When numbers fail: Do researchers agree on operationalisation of published research?

Matthias Haucke^{12*}, Rink Hoekstra³, and Don van Ravenzwaaij²

¹*Department of Clinical Psychology and Psychotherapy, Freie Universität Berlin, Berlin*

²*Department of Psychometrics and Statistics, University of Groningen, Groningen*

³*Department of Pedagogical and Educational Sciences, University of Groningen, Groningen*

Keywords: Operationalisation, Methodology, Construct validity, Measurement

1. Summary

Current discussions on improving the reproducibility of science often revolve around statistical innovations. However, equally important for improving methodological rigor is a valid operationalisation of phenomena. Operationalisation is the process of translating theoretical constructs into measurable lab-quantities. Thus, the validity of operationalisation is central for the quality of empirical studies. But do differences in the validity of operationalisation affect the way scientists evaluate scientific literature? To investigate this, we adapted three published studies and manipulated the strength of operationalisation. We preregistered our planned study as a registered report. For each study, we designed two tasks in which the validity of the used operationalization is manipulated and sent them to researchers via email. In the first task, researchers were presented with a summary of the Method and Result section from one of the studies and were asked to guess the hypothesis that was investigated via a multiple-choice questionnaire. In a second task, researchers were asked to rate the perceived quality of the study. Our results show that (1) researchers are better at inferring the underlying research question from empirical results if the operationalisation is more valid; but (2) the different validity is only to some extent reflected in a different judgment of the quality of the study. These results combined give partial evidence for the

notion that researchers' evaluations of research results are not affected by operationalisation validity.

*Corresponding author (email: matthias.haucke@gmx.de)

2. Introduction

"Numbers are a creation of our mind. Phenomena do not carry tags with numbers on them, nor do they possess 'quantity' as an intrinsic attribute."

(Roskam, 1989)

A cornerstone of the credibility of a scientific finding is its replicability (2-5). However, the replicability of scientific findings remains low in a range of disciplines, such as Psychology (6) and Biomedicine (7). So far, suggestions to enhance reproducibility of findings often revolve around the use of statistics, for example the false interpretation of p -values or a lack of control for statistical power (3-5). One important methodological concern related to replicability that has received little attention in this discussion is the validity of the operationalisation.

Operationalisation is the process of making a theoretical construct concrete or tangible so that it can be studied via empirical observations (8). Even if one successfully reduces statistical practices that increase false positives, one can still be faced with erroneous conclusions from falsely mapping empirical results onto constructs. That is, although the results may seem convincing, they do not necessarily warrant the conclusions drawn with respect to one's construct or theory. The goal of this paper is not to investigate which studies have a valid operationalisation, but to gather empirical evidence on the extent to which researchers consider the validity of operationalisation when drawing conclusions about empirical findings. This is an important question as a study can lead to convincing statistical results, while actually not operationalising the underlying concept well. Therefore, we believe researchers need to be attuned towards invalid operationalisations. The following study focuses on concepts, which are the basic blocks on which theories stand.

First, we will define the terms *concept*, *measurement* and *operationalisation*. A concept is represented by words often taken from everyday conversations (e.g., justice, creativity, intelligence). The more abstract a concept becomes, the less likely will researchers agree on appropriate measurement strategies (8). Thus, the issue of defining concepts and establishing their relationship to observations is especially relevant in the social sciences, which primarily depend on the investigation of abstract concepts, such as creativity or intelligence. Measurement, on the other hand, is the less abstract process of assigning numbers to observable variables, which can be nominal (e.g., sex), ordinal (e.g., level of education), or continuous (e.g., performance on an intelligence test; 8). Linking those empirical observations to unobservable concepts is called operationalisation (1). Operationalisation is a necessary step to study abstract concepts by defining them in terms of empirical observations. The validity of the theoretical conclusions drawn from data depends on the validity of the operationalisation.

To understand the gap between measure and construct, it can be helpful to distinguish between the concept-as-intended and the concept-as-determined (9). The **concept-as-intended** is one's framework or concept to be investigated (e.g., creativity). On the other hand, the **concept-as-determined** is the actual measurement or empirical variable (e.g., the number of colours used in a drawing task), which serves as the operational definition of the former. To provide empirical support for a theory, the concept-as-determined needs to be a valid representation of the concept-as-intended. If others cannot link constructs to measurement procedures, these procedures and their associated conclusions become less valid. In line with this notion, Feigl stated that "concepts which are to be of value to the factual sciences must be definable by operations which are... intersubjective and repeatable" (10). That is, the overlap between the concepts-as-intended and the concepts-as-determined influence the validity of the conclusions that researchers draw from observed data to theory.

Similarly, Cronbach and Meehl (11) adopt the term "nomological network" to designate the system of law-like relationships that hold between theoretical entities (e.g. intelligence) and their

observable indicators (e.g. IQ tests). A set of theoretical statements becomes a system due to the *semantic* overlap of shared terms. We would, for example, have a hard time deriving a person's intelligence from their favourite pop song. According to Meehl (12), a sign of proper operationalisation is a high interpersonal consensus on how the theoretical terms are linked to observations.

Consensus on the used operationalisation is not always desirable. For instance, Thomas Kuhn argues that scientific advancement happens through violation of consensus, i.e. a paradigm shift (13). Moreover, McLeod (14) argues that concepts do not only help to advance scientific knowledge by exactly representing aspects of the world, but through their open-endedness and epistemic vagueness. In his view concepts are not representing theoretical ideas frozen in time but are part of a continual development. In such a way an epistemically vague or fuzzy concept can inspire exact reformulations, as well as construction of experimental techniques for probing and testing it, if it relates to a representation specifying its structure and causal nature. We argue that, in order to advance from fuzzy to exact concept, individual researchers need to notice problematic operationalisations (i.e. there is no overlap between methodology and intended hypothesis) such that they should affect the conclusion researchers personally draw from study results. Otherwise, compelling empirical results might prevent researchers from detecting that the underlying concepts are still fuzzy, stifling theory development.

