Measurement Schmeasurement: Questionable Measurement Practices and How to Avoid Them

Jessica Kay Flake

Department of Psychology | Quantitative Psychology & Modelling

McGill University

@jkayflake, kayflake@gmail.com

Eiko I. Fried

Department of Clinical Psychology

Leiden University

@EikoFried, eikofried@gmail.com

Author contributions: JKF & EIF wrote the paper and contributed equally.

Corresponding author: JKF, kayflake@gmail.com

This manuscript was accepted for publication in *Advances in Methods and Practices in Psychological Science* in June of 2020. This [preprint](#) is the peer-reviewed, not yet copy edited version and may differ from the final version published in the journal. Previous versions were reviewed by Eric Hehman, Yoel Inbar, Kim Luijken, and Donald Robinaugh, who all provided thoughtful feedback that improved the manuscript. The content was originally developed for a workshop presented at the Society for the Improvement of Psychological Science annual meeting in 2018 and during a symposium of a similar name at the Association for Psychological Science meeting in 2018.

Abstract

In this paper, we define questionable measurement practices (QMPs) as decisions researchers make that raise doubts about the validity of the measures, and ultimately the validity of study conclusions. Doubts arise for a host of reasons including a lack of transparency, ignorance, negligence, or misrepresentation of the evidence. We describe the scope of the problem and focus on how transparency is a part of the solution. A lack of measurement transparency makes it impossible to evaluate potential threats to internal, external, statistical conclusion, and construct validity. We demonstrate that psychology is plagued by a *measurement schmeasurement* attitude: QMPs are common, hide a stunning source of researcher degrees of freedom, pose a serious threat to cumulative psychological science, but are largely ignored. We address these challenges by providing a set of questions that researchers and consumers of scientific research can consider to identify and avoid QMPs. Transparent answers to these measurement questions promote rigorous research, allow for thorough evaluations of a study's inferences, and are necessary for meaningful replication studies.

Introduction

A foundational part of science is defining and measuring what we study. Before we can study the etiology of depression or the efficacy of an intervention to treat it, we must define the construct under investigation and build instruments that measure it accurately. This is difficult because psychological constructs like depression, personality traits, attitudes, or cognitive abilities often do not permit direct observation. We can directly observe the height of a person next to us on the bus, but we often have limited insight into their psychological processes, states, and attributes. If our measures don't accurately capture depression, we will deny treatment to those who need it, while prescribing medication to others who do not need it, exposing them to potentially serious side-effects. Neither rigorous research design, advanced statistics, nor large samples can correct such false inferences.

Identifying and defining constructs is central to developing theory in psychology, because developing some initial way to measure constructs must occur before they can be studied empirically. Construct validation—collecting evidence that the instruments we build actually measure the constructs we claim they measure—is a difficult and necessary part of our research process (Cronbach & Meehl, 1955). It has taken many programs of research, thousands of studies, and decades of work to identify, define, and measure depression. And still, concerns remain as to the validity of depression as a construct, and the use of instruments that measure it (Fried, 2015, 2017; Parker, 2005).

Despite the foundational role that measurement plays in the validity of study conclusions, important information regarding measurement in scientific manuscripts in the social sciences is absent. Measures are used without reference to their source, they lack evidence that they

actually measure the purported constructs, and unjustified measurement flexibility abounds. Such questionable practices seem to go on unseen, but the few meta-scientific studies on measurement practices shine light on a grim state of affairs. Barry et al. (2014) reported that between 40-93% of measures used across 7 journals in educational behavior lacked validity evidence, and Weidmann et al. (2017) reported that among the 356 measurement instances coded in their review of emotions, 69% were used with no reference to prior research or a systematic development process. In their review of the relation between technology use and well-being, Orben and Przybylski (2019) report that researchers pick and choose within and between questionnaires, “making the pre-specified constructs more of an accessory for publication than a guide for analyses. (p. 181)”

When we lack validity evidence for measures, we lack the necessary information to evaluate the overall validity of a study’s conclusions. Further, recent research on commonly used measures in social and personality psychology showed that measures with less published validity evidence were less likely to show strong construct validity evidence when evaluated in new data (Hussey & Hughes, 2019; Shaw et al., in press). The lack of information about measures is a critical problem that could stem from underreporting, ignorance, negligence, misrepresentation, or some combination. Importantly, regardless of why the information is missing, a lack of transparency undermines the validity of psychological science. To shine more light on these issues, we introduce and define the term *questionable measurement practice* (QMP). Questionable measurement practices are decisions researchers make that raise doubts about the validity of measure use in a study, and ultimately the study’s final conclusions. We demonstrate that reporting all measurement decisions transparently—irrespective of whether

the lack of transparency stems from ignorance, negligence, or nefarious intent—is a crucial first step toward improving measurement practices. We start with transparency because a lack of it prevents the evaluation of all aspects of a study’s validity: internal, external, statistical conclusion, and construct validity (Shadish et al., 2002). We show that QMPs are ubiquitous, largely ignored in the literature, provide researchers with ample degrees of freedom that can be exploited to obtain desired results, and thus pose a serious threat to cumulative psychological science. We address these challenges by providing a list of questions that those planning, preregistering, conducting, reviewing, and consuming research can use to identify QMPs, avoid a lack of transparency regarding a study’s measurements, and confront egregious practices.

