Strengthening the foundation of educational psychology by integrating construct validation into open science reform

Jessica Kay Flake

Department of Psychology, McGill University

This paper was accepted for publication in *Educational Psychologist* volume 56, issue 2; scheduled for April 2021, as a part of a special issue titled, “Educational psychology in the open science era.”

This preprint has been peer reviewed, but not copy edited by the journal and may differ from the final published version. The DOI of the final published version is: 10.1080/00461520.2021.1898962. Once the article is published online, it will be available at the following permanent link: <https://doi.org/10.1080/00461520.2021.1898962>

Abstract

An increased focus on transparency and replication in science has stimulated reform in research practices and dissemination. As a result, the research culture is changing: the use of preregistration is on the rise, access to data and materials is increasing, and large-scale replication studies are more common. In this paper, I discuss two problems the methodological reform movement is now ready to tackle given the progress thus far and how educational psychology is particularly well suited to contribute. The first problem is that there is a lack of transparency and rigor in measurement development and use. The second problem is caused by the first; replication research is difficult and potentially futile as long as the first problem persists. I describe how to expand transparent practices into measure use and how construct validation can be implemented to bolster the validity of replication studies.

Introduction

Failed replications and pervasive misleading research practices have incited a crisis of confidence in science. Across many disciplines, researchers are questioning the validity of the conclusions of scientific research (Baker, 2016) and the psychological sciences are no exception. In fact, the larger discipline of psychology has fueled debates about a replication crisis (Gilbert et al., 2016; Open Science Collaboration, 2015), spurred developments in how to address it (Nosek et al., 2018), and is years into a period of self-reflection. This self-reflection has welcomed meta-scientific research, methodological innovation, and reformation of research practices.

While reforms are well underway in some areas of psychology, they are also on-going, revealing the next wave of methodological improvements psychologists need to make. As educational psychologists reform their practices, they can benefit from the lessons that have been learned, the discovery of new challenges, and the refinement of solutions. In this paper, I discuss two problems the methodological reform movement is now ready to tackle and how educational psychology is particularly well suited to contribute. The first problem is that there is a lack of transparency and rigor in measurement development and use. The second problem is caused by the first; replication research is difficult and potentially futile as long as the first problem persists.

First, a brief review the methodological reform and open science movement is provided, focusing on how the stage is set to move forward with reforming measurement practices. Then, I explain the how and why of the two connected problems: (1) transparency of measure development and use and (2) the utility and impact of replication research. I will end by describing solutions to these problems that educational psychologists are well suited to implement, as well as concrete next steps.

Methodological Reform and Open Science Movement

The methodological reform and open science movement is comprised of a large group of scientists at all career stages who share the desire to improve the validity of scientific output and they recognize that current practices need to change to meet that end. Though the movement is comprised of many different people working along different avenues of improvement, two core values are central: researchers need to be more transparent in the research process and they need to conduct more direct replications of findings. Articles in this special issue focus on the merit, conduct, and challenges of these central values thoroughly: open practices (Fleming et al., 2021), preregistration and registered reports (Reich, 2021), and replication (Plucker & Makel, 2021). My goal here is not to reiterate those points but to use them as a springboard to discuss progress made thus far and progress needed moving forward, particularly as it relates to measurement practices.

Increasing Analytic Transparency and Direct Replication Research

Transparency is a core part of the methodological reform and open science movement because it supports the rigorous conduct of science by allowing the community to thoroughly evaluate, challenge, reuse, and/or adapt existing scientific outputs to accumulate reliable knowledge (Nosek & Bar-Anan, 2012). Progress towards a more transparent psychology has been made in the past decade and is apparent in advances to infrastructure, changing policies at journals, and researchers changing how they work (Mellor, 2021/this issue). The infrastructure to support transparency in science is growing with the addition of multiple cloud-based options for sharing data, materials, and preregistrations. Just a few examples are: the Open Science Framework (<https://osf.io/>), Research Box with AsPredicted (<https://researchbox.org/>), and for the education researcher, the Society for Research on Educational Effectiveness and the University of Michigan's Inter-university Consortium for Political and Social Research (<https://sreereg.icpsr.umich.edu/sreereg/>). In addition to infrastructure, leading journals are acknowledging, supporting, and incentivising researchers to be more transparent. Many journals offer badges to acknowledge open practices, which are visible on a published paper and