Previous discussions of valid operationalisation have centred around the lack of a consistent use of theoretically founded measures, and the stimuli representativeness. Brunswik (15) proposed that "cues" (or stimuli) should be sampled from the participants' typical environment. He stated that the used cues allow drawing conclusions about the (non-observable) construct to the extent that the used cues correctly represent the population of environmental stimuli. But even when the stimuli are a valid sample from the stimuli population, previous researchers have noted that the employed measures can lack a clear theoretical foundation, thus increasing the resulting flexibility during data analysis (i.e. allowing selective reporting (16)). This lack of operationalisation clarity has

been shown and criticized in studies using self-reports (17), experimental studies (18) and fMRI studies (19-21). For example, Elson and colleagues (18) have demonstrated that the Competitive Reaction Time Task, which is a paradigm that measures aggressive behaviour, allows researchers to operationalise the severity of aggressive behaviour via a noise blast's volume, duration or a composite score of both. Elson and colleagues argue that this lack of a theoretically founded measure makes it easier to report those specific outcome variables that happened to be statistically significant, thus increasing the occurrence of false positive findings in the literature. In sum, previous research has concentrated on the representativeness of stimuli and the analytical flexibility that results from theoretically unfounded measures.

In the following study, we will not investigate the lack of clear operationalisation in specific studies, but we will empirically test researchers' interpretation of these operationalisations. Specifically, we will test whether researchers consider the validity of the operationalisation when drawing conclusions about the results of a study. We argue that even when stimuli correctly represent the population of environmental stimuli and there is no variability in operationalising the construct under investigation, there can be a gap between the measure and the construct that was intended to be measured.

Despite the importance of an appropriate operationalisation of concepts, we are not aware of previous studies that examine whether researchers explicitly consider operationalisation when evaluating research. To fill this gap, we have conducted a preregistered study in which we investigated to what extent researchers consider the validity of operationalisation when drawing conclusions about empirical findings. To do so, we developed three fictional scenarios, which are related to existing research in Psychology. The study consisted of two parts. In part 1, we presented empirical outcomes from different studies and assessed whether researchers can reverse engineer the construct under investigation from the used operationalisation. In part 2, we investigated whether the validity of the operationalisation is related to researchers' perception of the study's quality. Together, the two parts allow us to answer the question whether researchers consider the

study's operationalisation when drawing conclusions about the quality of empirical results. The Stage 1 registered report may be found at <https://osf.io/vsfbh/>.

We hypothesized that researchers are more capable to deduce an original hypothesis from more valid operationalizations (and less valid operationalizations lead to lower rates of correctly identifying the original hypothesis). However, we assumed that less valid operationalisations do not result in lower ratings of perceived quality of a study by researchers. In sum, we expect that relatively poor operationalizations have the direct effect of researchers being less capable of reverse-engineering what the underlying research question was, but we expect that those same poor operationalizations do *not* have the indirect effect of researchers rating the study of lower quality.

3. Method

Participants

We have sent emails with a link to an online questionnaire, made with the online survey platform Qualtrics, to the corresponding authors of all articles published in 2015, 2016, 2017, 2018, 2019, 2020 and 2021 from the following journals:

1. *Journal of Experimental Psychology: General*
2. *Psychological Science*
3. *Journal of Abnormal Psychology*
4. *Journal of Consulting and Clinical Psychology*
5. *Journal of Experimental Social Psychology*
6. *Journal of Personality and Social Psychology*

These journals were chosen to represent a sample of researchers in diverse fields of psychology (experimental, social, neuro, and clinical), a sampling strategy previously used by Cramer et al. (23). All duplicates were removed. Participants who did not respond after two weeks received a reminder. Participants were randomly assigned to each of three groups (groups are described in

the Procedure section below). Text of the invitation email, reminder email, preregistered sampling plan and power calculation can be found at <https://osf.io/vsfbh/>. We approached 2981 researchers. In total 325 (10,9%) participants started the online survey, and 66 were excluded due to incomplete responses, resulting in a final sample of 259 participants.

Materials

Research Scenarios

In order to test the influence of operationalisation validity on the perceived quality of the study, we adapted three published studies investigating false memory, mental fluency and cognitive dissonance. For each scenario we manipulate the validity of the operationalization by reducing the strength of the effect, which resulted in three versions that vary in operationalisation validity. Each scenario will be presented with a different validity condition to each participant (see Figure 1). Moreover, presentation of the different scenarios within each group is randomized.

Figure 1. Study design. The three groups will first answer a multiple-choice question and subsequently rate each of the three research scenarios. Scenarios presentation within each group will be randomized so that each scenario will appear in varying order.

Scenario 1 (false memory). In a first scenario, we modified a study by Roediger and McDermott (22). The study investigated the occurrences of false memory, based on how often participants rated an unseen word to be new or old. Participants were presented with two lists of words they had to

memorize. Afterwards participants had to rate words that were semantically related or unrelated to the words on the memorized list. Then, participants had to rate how likely they thought it was that they had already seen the word (see Figure 2).

The results showed that unseen words that were semantically related to the memorized list (i.e., the word ‘*doctor*’ which was not presented in a list of medically related words) increases a false impression of recognition, in comparison to semantically unrelated words. That is, the original study claims that participants falsely remembered ‘*doctor*’, because it’s highly similar to the words on the list. Aside from the existing version of the task (labelled ‘high validity’), we constructed two other versions of the task by varying the words on the second list to be increasingly semantically similar to the unseen word.”