Questioning the Validity of Psychological Science

In the last decade, psychology built a vernacular for describing aspects of research that undermine the conclusions of the study. These include, among others, researcher degrees of freedom, p-hacking, hypothesizing after results are known, motivated reasoning, fishing. These terms, often discussed under the umbrella term questionable research practices (QRP), tend to span three main issues that potentially co-occur: a lack of transparency, ignorance or negligence on the part of the researcher, and the intent to misrepresent the data. The general emphasis is that there are attributes of the study that raise questions, concerns, or grave doubts about the study’s conclusions.

QRPs have been defined as the exploitation of ambiguities in what is acceptable practice for the purpose of obtaining a desired result (Banks et al., 2016; John et al., 2012). Readers may question various design, analytic, and reporting practices because it is unclear if authors

presented biased evidence to favor a particular claim. QRP is a diffuse term, in that it has been used to describe overtly unethical practices carried out by researchers to mislead, but also to describe ambiguous practices. For example, there might be good reasons for excluding participants, but when reasons are not reported transparently, the practice is questionable. Such ambiguous QRPs are sometimes discussed as completely unintentional on the part of the researcher, rather than with nefarious intent. This conflation of intent and transparency in the discussions about QRPs led some to urge us to focus exclusively on transparency (Banks et al., 2016; Fanelli, 2013; Sijtsma, 2016); Wigboldus and Dotsch (2016) argued to use the term *questionable reporting practices* over questionable research practices.

A focus on transparency is necessary because research entails many decisions at all stages of the study pipeline, such as theory, design, data collection, analysis, and reporting. Even for the scientist with the best of intentions, these decisions can be difficult to navigate. Flexibility is inherent in the research process and exists regardless of the researcher's conscious intent to exploit it. This flexibility has been described as a 'garden of forking paths' (Gelman & Loken, 2013), with each decision taking one down a potentially different path. Critically, the decisions made along the way can lead to fundamentally different conclusions, even when researchers begin with the same raw data and have the same research question (Silberzahn et al., 2018; Steegen et al., 2016). Some may intentionally explore every forking path in the garden to find a significant result, whereas others may have simply gotten lost, taking a few wrong turns.

Taken together, questionable research or reporting practices can range from accidental omission of information to fraud, and include justifiable practices, ignorant and negligent

practices, practices meant to mislead and misrepresent, downright lying, and many combinations. Consistent with calls for disclosure (Simmons et al., 2011), preregistration (Cybulski et al., 2016; Nosek et al., 2019), and open science more generally (Aguinis et al., 2018; LeBel et al., 2017; Nosek & Bar-Anan, 2012), transparency goes a long way toward confirming or allaying doubts about the validity of study conclusions. These calls for transparency, as well as reporting guidelines (e.g., CONSORT), prescribe the transparent reporting of measurement practices. However, problematic measurement practices persist and there lacks a thorough examination of this problem and how to address it. Transparency is the first and necessary step toward getting questions answered about the validity of our work, and this pertains to measurement just as much as it pertains to statistical analyses.

Questionable Measurement Practices

Researchers make decisions in all phases of the research process, including how to measure the psychological constructs they are studying. Here we use the term ‘measurement’ to describe any approach that researchers take to create a number to represent a variable under study. On the one hand, this definition of measurement is broad, encompassing the entire process that requires decisions from theorizing about the nature of a given phenomenon to operationalization to analysis. On the other hand, we do not discuss cases of qualitative measurement that do not result in a number, which undoubtedly can also raise questions about validity. Issues of qualitative measurement and transparency are out of the scope of the current paper, but an important area for future work.

In psychology, the possible approaches to measurement are vast, and the measures in a study introduce a stunning source of flexibility with copious decisions to navigate. We can think

of no single psychological construct for which there exists exactly one rigorously validated measure that is universally accepted by the field, with no degrees of freedom regarding how to score it. Measurement in the social sciences usually comes with a garden of forking paths. For example, data collected with a 10-item questionnaire can be summarized in one sum-score 1,023 different ways by summing different subsets of items (e.g., using items 1-3 to calculate a sum-score or items 2-10 or items 6-8, and so on), and even more possibilities exist when allowing for multiple subscales or different analytic techniques to obtain a final score. The purpose of the current paper is to focus on the need for transparency regarding the measurement in empirical studies. We argue that transparency is a *necessary first step* if we are to confront ignorance, negligence, and misrepresentation of measure use and reform our measurement standards with the goal to improve psychological research.

We define questionable measurement practices as decisions researchers make that raise doubts about the validity of measure use in a study and ultimately the final conclusion. We introduce the term QMP to raise awareness of this foundational aspect of the research process that has largely been obscured in reporting, and threatens all aspects of a study's validity, as we demonstrate below. Like QRPs, QMPs can range from a lack of transparency all the way to the kinds of misleading practices the community generally regards as questionable.