indicate if researchers have made their data, materials, and/or a preregistration publicly available (Kidwell et al., 2016). Another step toward transparency in publishing is a new publication format in which authors can submit a paper before the data are collected and the paper is reviewed based on the theory and proposed methods and analysis (a registered report; Chambers et al., 2014). As of 2020, American Psychological Association journals, including the *Journal of Educational Psychology* offer badges and the registered report publication option, demonstrating educational psychology's progress toward transparency. With infrastructure and the support of journals, researchers are changing their practices. Use of sites like the Open Science Framework is growing exponentially and more researchers are preregistering their work each year (Nosek et al., 2018).

Transparent research practices dovetail with an increased focus on replication research. I will not present the debated value of replication research, which is discussed in depth by Plucker and Makel (2021) in this special issue. What is important is that a core principle of the open science and reform movement has been the conduct and support of replication research that is as close to the original study as possible, often referred to as a "direct" replication (Simons, 2014). The open science movement aims to increase the value placed on these direct replication studies which focus specifically on finding the same result, using the same method, as an original study. In psychology broadly, researchers have come together, often on a larger scale than we have seen in the past, to execute such replication studies that promote the related value of transparency. For example, the Many Labs projects involve hundreds of researchers working together to conduct direct replications of a set of studies that result in larger, pooled datasets than is typical and those datasets and associated materials are open for anyone to use. These projects are becoming a norm, at the time of writing ManyLabs has released their 5th set of studies (Ebersole et al., 2020). In another example, the Psychological Science Accelerator is a distributed laboratory network (Moshontz et al., 2018) that carries out large-scale studies with multisite data collection labs, publishes open data, and routinely conducts direct replication studies. These large-scale,

transparent, direct replication studies are getting journal support. Preregistered replication reports can be submitted to the Association for Psychological Science journal *Advances in Methods and Practices in Psychological Science*, which publishes research relevant to all areas of psychology (Simons, 2018).

Just as the broader discipline of psychology is increasing emphasis on replication research, education researchers are well positioned to do the same with existing supports. For example, there are specific funding streams to support multisite trials of educational research through the Institute for Education Sciences and education researchers have frameworks for thinking through the methodological complexities of such studies. For example, educational researchers have contributed to understanding issues of implementation fidelity and treatment contamination in large-scale educational experiments (O'Donnell, 2008; Rhoads, 2011). As such, educational psychologists have valuable insights to contribute to the methodology of large-scale collaborative replication studies.

Extending Analytic Transparency and Direct Replication to Measurement

Many of the reforms undertaken in psychology to increase transparency, including the ones I have just discussed, focus on statistics and data analysis practices (see Shrout & Rodgers, 2018 for a review). This is because there is an abundance of analytic flexibility in the data processing and analysis part of the research pipeline and that flexibility makes it possible for researchers to engage in practices that produce a desired and misleading result in support of their hypothesis. Some terms of the last decade for these practices are *p*-hacking, questionable research practices, researcher degrees of freedom, and exploring the garden of forking paths (Gelman & Loken, 2014; John et al., 2012; Simmons et al., 2011; Wicherts et al., 2016). The purpose of preregistration and registered reports is to prevent these practices, i.e., they limit the opportunities researchers have to engage in undisclosed analytical flexibility which might produce misleading results. When researchers transparently report all decisions and actions taken in the analysis process, that information enables the community to accurately

interpret statistical tests of key effects and calibrate claims according to the evidence (Nosek et al., 2019).