Figure 2. Scenario 1. The validity of the operationalisation was manipulated by varying the words of list 2 to be semantically unrelated, randomly selected or related to the unseen word.

Scenario 2 (perceptual fluency). In a second scenario, we modified a study by Reber and Schwarz (23). The study investigated the influence of perceptual fluency on the perceived truthfulness of a statement. Perceptual fluency was defined as the easiness to read a sentence,

which was manipulated by using a hard to read or an easy to read colour in a between-subjects design (see Figure 3).

The results showed that an easier to read font colour increased the perceptual fluency and thus the truthfulness rating of a sentence (i.e. “How true is this sentence?”). That is, the original study claims that participants rated an easy to read font as more truthful, because an easy to read font increases the speed by which textual information could be cognitively processed. Therefore, aside from the existing version of the task (labelled ‘high validity’), we constructed two other versions of the task by varying the font colour in the second condition to be increasingly more readable, compared to the first condition. The result is that Condition B becomes progressively more similar to read to Condition A.

Figure 3. Scenario 2. The validity of the operationalisation was manipulated by varying the perceptual visibility of the sentence in the second condition.

Scenario 3 (cognitive dissonance). In a third scenario, we modified a study by Festinger and Carlsmith (24) on cognitive dissonance. Participants were first asked to do two lengthy and monotonous tasks, which were supposed to create a negative opinion about the tasks. Afterwards,

participants were asked by the experimenter to convince another student that the tasks were pleasurable. For doing this, the participant either received a high reward (20\$) or a small reward (1\$). Finally, the participants had to rate how enjoyable they found the monotonous tasks (see Figure 4).

The results showed that participants who received a high reward rated the tasks as less enjoyable than participants who received a low reward. The authors concluded that participants experienced a dissonance between having to convince a stranger that the task was pleasurable, although the tasks were in fact very boring. The cognitive dissonance occurs because participants try to reconcile the fact that they did boring tasks with the fact that it did not even pay well, resulting in a relatively high rating of perceived enjoyability of the tasks in an attempt to reduce the cognitive dissonance. This dissonance should be smaller in the high reward condition, as the higher pay would be enough reason as and of itself to participate in the tasks without having to find it intrinsically enjoyable.

Aside from the existing version of the task (labelled 'high validity'), we constructed two other versions of the task by increasing the monetary reward in the second condition such that it was more similar to the monetary reward in the first condition. The result is that cognitive dissonance should be progressively more similar for Condition B compared to Condition A for lower validity.

Figure 4. Scenario 3. The validity of the operationalisation was manipulated by varying the monetary reward in the second condition.

Tasks

Task 1: Guess the Hypothesis. In the first task, we examined whether researchers were more likely to link the methods and results to the research question when the operationalisation is more valid via a multiple choice setting. The presentation consisted of a brief summary of the method and results, as well as their schematic depiction. Each of the three research scenario was presented as published research and the participants were asked to indicate what they would most likely conclude from the study. The exact questionnaire for each of the research scenarios and validity conditions may be found at <https://osf.io/vsfbh/>.

After reading the false memory research scenario, participants were asked to indicate what they thought was a valid conclusion from the presented experiment. They were given four answer choices, of which one was the actual hypothesis that was investigated, whereas the others were

incorrect but plausible alternatives: *“Words that are similar to each other can create false memories”* (correct), *“Words that are similar to each other can decrease attention to the task”*, *“Words that are similar to each other can increase creativity”* and *“Words that are similar to each other can lead to a wider definition of categories”*.

After reading the perceptual fluency scenario, participants were asked to indicate what they thought was a valid conclusion from the presented experiment. They were given four answer choices, of which one was the actual hypothesis that was investigated, whereas the others were incorrect but plausible alternatives: *“Easier to read sentences can increase truth judgments”* (correct), *“Red colour can increase truth judgments”*, *“Blue colour can increase suspicion”*, *“Geographical facts can increase truth judgments”*.

After reading the cognitive dissonance scenario, participants were asked to indicate what they thought was a valid conclusion from the presented experiment. They were given four answer choices, of which one was the actual hypothesis investigated, whereas the others were incorrect but plausible alternatives: *“Higher mental discrepancy can lead to a change in opinion”* (correct), *“Higher monetary rewards can increase negative affect.”*, *“Higher monetary rewards can lead to in-depth cognitive processing of information”*, *“Higher cognitive demands can increase negative affect”*.

Task 2: Rate the Study Quality. After each presented scenario, researchers were asked about the perceived quality of the presented study. Participants answered the following questions: *“How would you judge the quality of the presented study?”* and *“How would you judge the strength of support for the theoretical proposition?”* on a scale from 1 (very low) to 9 (very high).

Procedure

First, participants were asked to indicate their academic position (BSc, MSc, PhD, Postdoc, Assistant Professor, Associate Professor, Full Professor), and their research background (Social Psychology, Clinical Psychology etc.). Then, each participant was presented with three research

scenarios, described above. For each scenario, participants saw the research methodology and results, but not the research question. Each participant also evaluated the perceived quality of the study. Participation took approximately 15 minutes. After participants completed the study, we asked them for their familiarity with each presented scenario (1 = not at all, 9 = very high).

4. Analysis

We analysed the results of the two tasks with a Bayesian test of proportions (25) for Task 1 and a Bayesian analysis of variance (ANOVA) (26) for Task 2. In the Results section, we reported Bayes factor as well as mean and credible intervals. The Bayes factor (BF_{10}) is the relative ratio of the likelihood of the data, given the alternative hypothesis, and the likelihood of the data, given the null hypothesis. For instance, a BF_{10} of 10 indicates that the observed data are 10 times more likely under the alternative hypothesis than under the null hypothesis; a BF_{10} of 1 indicates that the observed data are equally likely under both hypotheses (i.e., the data does not favour one hypothesis over the other) and a BF_{10} of 1/10 indicates that the observed data are 10 times more likely under the null hypothesis than under the alternative hypothesis. In line with the guidelines of Lee and Wagenmakers (27), we interpret a BF_{10} between 1-3 as anecdotal evidence, a BF_{10} between 3-10 as moderate evidence, 10-30 as strong evidence, 30-100 as very strong and >100 as extreme evidence in favour of the alternative hypothesis.