QMPs Threaten All Aspects of Validity

QMPs are not just philosophical threats that researchers specialized in measurement should concern themselves with: they are ubiquitous and highly problematic. A lack of information about the measures in a study introduces uncertainty in all aspects of a study's validity, making QMPs broader than QRPs, which generally focus on distorting statistical results. Shadish et al.

describe four types of validity that contribute to the overall validity of a conclusion: internal validity, external validity, statistical conclusion validity, and construct validity (Shadish et al., 2002). Transparency is crucial because studies with more validity evidence can make stronger claims than studies with less validity evidence. Our discussion of validity is circumscribed, but we define each validity type and provide an example of how issues of opaque measurement can threaten it.

Internal validity captures the aspects of a study's design that support causal claims about the relations between variables. Measures can threaten internal validity. For example, if we lack the information to determine if the measurement properties in a study change across two treatment conditions or across time, we cannot thoroughly evaluate internal validity. External validity concerns the generalizability of findings across different populations and settings. A lack of transparency regarding measurement can threaten external validity when we lack the necessary information to evaluate if the measures are sample or population specific. Statistical conclusion validity captures whether conclusions from a statistical analysis are correct. This is difficult to evaluate when undisclosed measurement flexibility generates multiple comparisons in a statistical test, which can bias the statistical test results. We note that much of the discussion regarding questionable research practices focus on statistical conclusion validity. Finally, construct validity encompasses how the constructs in the study are operationalized and is threatened when we lack the requisite information to determine how well the measure captures the construct.

Every claim from a study has some number of threats to its validity, and with this validity typology we can categorize those threats into internal, external, statistical conclusion, or

construct. Detailed information about the measures in a study is needed to fully evaluate each type of validity and the study's conclusions. Scientists can agree or disagree about what constitutes the necessary evidence, what the best practices are, or what methodologies and designs are most rigorous. Three reviewers may have three different ideas about the most rigorous way to measure depression, but reporting measurement practices transparently *enables* the scientific community to vet such validity claims in the first place.

QMPs cover a whole range of issues including a lack of transparency, ignorance, negligence, and misrepresentation—all of which betray a *measurement schmeasurement* attitude. If we took the validity of our measure use more seriously we would report more information, provide better training to early career researchers, and demand rigorous practices during the review process. Given the breadth of the problem, this paper is focused on transparency and only touches on issues of ignorance, negligence, and misrepresentation. We cannot have a debate about how bad our measurement practices are if we are largely missing the information to evaluate them. Transparency does not automatically make our science more rigorous, it facilitates rigor by allowing us to evaluate the evidence thoroughly and accurately. Eliminating questions about how researchers measured the constructs in their study is a first and necessary step to evaluating and ultimately bolstering the validity evidence.

Questions to Answer that Promote Validity of Measure Use

When we lack information about the measures in a study, we lack the information needed to evaluate the validity of the study. In this section, we provide questions researchers should consider in order to identify QMPs and avoid them. These questions can be used to guide researchers in planning and preregistering their studies. Additionally, they enable

reviewers, editors, and consumers of psychological research to identify QMPs and facilitate the critical evaluation of a study's validity. They do not dictate the best or correct way to create or use a measure, how to evaluate the validity of a measure, or which validity theories or psychometric models to use. In other words, this paper is not a tutorial on scale development, validity theory, or psychometrics. Resources on those practices and methodologies are abundant, and we have curated a list of them on the Open Science Framework (<https://osf.io/zrkd4/>). Instead, the questions in the present paper are aimed at promoting validity by reporting the requisite information needed to evaluate how the measures in a study were selected and used.

How to Use These Questions

Here we introduce a list of questions that, when answered, help avoid QMPs by increasing transparency (*Table 1*). If these questions are unanswered readers may doubt the validity of the measure use and final conclusions of the study. They may also suspect ignorance, negligence, or misrepresentation of the data, though identifying those causes is not the focus of these questions. Answering all questions thoroughly does not automatically make study inferences valid but gives the scientific community information needed to critically vet scientific claims and support authors in calibrating them.

To confront the lack of transparency and validity of our measurement practices there are two levels to consider: decisions made by researchers while conducting individual studies, and systematic questionable practices that appear again and again in the literature. The questions we list first address how each individual researcher can bolster their transparency and dispel doubts about the validity of measure use in a given study. Then we consider how these

unanswered questions raise grave concerns about the validity of results in our field as a whole by presenting findings from the meta-scientific measurement literature. Meta research is often required to shine light on pervasive questionable practices. For example, a meta-analysis of clinical trials can reveal severe publication bias, indicating that many individual researchers may engage in questionable practices. As such, we list six questions and for each of the questions we outline the information needed to answer the question, advice for how to obtain that information, ramifications of leaving that question unanswered; and examples from the meta-science literature that indicate pervasive questionable practices in our field.