As researchers share more data, write more preregistrations, carry out more replication studies and take a closer look at existing research in the process, they learn how progress is being made, and where progress is needed. Transparency and rigor are

critical throughout the entire research process and necessary for meaningful replication studies, but the emphasis thus far has been on statistics. Figure 1 rearranges

aspects of the research pipeline into a hierarchy with the conclusion at the apex, supported by evidence from the research process. The philosophy and theoretical

approaches to research, the measurement of

hypothesized phenomena, and the design of studies are at

the foundation of the pyramid because decisions regarding those aspects of the research have serious repercussions at the higher levels of the process. Problems at the foundation of the pyramid can interact and cooccur: experimental manipulations can fail (research design), a survey can produce invalid scores for a key outcome (measurement), or a theoretically critical confound can be omitted (substantive theory). That lack of strong evidence travels up to the conclusion and it can not be assuaged by a confirmatory, preregistered statistical analysis.

Figure 1. Hierarchy of research processes that support valid study conclusions (Flake et al., 2020).

Expanding and developing transparent practices and replication research for the measurement aspects of the study pipeline is a foundational part of strengthening the evidence that supports the conclusion. Recent meta-science research suggests that transparency in the measurement process is the next problem the open science and methodological reform movement needs to address (Flake & Fried, 2020; Slaney, 2017). Educational psychology has a rich history of developing methods and practices for

measurement and is thus well positioned to implement transparent practices and conduct replication studies that promote rigorous measurement and integrate those practices into the larger framework of open science.

The First Problem: Transparency of Measure Development and Use

A part of methodological reform is uncovering problematic practices through reflection, meta-science, and simulation. If the community knows what the problems are, then they can develop methods for fixing them. One such problem, highlighted by the methodological reform movement, is the lack of transparency and rigor in measure development and use. In this section I will explain how the research community knows this is a problem and why it is a problem. Transparency and measurement are important to consider in both qualitative and quantitative research. Though this is written from a quantitative perspective, the points here could apply to either. Developing transparent practices for measurement in qualitative research is an important area of study. Next, I'll review the basics of construct validity theory to provide the necessary background to the problem (for a thorough reading list see: <https://osf.io/zrkd4/>, Fried & Flake, 2018).

Construct Validity: A Primer

Measurement is not a straightforward part of the research process. It requires making many consequential decisions early in the research pipeline about what to measure, how to measure, and how to use quantitative methods to incorporate those measurements into a downstream analysis. Often, what educational psychologists want to study is not directly observable. When I see 400 faces looking at me in my introductory statistics course, I can't readily tell which students are motivated, prepared, or struggling. I ask my students to fill out a short survey to gauge their motivation and I give them homework assignments to judge if they are mastering the material. I trust that students with higher scores on the motivation measure are more motivated and that those who perform better on the

homework assignments have in fact mastered more of the content than those who are performing worse. But I cannot just trust in these numbers, I require evidence to support their interpretation. Construct validation is the process by which researchers gather evidence to support the use and interpretation of scores (AERA, NCME, APA, 2014).

The Standards for Educational and Psychological Testing detail an in-depth process of construct validation, encompassing theoretical breadth and depth, as well as methodological complexity. In short, researchers should integrate evidence from multiple sources regarding a measure's content, psychometric properties, and how scores on the measure relate to outcomes and measures of other relevant constructs. The potential studies and approaches for addressing any of these aspects of the construct validation process are vast. An example of evidence regarding content is having experts review the item for content clarity and relevance (see Gehlbach & Brinkworth, 2011, for a thorough example). Factor and reliability analyses are common approaches to assessing psychometric quality (Zumbo & Chan, 2014), and quantifying a measure's predictive ability for a key outcome or association with a theoretically relevant construct can provide further evidence. A rigorous initial validation study for the development of a new measure would ideally incorporate multiple studies addressing these various aspects of construct validation by utilizing a collection of methodologies.