Test of three proportions: Task 1

We will modify the model proposed by Kass and Vaidyanathan (25) to account for three proportions:

$$\log\left(\frac{p1}{1-p1}\right) = \beta + \frac{\psi}{2}$$

$$\log\left(\frac{p2}{1-p2}\right) = \beta$$

$$\log\left(\frac{p_3}{1-p_3}\right) = \beta - \frac{\psi}{2}$$

$$y_1 \sim \text{Binomial}(n_1, p_1)$$

$$y_2 \sim \text{Binomial}(n_2, p_2).$$

$$Y_3 \sim \text{Binomial}(n_3, p_3).$$

In these equations, p_1 is the proportion of people that correctly pick the true research hypothesis in the highly valid condition, p_2 is the proportion of people that correctly pick the true research hypothesis in the medium valid condition, and p_3 is the proportion of people that correctly pick the true research hypothesis in the low valid condition. Proportions p_1 and p_3 are functions of model parameters β and ψ . Nuisance parameter β corresponds to the grand mean of the log odds, whereas the test-relevant parameter ψ corresponds to the log odds ratio of the two extreme groups. We assigned β a standard normal prior and used a zero-centered normal prior with standard deviation σ for the log odds ratio ψ . The analysis was conducted with Stan (28) and the `bridgesampling` R package (29). For ease of interpretation, the results will be shown on the odds ratio scale.

Our analysis primarily focused on testing and quantified the extent to which the data support the null hypothesis: $\psi = 0$ versus the one-sided alternative hypothesis : $\psi > 0$. This one-sided alternative indicates a positive value for ψ , resulting in $p_1 > p_2 > p_3$, thus reflecting our hypothesis that a less valid operationalisation leads to lower rates of deduction of the intended hypothesis. For the specification of the alternative hypothesis for each presented scenario, we assumed a normal distribution for ψ with mean $\mu=0$ and standard deviation $\sigma = 0.4$ (i.e., a mildly informative prior; 28), truncated at zero to take into account that the alternative hypothesis is one-sided, which gives $H_1: \psi \sim N(0, 0.4^2)$. We also conducted parameter estimation for the ψ parameter. For this analysis, we used a two-sided prior. The R code for the analysis can be found online at <https://osf.io/z4gab/>.

Bayesian ANOVA: Task 2

To test whether the validity of the operationalization is related to the perceived quality of the study, we conducted two univariate three- group between-subjects Bayesian ANOVA with a multivariate generalization of the default Cauchy prior, using the statistical software package JASP (30). Although we had not planned to do a post-hoc test, our results made this a valuable additional analysis, so we report these as exploratory. We also conducted an exploratory analysis (Bayesian ANCOVA) including the familiarity of the researcher with the presented scenario.

5. Results

Most respondents indicated that their current academic position is Postdoc (26%), followed by Assistant Professor (22%), Associate Professor (20%), Full Professor (17%), PhD (9%) and Other¹ (6%). Most participants indicated that they work in the field of Social Psychology (36%), followed by Clinical Psychology (19%), Cognitive Psychology (14%), Personality Psychology (9%), Other² (9%), Experimental Psychology (8%), Methodology and Statistics (2%), Neuroscience (2%), Biological Psychology (0,7%) and Medicine (0.3%).

Pre-registered analyses

Task 1: Guess the Hypothesis

The number of correct reverse engineered hypotheses per validity condition can be seen in Table 1.

Validity Condition	Correct Reverse Engineering		
	No	Yes	Total
High	56	203	259
Medium	90	169	259

¹ participants stated that they now work outside of academia.

² participants stated that they work in the field of Developmental/Educational Psychology, Social Work, Marketing and I/O Psychology

Low	117	142	259
-----	-----	-----	-----

Table 1. Correctly/Incorrectly reverse engineered research hypothesis in Task 1.

First, we present the results of the hypothesis test. The null hypothesis postulates that there is no difference in correctly deducing the hypothesis between operationalisation validity groups $H_0: \psi = 0$. The one-sided alternative hypothesis states that the lower the validity of the operationalisation, the lower the proportion of correctly deduced hypotheses $H_1: \psi > 0$. The Bayes factor indicates overwhelming evidence for H_1 , with a $BF_{10} = 666665$, which means that the data are over 600,000 times more likely under the alternative hypothesis than under the null hypothesis (see Figure 5, left side).

Second, we present the results of the parameter estimation. Of interest is the odds ratio of the high validity condition against the low validity condition, defined as (high correct/high incorrect)/(low correct/low incorrect). Figure 5 (right side) shows the median of the resulting posterior distribution of odds ratio in the population equals 2.908, with central 95% credible intervals ranging from 2.001 to 4.239, indicating that a population values of 1 (indicating equivalence between the conditions) is very unlikely.

Figure 5. Bayesian test of three proportions. The left panel shows the one-sided procedure for hypothesis testing and the right panel shows the two-sided procedure for parameter estimation.

To assess the robustness of the Bayes factor to our priori specification, Figure 6 shows BF_{10} as a function of the prior width σ . The Bayes factor appears to be very robust to a wide range of values for σ and shows strictly pro-alternative evidence across the entire range that was examined. In sum, the data supports our hypothesis that operationalization validity influences the researcher's ability to reverse-engineer the correct hypothesis.