Table 1. *Questions to promote transparency of measurement practices*

Question	Information to transparently report
What is your construct?	Define construct Describe theories & research supporting the construct
How do you operationalize your construct?	Describe measure and administration procedure Match measure(s) to construct(s)
Why do you select your measure?	Justify measure selection Report existing validity evidence
How do you quantify your measure?	Describe response coding & transformation Describe items/stimuli per score Describe calculation of score Describe all conducted (e.g. psychometric) analyses
Do you modify the measure? If so, how and why?	Describe any modifications Describe if modifications occurred before or after data collection Provide justification for modifications
Do you create the measure on the fly?	Justify why you do not use an existing measure Report all measurement details listed above Describe all available validity evidence and if there is no evidence, report that

1. What is your construct?

The first question to answer is: what construct are you trying to measure? What *is* it? Answering this question requires substantive expertise, involves reviewing the literature, and reading about the theories of the phenomenon under study. It may require, depending on the existing literature, further development of a theoretical model or paradigm. There are likely different facets of the construct of interest—they could be largely orthogonal or have correlated dimensions.

Reporting clearly what the construct is and how it is defined in relation to the theories that support the construct allows readers to agree or disagree about the theoretical underpinnings of the construct. For example, there are numerous definitions and theories about psychological constructs such as emotions, attitudes, mental disorders, and personality traits. If the construct is not defined clearly at the outset, the ambiguity in what is being measured will make navigating all the future measurement decisions difficult and present many opportunities to exploit this ambiguity, knowingly or otherwise. With no construct definition as a guide, it is easy to get lost in the garden of forking measurement paths.

After theoretically defining your construct, it needs to be operationalized. This builds on the work to theoretically define the construct by specifically defining it in such a way that it is measureable. As an example, Robinaugh et al. (2019) provided a formalized theory of panic disorder, which defined all variables and processes relevant to the construct of panic. Such theorizing helps to operationalize the construct and processes of panic disorder, and, subsequently, to select a particular measurement that is consistent with the stated theory.

Transparently reporting what construct is being measured and how it is operationalized goes a long way toward promoting rigor by encouraging thorough evaluation of the theory underling the construct. Reviews on a number of constructs including giftedness (Dai & Chen, 2013), charisma (Antonakis et al., 2016), and control (Skinner, 1996) demonstrate that constructs are often ill-defined, which leads to confusion and invalid inferences in the literature. A lack of conceptual clarity for a construct is the first step toward measurement heterogeneity, flexibility, and the profusion of untested measures. For example, Mansfield et al. (2020) identified , “conceptual confusion, methodological inconsistency and a lack of measures developed and psychometrically tested. (p. 11)” for the construct of mental health literacy.

2. Why and how do you select your measure?

The next question is why researchers select one or more specific measures to assess a construct. This is important because there are usually multiple methods and potential instruments to choose from. For instance, there are at least 280 scales for measuring depression (Santor et al., 2006) and at least 65 different scales to measure emotions with 19 devoted to specifically to anger (Weidman et al., 2016). Researchers should use theoretical, empirical, and practical evidence to transparently justify their scale selection, ideally prior to data analysis.

To answer this question, we suggest researchers consider how their theoretical definition of the construct aligns with potential measures and the existing validity evidence for those measures. First, researchers should consider the content and face validity of an item or stimulus: does the instrument appear, at face, to measure what you are studying? Having answered question 1 above, ‘what is your construct?’ provides necessary guidance when

selecting measures. Second, after identifying measures that appear to capture the construct as predefined, researchers should then consider the measure's published validity evidence and if that validity evidence can extend to the current population and context. In case there are prior construct validity and instrument development studies, they should be cited and summarized in the paper. If construct validity evidence is gathered from a new or unpublished study, it should be described in detail. Reviewers of the study may agree or disagree about the measure's content validity, psychometric quality, or predictive validity, but citing and describing all validity evidence, previous or current, *enables* reviewers to do so.

Transparency regarding the (theoretical and operational) definition of the construct—as well as regarding the motivation for selecting a measure—helps prevent thinking two instruments measure the same construct because they have similar names (i.e., the *jingle fallacy*), or assuming two measures assess different constructs because they have different names (i.e., the *jangle fallacy*). The profusion of instruments and lack of transparency regarding them contributes to widespread jingle-jangle and makes avoiding it difficult. Transparency helps us spot *jingle-jangle*, but the meta-scientific research focusing on it shines light on systemic measurement practices that introduce serious validity concerns.

There is clear evidence of *jingle* in the literature, highlighting the need to address questions of what constructs are and how they are measured at the individual study level. To list just a few: depression scales (Fried, 2017), emotion measurement (Weidman et al., 2016), behavioral tasks vs. self-reports of self-regulation (Eisenberg et al., 2019), theory of mind (Warnell & Redcay, 2019), and fear extinction retention (Lonsdorf et al., 2019) each measure different content and phenomena under the same construct name. The on-going debate as to whether

measures of grit (perseverance and passion for long-term goals) and conscientiousness capture the same construct, given their correlation of 0.84 in a recent large meta-analysis (Credé et al., 2017), is a case of potential *jangle*. These fallacies are not without consequence: the relationship between loneliness and extraversion has been shown to be moderated by the particular scales used, with correlation coefficients ranging from $r = -0.44$ to $r = -0.19$ depending on the scale used (Buecker et al., 2019). A lack of clarity in what these scales are measuring, despite their names, muddies the interpretation of the relationship between constructs such as loneliness and extraversion. When such measurement questions are neglected, evaluating validity can be impossible.