But the work doesn't stop there; construct validation is an on-going process. When I provide strong evidence that the motivation measure captures students' motivation in my university level, primarily English speaking statistics course, that evidence might not extend to the context of my colleagues in the Netherlands using it five years from now to measure students' motivation in their middle school science classrooms. As a result, construct validation encompasses two broad processes – integrating evidence to support the interpretation of scores when developing a measure *and* incorporating new evidence to support the valid use of a measure for a given context, population, or time (Messick, 1995). In addition to *The Standards*, there are hundreds of articles and dozens of text

references on how to conduct construct validation and provide strong evidence in support of the interpretation and use of a measure. The problem is not that we lack the methodologies, tools, or theory in the educational measurement community to develop and use measures validly. The problem is that there is a lack of transparency in the measurement part of the research process and reforms to increase transparency need to be expanded into the measurement process. This will enable the community to figure out which practices are problematic and develop or utilize methods to improve practices.

How do researchers know there is a problem with transparency in measurement?

Measures are developed and used without the reporting of strong validity evidence and this lack of transparency directly threatens the validity of research conclusions. A growing body of meta-scientific work on measurement practices in psychology more broadly (see Slaney, 2017 for a review) reveals problematic practices that span from aspects of the development process (e.g., conceptualization and definition) through to aspects of on-going validation for measure use (e.g., modification of existing and rigorously studied instruments).

Recent metascientific studies demonstrate that it is common for researchers to use scales without explicit reporting of **any** validity evidence. Reviews in areas adjacent to educational psychology suggest a systemic lack of transparency: social psychology (Flake et al., 2017), health education (Barry et al., 2014), emotions (Weidman et al., 2016), and American Psychological Association journals more broadly (Qualls & Moss, 1996). These studies reported that nearly half or more of scales are used without any reported validity evidence. A non-trivial percentage of scales appear to have been made up *on-the-fly* for use in that particular study and it is not clear if the measure was evaluated in any way before use. Further, unjustified measurement flexibility and modification of existing measures are common practices, with estimates ranging from researchers adapting nearly 10% to 30% of scales and

without any accompanying validity evidence (Barry et al., 2014; Flake et al., 2017; Weidman et al., 2016). This “casual scale usage” (Weidman et al., 2016) is not substantively or methodologically rigorous. This metascientific research is a guide to where reforms to increase transparency are needed next: the measurement process.

Why is a lack of transparency a problem?

A lack of transparency means that researchers do not report validity evidence supporting a measure, and this makes it impossible to evaluate the validity of the measure’s interpretation and use. The consequences are both theoretical and methodological. The theoretical consequence of ill-defined constructs and their willy-nilly measures are ubiquitous jingle-jangle fallacies. *Jingle* refers to erroneously assuming two measures with the same name are measuring the same construct (Thorndike, 1904) and *jangle* is the erroneous assumption that two measures assess different constructs because they have different names (Kelley, 1927). There is no shortage of these issues in research relevant to educational psychology (Murphy & Alexander, 2000). For example, there is an ongoing debate regarding if grit is a different construct from conscientiousness given how strongly they correlate (a jangle, Credé et al., 2017). In another example, Flake et al. (2015) , identified numerous scales to measure students’ psychological “cost”, and though these scales were reported all to measure “cost” item content differed substantially as well as researchers’ definitions of the construct. Some areas of research suffer from both problems at once because researchers sometimes use the same name to refer to different constructs and other times use different names to refer to the same constructs (see Christenson et al., 2012 for a review related to student engagement; for a review related to self-concept Marsh et al., 2019). In a study relevant to educational psychology, Orben and Przybylski (2019) investigated how using digital technology influences adolescents’ well-being, and note that the excess of well-being measures, how flexibly they are used to combine and break apart scales, makes the constructs seem more like an “accessory” to publication, not a theoretical foundation for further study. This

measurement morass challenges the theoretical continuity of our work and directly negates the reliable accumulation of knowledge about a construct. This problem in traditional empirical research compounds across the literature to limit the utility of studies to synthesize the literature, such as meta-analyses (see Patall, 2021/this issue for discussion on transparency and research synthesis).

