sensitivity of that measure to the manipulation. All in all, manipulation checks only add a superficial zest to hypothetico-deductivism—hence, overinterpreting such measures (in particular, those designed “on the fly”) may actually erect a barrier to theoretical progress in experimental psychology (see also Fayant et al., 2017).

5.2. Recommendations and concluding remarks

In the wake of Popper (e.g., 1959, 1972), repeated calls towards stronger (Platt, 1964), more risky (Meehl, 1978), and more severe (Mayo, 2018) hypothesis tests have appeared in the meta-scientific literature. The concepts “strong”, “risky” and “severe” share a common denominator—they emphasize the importance of putting a causal hypothesis or theory to a stringent appraisal, rather than merely confirming a prediction that some association between variables exists. This paper reviewed four methodological heuristics which may engender a more stringent appraisal of causal hypotheses: (1) assessing the specificity of manipulation, (2) assessing theoretical risk, (3) attempting to cast doubt, and (4) comparing alternative explanations using explanatory virtues.

Left unanswered is the question how to put these heuristics into practice. Although the heuristics are not meant as some canonical, or deontic scheme for appraising causal hypotheses, applying the heuristics more systematically to plan or evaluate research may be of value to theory testers in experimental psychology. There is however a logical order to the sequence in which these heuristics were presented. A natural starting point for hypothesis testers would be to first evaluate the effects of a manipulation and its specificity (1). A theoretical risk assessments (2) may in turn be beneficial to direct follow-up research to concrete alternative explanations—alternatives upon which researchers may want to cast some doubt (3). And lastly, considering theoretical virtues of particular alternatives (4) may help researchers and reviewers to evaluate their relative merits based on the accrued and available evidence-base.

Despite offering a more comprehensive appraisal of causal hypotheses, these heuristics are of course not problem-free. For one, the heuristics strongly hinge on researchers' judgment (rather than mere quantitative assessment) on whether an inferential leap from operational to causal hypothesis is warranted. Although steps towards a more formal evaluation of explanatory 'goodness' have been made (e.g., Thagard, 1989), these are unfortunately underdeveloped as of yet (see Haig, 2009). A renewed interest in formal methods to evaluate the theoretical virtues of rival hypotheses and theory would therefore be a welcome development. Notwithstanding this limitation, a more consistent application of the four methodological heuristics, including consideration of theoretical virtues using a selection of the reviewed criteria, allows for a more elaborate and rigorous appraisal of competing hypotheses (Haig, 2009).

In conclusion, the take-home message of this paper is that the cost of manipulation checks, in the absence of clear inferential benefits, warrants avoidance of such checks in research designed to test some other prediction. Inferential confidence might instead be gleaned from a more comprehensive appraisal of a causal hypothesis or theory—an appraisal which is hopefully tangibly assisted by the four heuristics reviewed in this paper.

Conflicts of Interest: “The author(s) declare that there were no conflicts of interest with respect to the authorship or the publication of this article.

Acknowledgements: The author is thankful to Rachel Brown and Bram Fleuren for critical comments on previous drafts of this paper. Gratitude is further expressed to Marie-Pierre Fayant, David Chester, Roger Giner-Sorolla, and various anonymous reviewers, who provided valuable suggestions for improvement of prior versions.