Figure 6. The Bayes factor robustness plot. The red dot indicates the prior width setting that results in the maximum BF_{10} . The grey dot indicates BF_{10} for the user specified prior ($\sigma = 0.4$).

Task 2: Rate the Study Quality.

A visual check indicates that perceived quality and perceived support of the theoretical proposition are normally distributed and that the assumption of homogeneity of variance holds. The means and standard deviation of both outcomes can be seen in **Table 2**.

Validity Condition	Perceived quality	Support of
--------------------	-------------------	------------

	(Mean(SD))	theoretical proposition (Mean(SD))
High	4.873 (1.866)	4.293 (2.034)
Medium	4.853 (1.949)	4.305 (1.944)
Low	4.108 (1.883)	3.525 (1.975)
Total	4.611 (1.931)	4.041 (2.016)

Table 2. Means (standard deviations in parentheses) of perceived quality and perceived support of proposition per validity condition.

First, we present the results of the hypothesis test with perceived quality of the study as the outcome. The null hypothesis postulates that there is no difference in perceived quality between operationalisation validity groups $H_0: \delta = 0$. The alternative hypothesis states that the validity of the operationalisation is related to the perceived quality $H_1: \delta \neq 0$. Contrary to our a-priori expectations, we obtained a $BF_{10} = 6154$ which means that the data are over 6,000 times more likely under the alternative hypothesis than under the null hypothesis. The posterior distributions under the alternative hypothesis for each validity condition with perceived quality as an outcome are shown in **Figure 7**, reflecting the pattern of the descriptives in Table 2 that perceived quality in the low validity condition differed substantially from the other two conditions.

Figure 7. The posterior distributions under the alternative hypothesis for each validity condition with perceived quality as an outcome.

Second, we present the results of the hypothesis test with perceived support for the theoretical proposition as the outcome. The null hypothesis postulates that there is no difference in perceived support for the theoretical proposition between operationalisation validity groups $H_0: \delta = 0$. The alternative hypothesis states that the validity of the operationalisation is related to the support for the theoretical proposition $H_1: \delta \neq 0$. Contrary to our a-priori expectations, we obtained a $BF_{10} = 3861$ which means that the data are over 3,000 times more likely under the alternative hypothesis than under the null hypothesis. The posterior distributions under the alternative hypothesis of the average perceived support for the theoretical proposition in the population are shown in **Figure 8**, reflecting the pattern of the descriptives in Table 2 that perceived support for the theoretical proposition in the low validity condition differed substantially from the other two conditions.

Figure 8. The posterior distributions under the alternative hypothesis for each validity condition with perceived support for theoretical proposition as an outcome.

Finally, using JASP we conducted two Bayesian ANCOVA with the covariate ‘familiarity of with presented scenario’ and perceived quality and well as support for theoretical proposition as outcomes. We report the inclusion Bayes factor (BF_{incl}), which quantifies the change from prior inclusion odds to posterior inclusion odds and can be interpreted as the evidence in the data for including a predictor (i.e., In this case familiarity with the presented scenario)(31). With perceived quality as an outcome, for Validity Condition we found $BF_{incl} = ((0.997+5.46e-27)/(0.003+1.41e-30)) / ((0.25+0.25)/(0.25+0.25)) = 332.3$ and for Familiarity $BF_{incl} = ((0.997+0.003)/(5.46e-27+1.41e-30)) / ((0.25+0.25)/(0.25+0.25)) = 1.8e+26$. With perceived support for theoretical proposition as an outcome, for Validity Condition we found $BF_{incl} = ((0.998+2.18e-17)/(0.002+3.54e-21)) / ((0.25+0.25)/(0.25+0.25)) = 499$ and for Familiarity $BF_{incl} = ((0.998+0.002)/(2.18e-17+3.54e-21)) / ((0.25+0.25)/(0.25+0.25)) = 4.6e+16$. Therefore, we can conclude that operationalisation validity condition as well as familiarity influenced perceived

quality of the study.

Interim Conclusion

Overall, we found partial support for our a-priori expectations. Our data are in line with our first hypothesis: we found that when we reduced the validity of the operationalization, researchers are less able to reverse engineer what the underlying hypothesis was. Contrary to our expectations, we found that researchers also rate the study of lower quality when the validity of the operationalization is reduced. Our original expectation was that compelling results can lead to the perception of a high-quality study even if the mapping of the experiment to the underlying research question is poor.

After observing the data, we do notice an interesting pattern in the results. For Task 1 ‘Guess the Hypothesis’, inspection of the descriptives suggests results are different for each of the three conditions. For Task 2 ‘Rate the Study Quality’ however, there seems to be a qualitative divide between the low validity condition on the one hand and the other two conditions on the other hand. It is possible that the manipulation might have been too strong and the unexpected result for our second task might have been driven entirely by the low validity condition. To test this post-hoc explanation, we conducted a set of exploratory analyses, in which we essentially repeat our analysis strategy for both Tasks, but now focusing exclusively on the comparison of the medium and high validity conditions.

Exploratory analyses

For the first task, we test whether the proportion of correctly deduced hypotheses differs between the medium and high valid condition. The two-condition analogue to our preregistered analysis is the Bayesian 2x2 chi square test, which we conducted using the R package “BayesFactor”(32). The null hypothesis states that the medium and high validity condition do not differ in deduced hypothesis $H_0: \delta = 0$. The alternative hypothesis states that the validity of the operationalisation is related to the proportion of correctly deduced hypotheses $H_1: \delta \neq 0$. Under

the null hypothesis, the prior for both proportions is a single uniform distribution. Under the alternative hypothesis, each proportion separately has a uniform prior (33). The Bayes factor indicates strong evidence for the alternative hypothesis, with a $BF_{10} = 25$ which means that the data are approximately 25 times more likely under the alternative hypothesis than under the null hypothesis.