In addition to the negative theoretical consequences of *jingle-jangle* fallacies and a lack of conceptual continuity in the measure development and use process, a lack of transparent validity evidence for measures facilitates methodological flexibility, which undermines validity. Educational and psychological measurement necessitates many decisions and without transparency, justification, and empirical evidence in support of those decisions, it is unclear if the numbers represent anything meaningful. A lack of reported validity evidence for a measure raises questions. Some examples from my own experience reviewing papers for educational psychology journals: Why were different measures swapped in and out across studies of the same construct? What motivated item removal and addition? Why did the authors score the scale in that way? With so many related, yet slightly different constructs, researchers have seemingly endless flexibility to combine, create, and remove variables to tell a story with their analysis that fits their hypothesis.

A lack of transparency in the reporting of decisions related to the measures in a study raises doubts about the validity of the measure use and ultimately the final conclusion of the study. A measure could have strong, yet unreported validity evidence **or** simply not produce valid scores or interpretations, the key issue is that readers do not know. Flake and Fried (2020) define such reporting practices for measures as Questionable Measurement Practices (QMPs). QMPs should be unacceptable in the published literature, as they prevent a thorough evaluation of the construct validation process. What's more, a lack of detailed information and justified evidence about our measures makes replicating our work nearly impossible and potentially futile. How can we replicate our findings if we don't even know what we are measuring?

The Second Problem: A Lack of Construct Validity Evidence Stymies Replication Research

Fundamentally science is a team endeavor. If, in my work I discover a principle, process, or phenomenon, I hope another scientist will come along and put that knowledge to the test via application, replication, and extension. Ideally, through multiple studies and perspectives scientists generate reliable knowledge. Direct replication is one of many avenues for generating that reliable knowledge and is a core part of the open science and methodological reform movement. Large-scale replication studies for which hundreds of researchers contribute, are becoming a norm. Replication studies are logistically complex and present methodological challenges, particularly measurement challenges. This section explains how the research community knows that a lack of transparency of measurement stymies replication studies and why that is a problem. But first, I provide some background regarding large scale replication studies in psychology, which have facilitated a thorough understanding of this problem.

Large Scale Replication Studies in Psychology

Psychologists have undertaken large-scale replication projects in which researchers from different institutions work together to collect data to replicate an original study as directly as possible. These projects provide insights into the benefits and challenges of conducting and interpreting large scale replication studies that only hindsight can afford. I'll briefly describe two such projects for context: The Reproducibility Project: Psychology (RPP: Open Science Collaboration, 2015) and the Many Labs 2 (ML2: Klein et al., 2014).

The RPP set out to replicate 100 studies published in *The Journal of Experimental Psychology: General* (JEP:G), *Psychological Science* (PS), and *The Journal of Personality and Social Psychology* (JPSP). The goal of the replicators was to adequately power the replications and conduct them as similarly to the original studies as possible. The findings from the RPP solidified a mounting concern about a

replication crisis and spurred many of the reforms I have discussed. A minority (36/100) of the replication studies demonstrated the key significant effect of the original at the .05 level of significance and only 44% of the replicated effects were within the 95% confidence interval of the original effect size (Open Science Collaboration, 2015). Additionally, 82 of the replicated effect sizes were smaller than those reported in the original study.

Similar to the RPP, the Many Labs coordinators selected a set of studies to replicate, but in contrast to the RPP in which one lab conducted one replication study each, the goal was to conduct replication studies with 'many labs' collecting data for the same study. The ML2 researchers conducted replication studies of 28 effects with laboratories running more than one study and many laboratories running the same study. As a result, 125 samples of data were collected by labs from all six populated continents. By design, the ML2 was a heterogeneous data collection in that samples were diverse, and studies were conducted in many different languages. This emphasis on getting larger and more representative samples for replication studies has continued and is a key feature of other large-scale replication studies (Moshontz et al., 2018).

How do researchers know that measurement in replication studies is a problem?

Replication research is challenging, particularly large-scale replication research, and only through doing it and then reflecting on how it was done, will the research community learn how to improve the process. Because these large-scale replication studies promote the value of transparency, the research community can review the procedures, materials, and results systematically to facilitate that learning. Here I'll discuss key conclusions from two systematic measurement reviews of the RPP (Flake et al., 2020) and ML2 (Shaw et al., 2020), which shared the goal of understanding how a lack of transparency and validity evidence for measures in original research influences the process of conducting and interpreting replication research.