References

- Aronow, P. M., Baron, J., & Pinson, L. (2019). A Note on Dropping Experimental Subjects who Fail a Manipulation Check. *Political Analysis*, 27(4), 572–589. <https://doi.org/10.1017/pan.2019.5>
- Becker, M. W., & Leinenger, M. (2011). Attentional selection is biased toward mood-congruent stimuli. *Emotion*, 11(5), 1248–1254. <https://doi.org/10.1037/a0023524>
- Bless, H., & Burger, A. M. (2016). A Closer Look at Social Psychologists' Silver Bullet. *Perspectives on Psychological Science*, 11(2), 296–308. <https://doi.org/10.1177/1745691615621278>
- Borsboom, D. (2006). The attack of the psychometricians. *Psychometrika*, 71(3), 425–440. <https://doi.org/10.1007/s11336-006-1447-6>
- Borsboom, D., Cramer, A. O. J., Kievit, R., Zand Scholten, A., & Franic, S. (2009). The end of construct validity. In *The concept of validity* (pp. 135–170). Information Age publishing.
- Borsboom, D., Mellenbergh, G. J., & van Heerden, J. (2004). The Concept of Validity. *Psychological Review*, 111(4), 1061–1071. <https://doi.org/10.1037/0033-295X.111.4.1061>
- Broers, N. J. (2021). When the Numbers Do Not Add Up: The Practical Limits of Stochastologicals for Soft Psychology. *Perspectives on Psychological Science*, 174569162097055. <https://doi.org/10.1177/1745691620970557>
- Campbell, D. T. (1957). Factors relevant to the validity of experiments in social settings. *Psychological Bulletin*, 54(4), 297–312. <https://doi.org/10.1037/h0040950>
- Chester, D., & Lasko, E. (2020). Construct Validation of Experimental Manipulations in Social Psychology: Current Practices and Recommendations for the Future. *Perspectives on Psychological Science*, in press. <https://doi.org/10.31234/osf.io/t7ev9>

- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49(12), 997–1003.
<https://doi.org/10.1037/0003-066X.49.12.997>
- Crandall, C. S., & Sherman, J. W. (2016). On the scientific superiority of conceptual replications for scientific progress. *Journal of Experimental Social Psychology*, 66, 93–99.
<https://doi.org/10.1016/j.jesp.2015.10.002>
- Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52(4), 281–302. <https://doi.org/10.1037/h0040957>
- Dienes, Z. (2008). *Understanding psychology as a science: An introduction to scientific and statistical inference*. Macmillan International Higher Education.
- Earp, B. D., & Trafimow, D. (2015). Replication, falsification, and the crisis of confidence in social psychology. *Frontiers in Psychology*, 6(May), 621.
<https://doi.org/10.3389/fpsyg.2015.00621>
- Ejelöv, E., & Luke, T. J. (2020). “Rarely safe to assume”: Evaluating the use and interpretation of manipulation checks in experimental social psychology. *Journal of Experimental Social Psychology*, 87, 103937. <https://doi.org/10.1016/j.jesp.2019.103937>
- Eronen, M. I. (2020). Causal discovery and the problem of psychological interventions. *New Ideas in Psychology*, 59, 100785. <https://doi.org/10.1016/j.newideapsych.2020.100785>
- 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(1), 125. <https://doi.org/10.5334/irsp.102>
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford university press.
- Fiedler, K., McCaughey, L., & Prager, J. (2021). Quo Vadis, Methodology? The Key Role of Manipulation Checks for Validity Control and Quality of Science. *Perspectives on Psychological Science*, 174569162097060. <https://doi.org/10.1177/1745691620970602>