For the second task, the two-condition analogue to our preregistered analysis is the Bayesian t -test, which we conducted using the R package “BayesFactor”(32). The first t -test compared perceived quality of the study between the medium and high validity conditions. The null hypothesis postulates that there is no difference in perceived quality perception between the two operationalisation validity groups $H_0: \delta = 0$. The two-sided alternative hypothesis states that the validity of the operationalisation is related to the perception of quality $H_1: \delta \neq 0$. We used a default Cauchy prior distribution with $r = 1/\sqrt{2}$ for effect size parameter δ . In line with our post-hoc expectation, we found a $BF_{10} = 0.10$ ($BF_{01} = 1/BF_{10} = 10.2$), which means that the data are approximately 10 times more likely under the *null* hypothesis than under the alternative hypothesis.

The second t -test compared perceived support of the theoretical proposition between the medium and high validity conditions. The null hypothesis postulates that there is no difference in perceived support of the theoretical propositions between the two operationalisation validity groups $H_0: \delta = 0$. The two-sided alternative hypothesis states that the validity of the operationalisation is related to the perception of support for the theoretical proposition $H_1: \delta \neq 0$. We used a default Cauchy prior distribution with $r = 1/\sqrt{2}$ for effect size parameter δ . In line with our post-hoc expectation, we found a $BF_{10} = 0.10$ ($BF_{01} = 1/BF_{10} = 10.2$), which means that the data are approximately 10 times more likely under the *null* hypothesis than under the alternative hypothesis.

6. Discussion

Many findings from the psychological literature are not replicable (6, 34), indicating a

structural and methodological problem. Most recommendation to address this pertain to publication methods, data collection or data analysis (35-38). Yet, even the most advanced statistical inferences become redundant if operationalisation, the process of translating theoretical constructs into measurable lab-quantities, fails. Therefore, researchers need to be attuned towards invalid operationalisation. In this study we investigated to which extent researchers consider the validity of operationalisation when drawing conclusions about empirical findings.

A sign of proper operationalisation is a high interpersonal consensus on how the theoretical terms are linked to observations (12), in our study we found that researchers are better at inferring the underlying research hypothesis from empirical results in more valid operationalisation scenarios. Therefore, we conclude that we have successfully manipulated the validity of the operationalisation. We found mixed evidence for the extent to which the validity of the operationalization affects the perceived quality of the study. Specifically, we showed that in our Medium and High validity conditions, researchers were less capable to deduce what the tested hypothesis was, but this did not influence either their perception of the quality of the study nor the support for its theoretical proposition. Thus, we found some support for the notion that the validity of the operationalisation does not affect researchers' evaluation of an empirical result.

Yet, an extreme change of operationalisation validity does change the evaluation of empirical results. We found the effect of operationalisation validity on judgment of empirical results only between the medium and high operationalisation validity condition, not when we included the low validity condition. In the low validity operationalisation condition, we reduced the experiment to absurdity (i.e., using the same word category in list A and B in the false memory scenario, using the same colour strength in the perceptual fluency scenario, or giving almost the exact same amount of money in the cognitive dissonance scenario). In the medium validity operationalisation condition, on the other hand, the inherent logic of the experiment was preserved (i.e., using different word categories in list A and B in the false memory scenario, using different colour strength in the perceptual fluency scenario, giving substantially different amounts of money in the cognitive dissonance scenario). Thus, we think it is very likely that our manipulation in the low validity

condition was too strong, turning the research scenario into absurdity.

A reason for the lack of researchers' attention to operationalisation could be the shared system of beliefs about how Psychology works as a science, which influences the dominant methodological practices and the content of methodological education. Psychology is strongly embedded in the experimental tradition, going back to Wilhelm Wundt, which emphasizes empiricism (39). One way of empiricism to gather objective knowledge about phenomena is the translation of observation into numbers (i.e., making a psychological attribute quantifiable). In line with this, Fechner (40), a pioneer of experimental psychology, stated "*As an exact science psychophysics, like physics, must rest on experience and the mathematical connection of those empirical facts* (Fechner, 1860, p. xxvii)."

However, the link between concept-as-determined (i.e., actual measurement) and concept-as-intended is non-numerical and subjective. This problem is especially relevant to a range of psychological studies, in which, most concepts-as-intended are highly abstract (e.g., intelligence, empathy, depression). The problem of imprecise concepts-as-intended is further exacerbated by the fact that a majority of published articles do not report the validity of their measured construct, but instead focus on psychometric properties that can be numerically quantified (e.g. reliability measures such as Cronbach's α) (41, 42). Moreover, scientists might be hesitant to question established psychological constructs, as construct are often embedded in "generative entrenchment", meaning that once a concept has been established (e.g., social anxiety disorder) many other concepts (e.g., fear conditioning), theories (e.g., reinforcement learning) or practices (e.g., cognitive behaviour therapy) depend on it (43). Finally, there is also a lack of attention on operationalisation during the education of psychological researchers. A review of graduate training in psychology has shown that few departments offered a full course on measurement, such as classical test theory (20–24% depending on the topic) (44).

Already in 1967 Paul Meehl stated that there is little theoretical progress in psychological science, with theories tending to come and go without ever being decisively refuted or accepted (12, 44). Michell (45) further argued that the implicit claim that a psychological attribute (e.g., a

personality trait, cognitive ability, or mental disorder) is quantitative and relates to other attributes quantitatively, needs to be treated as a falsifiable theory. For example, if a researcher tries to predict a person's creativity via their educational background and therefore measure units of creativity, they are accepting the hypothesis that creativity has a quantitative structure (e.g., a creativity unit can be five times another creativity unit); an assumption that may well be false. Therefore, a numerical assignment procedure alone cannot produce scientific measurement. The meaning of a scientific concept and is set by the operations (i.e., the measurement procedures) that were used to identify them (46). If problematic operationalisations go unnoticed, Psychology as a science might get stuck in a process of producing statistically significant results that barely link to their intended construct.