3. What measure(s) did you use?

Once constructs are defined and measures selected, we must transparently report the aspects of the measure used. This requires describing each measure in detail, including where the instrument comes from, the exact number of items or stimuli, item wording, response format, whether it is the short version or long version, whether it is the English or Russian version, and how the instrument or task was presented to participants. For some types of measures (e.g., reaction time or neuroimaging) specification of hardware and software employed is needed.

The Hamilton Rating Scale for Depression, for instance, exists in versions with 6, 17, 21, 24 and 29 items, and has been translated into numerous languages (Bagby et al., 2004). Without information on what instrument was used and how it was presented to participants, we cannot evaluate the validity of the study, and we will struggle to conduct meaningful replications. Many of the questions about the details of a study's measures can be answered if authors make

their materials publicly available, and some journals incentivize such transparency by awarding open materials badges (e.g., Eich, 2014). These details about the measures used in a study matter because differences in measure presentation and administration introduce methodological variability that can influence results across a literature. Steinberg (1994) notes that the generalizability of personality assessment is limited by variability in administration context and order of items. In another example, the number of response options can systematically influence the variance of the response, confounding the results of statistical tests (Dawes, 2008).

It is common practice for measurement details to go unreported in the published literature. Of 433 scales used in studies published in *JPSP*, 40% had no accompanying information regarding their source, 19% were reported without the number of items, and 9% without the response scale (Flake, et al., 2017). This ambiguity contributes to a lack of continuity in the published literature, because studies that appear to be using the same scale could differ in item wording, response format, or number of items, raising questions about the validity of the study.

4. How do you quantify your measure?

Next, researchers usually score their instrument to produce one or more final numbers per person. There are at least three avenues that require transparent decisions: transformations of the responses; which items or stimuli form which scores; and how each score is calculated.

Transformations of response scales can be appropriate for various reasons, such as reverse scoring a negatively worded item, or collapsing categories with sparse responding. But researchers need to disclose all such transformations and justify them in the paper, such that they could be reproduced by readers. A next common set of questions is: will certain items or

stimuli be removed and not go into the final score; will items or stimuli be grouped, and if so, how? Developing a priori decision rules for removing items or stimuli and/or grouping them is ideal. Reporting those decision rules in the paper facilitates others being able to reproduce and evaluate the work. Finally, there are numerous ways to go from responses to scores in a dataset: averaging items, calculating a standardized score, estimating a factor or component score, or grouping study participants into categories. In sum, we need to know *how* the score(s) were calculated (transparency), and *why* they were calculated in that manner (justification). Note that inherent flexibility in quantifying raw data is not isolated to surveys and abounds in others approaches to measurement such as neuroimaging (Carp, 2012).

In the absence of a clear set of a priori rules (which can be documented in a preregistration), researchers might be tempted to wander down the multiverse of different forking paths presented during scoring (item average vs. factor score, 2 vs. 3 subscales), which can lead to different outcomes. But if researchers report ample evidence that supports the decisions made regarding scoring, we can rule out threats to validity and be more confident that measurement flexibility was not exploited to obtain the final results. Examples of transparency are preregistering all analyses, and sharing the analytic code, preferably using free open-source software such as R (R Core Team, 2019). Examples of justifications are citing a validation study that shows empirical evidence for a certain scoring approach and conducting a psychometric evaluation from the data collected in the current study. The psychometric models that can be used to make these and related decisions are out of the scope of this paper; see <https://osf.io/zrkd4/> for a resource list including psychometric models.

These are not hypothetical scenarios that lack consequence, but are at the core of what many consider central to the notion of questionable practice and *p*-hacking. Simmons et al. (Simmons et al., 2011) demonstrated that when researchers have two instead of one items at their disposal, the measurement flexibility can increase the Type I error rate considerably. In analyses of clinical trials published in psychology journals, about 1 in 3 trial registrations did not contain specific information about measurement (Cybulski et al., 2016), threatening the statistical conclusion validity of clinical interventions through measurement flexibility. Further, reviews of preregistered protocols in clinical trials have revealed that outcome switching is pervasive (Chan et al., 2004), and Ramagopalan et al. (2018) found that nearly a third of 89,204 registered studies had their primary outcome changed after finishing the registration. The methodological and meta-science literature is clear: measurement flexibility is used to misrepresent and mislead.

5. Do you modify the scale? And if so, how and why?

Next, we want to know if researchers modified their measures. Using an existing measure reduces researcher degrees of freedom because it is likely that many decisions such as scoring rules are already made. In the best case, such decisions have been justified in previous validation studies, which also support the validity of the measure and overall study conclusions. For example, the response scale and a mapping of items to subscales is likely outlined in a scale development study. However, the use of an existing measure does not eliminate flexibility and threats to validity, as it is common practice to modify measures (Flake et al., 2017; Weidman et al., 2016). Transparency regarding measure modification is not only an issue restricted to

survey measures: results from the competitive reaction time task, for instance, have been quantified in over 150 different ways across 130 publications (Elson, 2016; Elson et al., 2014).