- Flake, J. K., Pek, J., & Hehman, E. (2017). Construct Validation in Social and Personality Research. *Social Psychological and Personality Science*, 8(4), 370–378. <https://doi.org/10.1177/1948550617693063>
- Groenwold, R. H. H., Donders, A. R. T., Roes, K. C. B., Harrell, F. E., & Moons, K. G. M. (2012). Dealing With Missing Outcome Data in Randomized Trials and Observational Studies. *American Journal of Epidemiology*, 175(3), 210–217. <https://doi.org/10.1093/aje/kwr302>
- Grujters, S. L. K. (2016). Baseline comparisons and covariate fishing: Bad statistical habits we should have broken yesterday. *European Health Psychologist*, 18(5), 205–209.
- Grujters, S. L. K. (2017). The reasoned actions of an espresso machine: a comment on Peters and Crutzen (2017). *Health Psychology Review*, 11(2), 125–129. <https://doi.org/10.1080/17437199.2017.1306716>
- Grujters, S. L. K., & Fleuren, B. P. I. (2018). Measuring the Unmeasurable: The Psychometrics of Life History Strategy. *Human Nature*, 29(1), 33–44. <https://doi.org/10.1007/s12110-017-9307-x>
- Grujters, S. L. K., & Peters, G.-J. Y. (2020). Meaningful change definitions: sample size planning for experimental intervention research. *Psychology & Health*, 1–16. <https://doi.org/10.1080/08870446.2020.1841762>
- Haig, B. D. (2005). An Abductive Theory of Scientific Method. *Psychological Methods*, 10(4), 35–64. <https://doi.org/10.1037/1082-989X.10.4.371>
- Haig, B. D. (2009). Inference to the best explanation: a neglected approach to theory appraisal in psychology. *The American Journal of Psychology*, 122(2), 219–234. <http://www.ncbi.nlm.nih.gov/pubmed/19507428>
- Haig, B. D. (2014). *Investigating the psychological world: Scientific method in the behavioral sciences*. MIT Press.

Harman, G. H. (1965). The Inference to the Best Explanation. *The Philosophical Review*, 74(1), 88.

<https://doi.org/10.2307/2183532>

Hauser, D. J., Ellsworth, P. C., & Gonzalez, R. (2018). Are Manipulation Checks Necessary?

Frontiers in Psychology, 9. <https://doi.org/10.3389/fpsyg.2018.00998>

Hauser, D. J., & Schwarz, N. (2015). It's a Trap! Instructional Manipulation Checks Prompt

Systematic Thinking on "Tricky" Tasks. *SAGE Open*, 5(2).

<https://doi.org/10.1177/2158244015584617>

Hood, S. B. (2009). Validity in Psychological Testing and Scientific Realism. *Theory &*

Psychology. <https://doi.org/10.1177/0959354309336320>

Hussey, I., & Hughes, S. (2020). Hidden Invalidity Among 15 Commonly Used Measures in

Social and Personality Psychology. *Advances in Methods and Practices in Psychological*

Science, 3(2), 166–184. <https://doi.org/10.1177/2515245919882903>

Kuhn, T. S. (1977). *The essential tension: Selected studies in scientific tradition and change*.

University of Chicago Press.

Kung, F. Y. H., Kwok, N., & Brown, D. J. (2018). Are Attention Check Questions a Threat to

Scale Validity? *Applied Psychology*, 67(2), 264–283. <https://doi.org/10.1111/apps.12108>

Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science*,

355(6325), 584–585. <https://doi.org/10.1126/science.aal3618>

Lovasz, N., & Slaney, K. L. (2013). What makes a hypothetical construct "hypothetical"?

Tracing the origins and uses of the 'hypothetical construct' concept in psychological

science. *New Ideas in Psychology*, 31(1), 22–31.

<https://doi.org/10.1016/j.newideapsych.2011.02.005>

MacCorquodale, K., & Meehl, P. E. (1948). On a distinction between hypothetical constructs

and intervening variables. *Psychological Review*, 55(2), 95–107.

<https://doi.org/10.1037/h0056029>

Mayo, D. (2018). *Statistical inference as severe testing: how to get beyond the statistics wars*. Cambridge University Press.

Meehl, P. E. (1967). Theory-Testing in Psychology and Physics: A Methodological Paradox. *Philosophy of Science*, 34(2), 103–115. <https://doi.org/10.1086/288135>

Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*. <https://doi.org/10.1037//0022-006x.46.4.806>

Meehl, P. E. (1990a). Appraising and Amending Theories: The Strategy of Lakatosian Defense and Two Principles That Warrant It. *Psychological Inquiry*. https://doi.org/10.1207/s15327965pli0102_1

Meehl, P. E. (1990b). Why Summaries of Research on Psychological Theories are Often Uninterpretable. *Psychological Reports*, 66(1), 195–244. <https://doi.org/10.2466/pr0.1990.66.1.195>

Michell, J. (1999). *Measurement in psychology*. Cambridge University Press.