Many voices have been raised calling for theories to be more formal and precise, strengthening the link between theory, construct, and hypothesis through mathematical formulation (47-50) and improving publication methods (37). Yet, mathematical models are simply abstractions of scientific problems, thus they can aid scientific inference only to the extent that the abstraction is appropriate to the theory and to the concept-as-intended (51, 52). Using more and more precise theoretical formulations can therefore only solve part of the problem and, in the worst case, let us falsely believe that we can somehow “overcome” the inherent abstract nature of the constructs under investigation. Moreover, improving publication methods does not compensate for weak theory building or low operationalisation validity (53). For instance, why should we judge an operationalization as more valid only because it was preregistered? We believe the real solution is for researchers, editors, and reviewers to become more attuned to problematic operationalization.

Limitations and future studies

In the ‘Guess the Hypothesis’ task, we used multiple choice items which impose a context in which the hypotheses are reverse engineered. Thus, we cannot exclude the possibility that other multiple-choice items would have interacted differently with the manipulation of operationalisation validity and thus would have changed researchers’ capacity to reverse engineer the hypothesis. However, given that our interest was in a relative effect, a difference in proportion

correct between different validity conditions, and not in an absolute effect, we do not believe this affects our conclusions.

Related to the previous limitation, researchers who thought that none of the options were valid were not given an option to indicate this. A way around this would have been to include open answer formats. We were concerned that this would make the analysis more subjective, as it would have required us to make judgment calls on whether the descriptions of participants did or did not match the intended hypothesis. However, with multiple raters this may well have been a preferable set-up.

Our study results are limited by the three chosen research scenarios, and it could be that the observed effect is limited to these scenarios. Although we found the reported effects consistently for each of the three research scenarios, a fruitful avenue for future research on this topic might be to have a wider range of experimental empirical scenarios to present to participants. In addition, these could include a wider range of operationalisation validity levels than the ones used in our study. Our findings might not be restricted to Psychology but apply to any scientific claim that is made using statistical procedures. Therefore, it would be interesting for future studies to investigate whether a similar “blind eye” to operationalisation exists in other disciplines (e.g., economics, medicine).

Conclusion

Progressing psychological science does not only require advancing statistical inference, but also the successful translation of theoretical constructs into measurable lab-quantities (i.e. operationalisation). Therefore, researchers need to be attuned to empirical studies that have an invalid operationalisation. In this study we found that researchers are better at inferring the underlying research question from empirical results, if the operationalisation is more valid. Moreover, only an extreme change of operationalisation validity lowers the perceived quality of a presented study. Thus, even if the operationalisation of an empirical finding is lowered, it may not affect researchers’ perception of the study’s quality. Our study indicates that most researchers are not considering the validity of operationalisation when evaluating scientific findings.

Ethics statement

This study was approved by the ethics commission of the University of Groningen (Code 17212-O).

Data accessibility

The article's data, materials and statistical/qualitative analysis are available on the open science framework at <https://osf.io/vsfbh/>.

Competing interests

The authors of this study have no competing interests.

Authors' contributions

MH drafted the manuscript, prepared the material, designed the study, gathered the data, and analysed the data; RH drafted the manuscript, prepared the material and participated in designing the study; DvR drafted the manuscript, prepared the material, designed the study and analysed the data.

Funding statement

DvR was supported by a Dutch scientific organization VIDI fellowship grant

(016.Vidi.188.001).

References

1. Roskam EE. Operationalization, a superfluous concept. *Quality and Quantity*. 1989;23(3-4):237-75.
2. Bell P, Staines P, Michell J. *Evaluating, doing and writing research in psychology: A step-by-step guide for students*: Sage; 2001.
3. Greenland S, Senn SJ, Rothman KJ, Carlin JB, Poole C, Goodman SN, et al. Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations. *European journal of epidemiology*. 2016;31(4):337-50.
4. Sterne JA, Smith GD. Sifting the evidence—what's wrong with significance tests? *Physical Therapy*. 2001;81(8):1464-9.
5. Vankov I, Bowers J, Munafò MR. Article Commentary: On the Persistence of Low Power in Psychological Science. *Quarterly journal of experimental psychology*. 2014;67(5):1037-40.
6. Collaboration OS. Estimating the reproducibility of psychological science. *Science*. 2015;349(6251):aac4716.
7. Begley CG, Ellis LM. Drug development: Raise standards for preclinical cancer research. *Nature*. 2012;483(7391):531.
8. Mueller CW. Conceptualization, operationalization, and measurement. *Encyclopedia of social science research methods*. 2004;1(163):1-8.
9. de Groot AD. *Methodologie*: Hague: Mouton; 1969.
10. Feigl H. Operationism and scientific method. *Psychological Review*. 1945;52(5):250.
11. Cronbach LJ, Meehl PE. Construct validity in psychological tests. *Psychological bulletin*. 1955;52(4):281.
12. Meehl PE. Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of consulting and clinical Psychology*. 1978;46(4):806.
13. Kuhn TS. *The structure of scientific revolutions*: University of Chicago press; 2012.
14. MacLeod M. Rethinking scientific concepts for research contexts: The case of the classical gene. *Scientific Concepts and Investigative Practice*. 2012;3:47.
15. Brunswik E. *Perception and the representative design of psychological experiments*: Univ of California Press; 1956.
16. John LK, Loewenstein G, Prelec D. Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science*. 2012;23(5):524-32.
17. Weinstein Y. Mind-wandering, how do I measure thee with probes? Let me count the ways. *Behavior research methods*. 2017:1-20.
18. Elson M, Mohseni MR, Breuer J, Scharkow M, Quandt T. Press CRTT to measure aggressive behavior: The unstandardized use of the competitive reaction time task in aggression research. *Psychological Assessment*. 2014;26(2):419.
19. Carp J. On the plurality of (methodological) worlds: estimating the analytic flexibility of fMRI experiments. *Frontiers in neuroscience*. 2012;6:149.
20. Carp J. The secret lives of experiments: methods reporting in the fMRI literature. *Neuroimage*. 2012;63(1):289-300.
21. Poldrack RA, Yarkoni TJ. From brain maps to cognitive ontologies: informatics and the search for mental structure. 2016;67:587-612.
22. Roediger HL, McDermott KB. Creating false memories: Remembering words not presented in lists. *Journal of experimental psychology: Learning, Memory, and Cognition*. 1995;21(4):803.
23. Reber R, Schwarz NJC, cognition. Effects of perceptual fluency on judgments of truth. 1999;8(3):338-42.
24. Festinger L, Carlsmith JMT. *Psychology*. Cognitive consequences of forced compliance. 1959;58(2):203.
25. Kass RE, Vaidyanathan SK. Approximate Bayes factors and orthogonal parameters, with application to testing equality of two binomial proportions. *Journal of the royal statistical society Series B (Methodological)*. 1992:129-44.
26. Rouder JN, Morey RD, Speckman PL, Province JMJ. Default Bayes factors for ANOVA