Measures can be modified in a myriad of ways, before and after data collection, and these modifications challenge continuity in the published literature. Before data collection, potential degrees of freedom are a) changing the response type (e.g. 1-7 Likert to 1-100 visual analogue), b) changing the response style or options (1-7 Likert to 1-5 Likert), c) including additional items, d) dropping items, and e) changing item wording or content (in the case of tasks: stimuli, presentation time, etc.). After data collection, in addition to b) recoding items and d) dropping items, researchers can also f) adapt the scoring (e.g., which items map to which factors, conducting single-item analyses instead of averaging items, type of score calculated). Again, the main point is not that such modifications are never called for, but that if they remain undeclared and unjustified, they threaten the validity of the inferences we can draw from a given study.

Modifications offer large sources of flexibility which introduce uncertainty about the construct validity evidence for the interpretation of the scores. The meta-science literature suggests questionable measurement modification, for example, Flake et al., (2017) noted, in their review of 500 measures used in articles published in *JPSP*, 18 measures were the result of combining two separate measures into one on the basis of the reliability coefficient α . All modifications, either before or after data collection, need to be reported transparently, justified with all evidence available to the researcher, and ideally decided on in the study planning process. If researchers follow all these steps, it allows us to evaluate the validity of the study comprehensively and rule out alternative explanations for the final result.

6. Do you create a measure on the fly?

Sometimes, there are no scales available to study a construct of interest, or there may be good reasons to not use existing scales. When using existing scales, many decisions have already been made, such as the response format, the number of items or stimuli, item wording or stimuli presentation, scoring instructions, and clearly defined subscales. Using a measure exactly as described elsewhere, and providing a citation to the original source, decreases ambiguity and the potential for flexibility.

When researchers create a new measure, they need to address the 5 main questions we detailed above, and justify decisions with as much evidence as possible. Without transparent justification for how a scale is created and used, many questions regarding the validity of the study remain unanswered. Further, without guidance from the published literature as to how to answer the above questions, the researcher could be tempted to wander down many paths in the measurement garden, which is unlikely to result in a valid interpretation of the score meant to measure the construct, and therefore, unlikely to lead to a valid study conclusion. Of course, the more validity evidence and reasoned justification for creating and using a scale, the more threats to validity readers can rule out. But, at a minimum, researchers should disclose why they created the measure, and justify why they use an ad-hoc measure over using an existing measure in case such a measure exists. Further, if there is no or little validity evidence for that measure, researchers should discuss it as a limitation of the study.

A clear pattern is emerging from the meta-scientific measurement literature: it is common practice to create and use measures with no evidence of systematic development. Emotions, 69% (Weidman et al., 2016), education and behavior, 40-90% (Barry, Chaney, Piazza-Gardner,

et al., 2014), social and personality, 40% (Flake et al., 2017). Most recently, Shaw et al. (in-press) reviewed all of the measures used in the Many Labs 2 studies and of the 43 item-based scales they reviewed, 34 (79%) appeared to be ad-hoc and were used with no other supporting validity information in the original study. If a study uses undeveloped measures with no validity evidence, we as a community, have a lot of unanswered questions and serious doubts about the validity of the study conclusions.

Summary

In this paper, we defined QMPs as decisions researchers make that raise doubts about the validity of measure use in a study and ultimately the final conclusion. We explain that QMPs can threaten all aspects of a study's validity (internal, external, statistical, and construct) and focus on transparency of measurement as a first, necessary step to improving measurement practices. A lack of transparency makes it impossible for the scientific community to identify potential threats to the validity of a study's conclusions. We provide examples from the published literature that demonstrate the ubiquity of QMPs, and show that QMPs promote researcher degrees of freedom, and threaten the validity and replicability of psychological science. We list a set of questions that researchers, reviewers, and readers of scientific work can consider when planning, preregistering, conducting, or consuming research.

Answering these questions transparently facilitates the rigorous evaluation of the validity of research and enables meaningful replication studies. When such transparency is absent, we are left wondering what a significant or non-significant effect or replication or non-replication of an effect means. After all, measurement is the foundation on which all empirical science rests, and if important questions about the measurement are left unanswered, there is little that can be

concluded from a study. In addition to promoting transparency, researchers may find that engaging with these questions seriously throughout the research process inspires new lines of inquiry, results in research that strengthens the validity of their measures, ultimately enhancing the quality of their work.

The increased awareness and emphasis on QRPs, such as *p*-hacking, have been an important contribution to improving psychological science. We echo those concerns, but also see a grave need for broadening our scrutiny of current practices to include measurement (Fried & Flake, 2018). Recalling our example of depression at the outset, even if we increase the sample size of our depression trials, adequately power our studies, pre-register our analytic strategies, and stop *p*-hacking — we can still be left wondering if we were ever measuring depression at all.