Michell, J. (2005). The logic of measurement: A realist overview. *Measurement*, 38(4), 285–294. <https://doi.org/10.1016/j.measurement.2005.09.004>

Miller, G. A. (2003). The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences*, 7(3), 141–144. [https://doi.org/10.1016/S1364-6613\(03\)00029-9](https://doi.org/10.1016/S1364-6613(03)00029-9)

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4), 867–872. <https://doi.org/http://dx.doi.org/10.1016/j.jesp.2009.03.009>

Orben, A., & Lakens, D. (2020). Crud (Re)Defined. *Advances in Methods and Practices in Psychological Science*, 3(2), 238–247. <https://doi.org/10.1177/2515245920917961>

- Peters, G.-J. Y., & Crutzen, R. (2017). Pragmatic Nihilism: How a Theory of Nothing can Help Health Psychology Progress. *Health Psychology Review*, 1–39. <https://doi.org/10.1080/17437199.2017.1284015>
- Platt, J. R. (1964). Strong Inference: Certain systematic methods of scientific thinking may produce much more rapid progress than others. *Science*, 146(3642), 347–353. <https://doi.org/10.1126/science.146.3642.347>
- Popper, K. R. (1959). *The logic of scientific discovery*. Routledge.
- Popper, K. R. (1972). *Conjectures and Refutations: The growth of scientific knowledge* (fourth). Routledge and Kegan Paul.
- Richards, A., French, C. C., Johnson, W., Naparstek, J., & Williams, J. (1992). Effects of mood manipulation and anxiety on performance of an emotional Stroop task. *British Journal of Psychology*, 83(4), 479–491. <https://doi.org/10.1111/j.2044-8295.1992.tb02454.x>
- Rozeboom, W. W. (1997). Good science is abductive, not hypothetico-deductive. In *What If There Were No Significance Tests?* <https://doi.org/10.4324/9781315629049-28>
- Schmidt, F. L., & Hunter, J. E. (1999). Theory Testing and Measurement Error. *Intelligence*, 27(3). [https://doi.org/10.1016/S0160-2896\(99\)00024-0](https://doi.org/10.1016/S0160-2896(99)00024-0)
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston: Houghton Mifflin.
- Siemer, M. (2005). Mood-congruent cognitions constitute mood experience. *Emotion*, 5(3), 296–308. <https://doi.org/10.1037/1528-3542.5.3.296>
- 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(3), 218–226. https://doi.org/10.1207/s15327957pspr0203_5
- Simons, D. J., & Holcombe, A. O. (2014). *Registered Replication Reports*. The Observer.

<https://www.psychologicalscience.org/observer/registered-replication-reports>

Smith, G. T. (2005). On Construct Validity: Issues of Method and Measurement. *Psychological Assessment*, 17(4), 396–408. <https://doi.org/10.1037/1040-3590.17.4.396>

Strauss, M. E., & Smith, G. T. (2009). Construct Validity: Advances in Theory and Methodology. *Annual Review of Clinical Psychology*, 5(1), 1–25. <https://doi.org/10.1146/annurev.clinpsy.032408.153639>

Stroebe, W., & Strack, F. (2014). The Alleged Crisis and the Illusion of Exact Replication. *Perspectives on Psychological Science*, 9(1), 59–71. <https://doi.org/10.1177/1745691613514450>

Tamir, M., & Robinson, M. D. (2007). The Happy Spotlight: Positive Mood and Selective Attention to Rewarding Information. *Personality and Social Psychology Bulletin*, 33(8), 1124–1136. <https://doi.org/10.1177/0146167207301030>

Thagard, P. R. (1978). The Best Explanation: Criteria for Theory Choice. *The Journal of Philosophy*. <https://doi.org/10.2307/2025686>