designs. 2012;56(5):356-74.

27. Lee MD, Wagenmakers E-J. Bayesian cognitive modeling: A practical course: Cambridge university press; 2014.
28. Team SD. RStan: the R Interface to Stan. R package version 2.17. 3. 2018.
29. Gronau QF, Singmann H, Wagenmakers E-J. Japa. Bridgesampling: An R package for estimating normalizing constants. 2017.
30. Team J. JASP (version 0.9)[computer software]. 2018.
31. van den Bergh D, Van Doorn J, Marsman M, Draws T, Van Kesteren E-J, Derks K, et al. A tutorial on conducting and interpreting a Bayesian ANOVA in JASP. *L'Année psychologique*. 2020;120(1):73-96.
32. Morey RD, Rouder JN, Jamil T, Morey MRD. Package 'bayesfactor'. URL <http://cran.r-project.org/web/packages/BayesFactor/BayesFactor.pdf> (accessed 10/06/15). 2015.
33. Gunel E, Dickey J. Bayes factors for independence in contingency tables. *Biometrika*. 1974;61(3):545-57.
34. Hagger MS, Chatzisarantis NL, Alberts H, Anggono CO, Batailler C, Birt AR, et al. A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*. 2016;11(4):546-73.
35. Munafò MR, Nosek BA, Bishop DV, Button KS, Chambers CD, Du Sert NP, et al. A manifesto for reproducible science. *Nature human behaviour*. 2017;1(1):1-9.
36. Asendorpf JB, Conner M, De Fruyt F, De Houwer J, Denissen JJ, Fiedler K, et al. Recommendations for increasing replicability in psychology. *European journal of personality*. 2013;27(2):108-19.
37. Nosek BA, Beck ED, Campbell L, Flake JK, Hardwicke TE, Mellor DT, et al. Preregistration is hard, and worthwhile. *Trends in cognitive sciences*. 2019;23(10):815-8.
38. Muthukrishna M, Henrich J. A problem in theory. *Nature Human Behaviour*. 2019;3(3):221-9.
39. Blumenthal AL, Danziger K. Wilhelm Wundt in history: The making of a scientific psychology: Springer Science & Business Media; 2001.
40. Fechner GT. *Elemente der psychophysik*: Breitkopf u. Härtel; 1860.
41. Flake JK, Pek J, Hehman E. Construct validation in social and personality research: Current practice and recommendations. *Social Psychological and Personality Science*. 2017;8(4):370-8.
42. Zumbo B, Chan E. *Validity and validation in social, behavioral, and health sciences (Social Indicators Research Series, Vol. 54)*. Cham: Springer International Publishing. 2014.
43. Eronen MI, Bringmann LF. The theory crisis in psychology: How to move forward. *Perspectives on Psychological Science*. 2021;1745691620970586.
44. Meehl PE. Theory-testing in psychology and physics: A methodological paradox. *Philosophy of science*. 1967;34(2):103-15.
45. Michell J. Quantitative science and the definition of measurement in psychology. *British journal of Psychology*. 1997;88(3):355-83.
46. Bergmann G, Spence KW. The logic of psychophysical measurement. *Psychological Review*. 1944;51(1):1.
47. Gigerenzer G. Personal reflections on theory and psychology. *Theory & Psychology*. 2010;20(6):733-43.
48. Oberauer K, Lewandowsky S. Addressing the theory crisis in psychology. *Psychonomic bulletin & review*. 2019;26(5):1596-618.
49. van Rooij I, Baggio G. Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *Perspectives on Psychological Science*. 2020;1745691620970604.
50. Devezir Berna, Navarro Danielle J., Vandekerckhove Joachim and Ozge Buzbas Erkan. The case for formal methodology in scientific reform. *R. Soc. open sci*. 2021;8200805200805. <http://doi.org/10.1098/rsos.200805>
51. Navarro DJ. Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior*. 2019;2(1):28-34.
52. Kellen D. A model hierarchy for psychological science. *Computational Brain & Behavior*. 2019;2(3):160-5.

53. Szollosi, Aba, David Kellen, Danielle Navarro, Rich Shiffrin, Iris van Rooij, Trisha Van Zandt, and Chris Donkin. 2019. "Is Preregistration Worthwhile?." PsyArXiv. October 31. doi:10.1016/j.tics.2019.11.009.