References

- Aguinis, H., Ramani, R. S., & Alabduljader, N. (2018). What you see is what you get? Enhancing methodological transparency in management research. *Academy of Management Annals*, 12(1), 83–110. <https://doi.org/10.5465/annals.2016.0011>
- Antonakis, J., Bastardoz, N., Jacquart, P., & Shamir, B. (2016). Charisma: An ill-defined and ill-measured gift. *Annual Review of Organizational Psychology and Organizational Behavior*, 3(1), 293–319. <https://doi.org/10.1146/annurev-orgpsych-041015-062305>
- Bagby, R. M., Ryder, A. G., Schuller, D. R., & Marshall, M. B. (2004). Reviews and overviews The Hamilton Depression Rating Scale: Has the gold standard become a lead weight? *American Journal of Psychiatry*, 161(12), 2163–2177. <https://doi.org/https://doi.org/10.1176/appi.ajp.161.12.2163>
- Banks, G. C., Boyle, E. H. O., Pollack, J. M., White, C. D., Batchelor, J. H., Whelpley, C. E., Bennett, A. A., & Adkins, C. L. (2016). Questions about questionable research practices in the field of management: A guest commentary. *Journal of Management*, 42(1), 5–20. <https://doi.org/10.1177/0149206315619011>
- Barry, A. E., Chaney, B. H., Piazza-gardner, A. K., & Chavarria, E. A. (2014). *Validity and Reliability Reporting Practices in the Field of Health Education and Behavior : A Review of Seven Journals*. <https://doi.org/10.1177/1090198113483139>
- Barry, A. E., Chaney, B., Piazza-Gardner, A. K., & Chavarria, E. A. (2014). Validity and reliability reporting practices in the field of health education and behavior. *Health Education & Behavior*, 41(1), 12–18. <https://doi.org/10.1177/1090198113483139>

- Buecker, S., Maes, M., Denissen, J., & Luhmann, M. (2019). *Loneliness and the Big Five Personality Traits: A Meta-Analysis*. <https://doi.org/10.22201/fq.18708404e.2004.3.66178>
- Carp, J. (2012). On the plurality of (methodological) worlds: Estimating the analytic flexibility of fMRI experiments. *Frontiers in Neuroscience*, 6(October), 1–13.
<https://doi.org/10.3389/fnins.2012.00149>
- Chan, A.-W., Hrobjartsson, A., Haahr, M. T., Gøtzsche, P. C., & Altman, D. G. (2004). Empirical Evidence for Selective Reporting. *Journal of the American Medical Association*, 291(20), 2457–2465. <https://doi.org/10.1001/jama.291.20.2457>
- Credé, M., Tynan, M. C., & Harms, P. D. (2017). Much ado about grit: A meta-analytic synthesis of the grit literature. *Journal of Personality and Social Psychology*, 113(3), 492–511.
- Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52(4), 281–302. <https://doi.org/10.1037/h0040957>
- Cybulski, L., Mayo-wilson, E., Grant, S., Corporation, R., & Monica, S. (2016). Improving Transparency and Reproducibility Through Registration : The Status of Intervention Trials Published in Clinical Psychology Journals. *Journal of Consulting and Clinical Psychology*, 84(9), 753–767.
- Dai, D. Y., & Chen, F. (2013). Three paradigms of gifted education: In search of conceptual clarity in research and practice. *Gifted Child Quarterly*, 57(3), 151–168.
<https://doi.org/10.1177/0016986213490020>
- Dawes, J. (2008). Do data characteristics change according to the number of scale points Used?

- An experiment using 5-Point, 7-Point and 10-point scales. *International Journal of Market Research*, 50(1), 61–104. <https://doi.org/10.1177/147078530805000106>
- Eich, E. (2014). Business not as usual. *Psychological Science*, 25(1), 3–6.
<https://doi.org/10.1177/0956797613512465>
- Eisenberg, I. W., Bissett, P. G., Zeynep Enkavi, A., Li, J., MacKinnon, D. P., Marsch, L. A., & Poldrack, R. A. (2019). Uncovering the structure of self-regulation through data-driven ontology discovery. *Nature Communications*, 10(1). <https://doi.org/10.1038/s41467-019-10301-1>
- Elson, M. (2016). *Flexiblemeasures.com: Competitive reaction time task*.
<https://doi.org/10.17605/osf.io/4g7fv>
- Elson, M., Mohseni, R. M., Breuer, J., Scharnow, M., & Quandt, T. (2014). Press CRTT to measure aggressive behavior: The unstandardized use of the competitive reaction time task in aggression research. *Psychological Assessment*, 26(2), 419–432.
<https://doi.org/10.1037/a0035569>
- Fanelli, D. (2013). Redefine misconduct as distorted reporting. *Nature*, 494(7436), 149.
<https://doi.org/10.1038/494149a>
- 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>
- Fried, E. I. (2015). Problematic assumptions have slowed down depression research: why

symptoms, not syndromes are the way forward. *Frontiers in Psychology*, 6(306), 1–11.

<https://doi.org/10.3389/fpsyg.2015.00309>

Fried, E. I. (2017). The 52 symptoms of major depression: Lack of content overlap among seven common depression scales. *Journal of Affective Disorders*, 208, 191–197.