Thagard, P. R. (1989). Explanatory coherence. *Behavioral and Brain Sciences*. <https://doi.org/10.1017/S0140525X00057046>

Trafimow, D. (2009). The Theory of Reasoned Action. *Theory & Psychology*, 19(4), 501–518. <https://doi.org/10.1177/0959354309336319>

Trafimow, D., & Rice, S. (2009). What If Social Scientists Had Reviewed Great Scientific Works of the Past? *Perspectives on Psychological Science*, 4(1), 65–78. <https://doi.org/10.1111/j.1745-6924.2009.01107.x>

Whitley Jr, B. E., & Kite, M. E. (2012). *Principles of Research in Behavioral Science*. Routledge.

Wilson, T. D., Aronson, E., & Carlsmith, K. (2010). The Art of Laboratory Experimentation. In *Handbook of Social Psychology*. Wiley. <https://doi.org/10.1002/9780470561119.socpsy001002>

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University

Press.

Yarkoni, T. (2020). Implicit Realism Impedes Progress in Psychology: Comment on Fried

(2020). *Psychological Inquiry*, 31(4), 326–333.

<https://doi.org/10.1080/1047840X.2020.1853478>

Footnotes.

1. It seemed pertinent to relegate discussion of a metaphysical issue to a footnote. In the case of hypothetical constructs, it is worth noting that MacCorquodale & Meehl (1948), and Cronbach & Meehl (1955) consider such constructs to refer to attributes that are—in some way—"possessed" by individuals. Thus, some realistic interpretation is attached to psychological constructs in that they designate (or describe) processes, states, or dispositions of people (rather than of the researcher). This is also the interpretation of constructs used in this paper, as the alternative approach of not requiring of constructs to exist in some meaningful sense—for instance viewing them merely as instrumental "useful fictions" (e.g., Peters & Crutzen, 2017; Yarkoni, 2020) makes it difficult to understand how such fictions can be measured (e.g., Michell, 1999, 2005), manipulated (e.g., Gruijters, 2017), or be used as potential explanations for phenomena. This does not imply, of course, that any hypothetical construct therefore *adequately* captures such a process, state, or disposition. For helpful discussion on the nature hypothetical constructs see the classic paper by MacCorquodale & Meehl (1948), and more recently by Lovasz & Slaney (2013).

2. As noted by a reviewer, in some research contexts the primary goal may merely be to examine the effect of an intervention. As is often the case in applied, problem-oriented, research. In such cases, researchers may only be interested in causal description and not in causal explanation. This paper is aimed at the more basic research questions, where manipulation checks play their foremost role, and basic researchers often do want to arrive at causal explanations of phenomena.

3. In their discussion of causal explanation, Shadish, Cook and Campbell emphasize (perhaps overly) unravelling the causal chain connecting a particular cause to an effect—the mechanism. However, inferring the existence of a hypothetical construct, as a cause of a

phenomenon, is often also considered as a causal explanation (though perhaps an unelaborated one). Woodward refers such unelaborated causal variables as “total causes” (Woodward, 2003).

4. Interestingly, while Meehl (e.g., 1990a, 1990b) discusses various obfuscating factors in his papers—including the Crud-factor—Meehl considered the Crud-factor to be mainly problematic for non-experimental research designs. But, while experimental control (e.g., randomization) marginalizes the role of spuriousness (the role of confounding variables) *prior* to intervention, it does not provide control over what happens to which hypothetical constructs *after* intervention. In other words, randomization does not engender specificity of a manipulation (that only the focal construct is affected), nor does it imply the Crud-factor to be without consequence in the context of experiments. See for instance the discussion on nomological shockwaves by Chester & Lasko (2020).

5. Thagard (1978) initially uses the label *consilience*, which (in my read at least) coincides with Kuhn’s *scope* and Popper’s (1972) *satisfactoriness* criteria. Thagard (1989) shifts from *consilience* towards using *explanatory breadth* to denote the different classes of empirical findings explained by a causal hypothesis or theory.