<https://doi.org/10.1016/j.jad.2016.10.019>

Fried, E. I., & Flake, J. K. (2018, March). Measurement matters. *Observer*.

<https://www.psychologicalscience.org/observer/measurement-matters>

Gelman, A., & Loken, E. (2013). *The garden of forking paths: Why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time* (pp. 1–17). Department of Statistics, Columbia University.

<https://doi.org/dx.doi.org/10.1037/a0037714>

Hussey, I., & Hughes, S. (2019). Hidden invalidity among fifteen commonly used measures in social and personality psychology. *Advances in Methods and Practices in Psychological Science*, 1–27. <https://doi.org/10.31234/osf.io/7rbfp>

John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23(5), 524–532.

<https://doi.org/10.1177/0956797611430953>

LeBel, E. P., Campbell, L., & Loving, T. J. (2017). Benefits of open and high-powered research outweigh costs. *Journal of Personality and Social Psychology*, 113(2), 230–243.

<https://doi.org/10.1037/pspi0000049>

- Lonsdorf, T. B., Merz, C. J., & Fullana, M. A. (2019). Fear Extinction Retention: Is It What We Think It Is? *Biological Psychiatry*, 85(12), 1074–1082.
<https://doi.org/10.1016/j.biopsych.2019.02.011>
- Mansfield, R., Patalay, P., & Humphrey, N. (2020). A systematic literature review of existing conceptualisation and measurement of mental health literacy in adolescent research : current challenges and inconsistencies. *BMC Public Health*, 20(607), 1–14.
- Nosek, B. A., & Bar-Anan, Y. (2012). Scientific utopia: I. Opening scientific communication. *Psychological Inquiry*, 23(3), 217–243. <https://doi.org/10.1080/1047840X.2012.692215>
- Nosek, B. A., Beck, E. D., Campbell, L., Flake, J. K., Hardwicke, T. E., Mellor, D. T., van 't Veer, A., & Vazire, S. (2019). Preregistration is hard, and worthwhile. *Trends in Cognitive Sciences*, 23(10), 815–818. <https://doi.org/https://doi.org/10.1016/j.tics.2019.07.009>
- Orben, A., & Przybylski, A. K. (2019). The association between adolescent well-being and digital technology use. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-018-0506-1>
- Parker, G. (2005). Beyond major depression. *Psychological Medicine*, 35(4), 467–474.
- R Core Team. (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Robinaugh, D. J., Haslbeck, J. M. B., Waldorp, L. J., Kossakowski, J. J., Fried, E. I., Millner, A. J., McNally, R. J., van Nes, E. H., Scheffer, M., Kendler, K. S., & Borsboom, D. (2019). *Advancing the Network Theory of Mental Disorders: A Computational Model of Panic Disorder*. 1–76. <https://doi.org/10.31234/osf.io/km37w>

Santor, D. A., Gregus, M., & Welch, A. (2006). Eight Decades of Measurement in Depression. *Measurement*, 4(3), 135–155.

https://doi.org/http://www.tandfonline.com/doi/abs/10.1207/s15366359mea0403_1

Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference* (2nd ed.). Wadsworth Publishing.

Shaw, M., Cloos, L. J. R., Luong, R., Elbaz, S., & Flake, J. K. (n.d.). Measurement practices in large-scale replications : Insights from Many Labs 2. *Canadian Psychology*.

Sijtsma, K. (2016). Playing with data-Or how to discourage questionable research practices and stimulate researchers to do things right. *Psychometrika*, 81(1), 1–15.

<https://doi.org/10.1007/s11336-015-9446-0>

Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., Bahnik, S., Bai, F., & Bannard, C. (2018). Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science*, 1(3), 337–356. <https://doi.org/10.1177/2515245917747646>

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant.

Psychological Science, 22(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>

Skinner, E. A. (1996). A guide to constructs of control. *Journal of Personality and Social Psychology*, 71(3), 549–570. <https://doi.org/10.1037//0022-3514.71.3.549>

Steege, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency

through a multiverse analysis. *Perspectives on Psychological Science*, 11(5), 702–712.

<https://doi.org/10.1177/1745691616658637>

Steinberg, L. (1994). Contexts and serial-order effects in personality measurement: limits on the generality of measuring changes the measure. *Journal of Personality and Social Psychology*, 66(2), 341–349.

Warnell, K. R., & Redcay, E. (2019). Minimal coherence among varied theory of mind measures in childhood and adulthood. *Cognition*, 191(December 2017), 103997.

<https://doi.org/10.1016/j.cognition.2019.06.009>

Weidman, A. C., Steckler, C. M., & Tracy, J. L. (2016). The jingle and jangle of emotion assessment: Imprecise measurement, casual scale usage, and conceptual fuzziness in emotion research. *Emotion*.

Wigboldus, D. H. J., & Dotsch, R. (2016). Encourage Playing with Data and Discourage Questionable Reporting Practices. *Psychometrika*, 81(1), 27–32.

<https://doi.org/10.1007/s11336-015-9445-1>