The first conclusion about measurement practices that can be drawn from these large-scale replication studies is that item-based, self-report measures are a core part of the replication process because they are the most frequent approach to measurement. The RPP ultimately utilized 362 measures, and the ML2 60 measures. In both the RPP and ML2, item-based measures such as surveys and questionnaires, on which participants' responses to items were scored to represent a primary variable of interest, were the most common approach to measurement: RPP = 53% (193 item-based scales, 50% were a single item) , ML2 = 72% (43 item based scales, 77% of those were a single item). Measures such as reaction time, coding of interview data, and laboratory tasks were less common. Consistent with the meta-scientific research discussed above, in the original studies the item-based measures lacked transparent validity evidence. Only a minority of measures were reported with a source in the original studies, 29% in the RPP and 19% in the ML2. It was common for measures to be used without reporting any supporting reliability or validity information: no source, no reliability measures, no psychometric analyses. In the RPP, 20% of measures were used with no information and 30% in the ML2.

QMPs from original research transfer to the associated replication research. Replicators do not know if the measures have strong, yet unreported validity evidence **or** simply do not produce valid scores and interpretations. Before the replication study even begins there is a challenging set of questions. Was there strong validity evidence in the original study that just wasn't reported? If so, does that evidence generalize to the replication study? If not, what will the replication study results mean if the measures produce invalid scores? It was not common or standard practice for replication reports, written as a part of the RPP and ML2, to include thorough measure evaluation as a part of the replication process to address these questions. For the limited number of reports that included some information regarding the validity of measurement, replicators noted lower reliability coefficients and different factor structures than those reported in the original study.

The ML2 study, conducted with transparency as a central goal, is completely open access with adequate sample sizes for evaluating the psychometric properties of measures used in the replication studies. This transparency meets one of the goals of the open science movement, it allows for further scrutiny and evaluation of the study materials and results. Generally, the measures, which had limited validity evidence in the original studies, exhibited poor psychometric properties in the replication study (Shaw et al., 2020). Of the item-based measures that were appropriate for factor and reliability analyses (i.e., had enough items and were assumed unidimensional), none uniformly met model fit criteria often used to indicate a “good” fitting model, though given the controversial and arbitrary nature of fit statistics cut-offs, some scales may be considered to have acceptable fit of the assumed factor structure. Across all of the measures, the minimum reliability coefficient as measured by alpha was .42 and the maximum .86. There was substantial variability of reliability coefficients across data collection labs and translated versions of measures consistently had lower reliability than original versions. Poor model fit of assumed factor structure, low reliability, and variable reliability across scale versions and labs, all point to invalid score use and interpretation in the replication study. The lack of transparency and validity evidence for measures from the original studies carries over to the replication study.

Why is measurement in replication studies a problem?

Ideally a replication lends conclusive evidence about the strength of a theory or claim. If the replication fails to find the original effect, usually researchers desire a simple interpretation: theory failure – conclusive evidence that the original effect is not valid. In the context of a direct replication, the methods and procedures are as close to the original studies as possible to rule out confounding variables that might complicate a simple interpretation of the replication. However, a task that might seem straightforward at the face: running the same study in the same way to see if you get the same result, is complicated. There are many possible confounds that can prevent the desired interpretation. For example, unforeseen or unknown population, manipulation, and procedural differences can

introduce uncertainty in the interpretation of an effect observed in a replication study (Fabrigar & Wegener, 2016). In the context of a large-scale replications with many laboratories involved, implementation heterogeneity of procedures and inevitable cultural variability introduce further confounds (Klein et al., 2014).

A lack of validity evidence for measures in replication studies is one such confound that completely undermines the desired interpretation of a replicated effect. An implicit assumption of a replication study is that the measures represent the constructs of interest **and** do so in the same way as the original study (Fabrigar & Wegener, 2016; Flake et al., 2020). The RPP and ML2 demonstrate that it is common for researchers to begin a replication study without knowing if the measures represented the construct of interest in the original work. Then, even if the original research produced valid scores from measures, it is possible the validity evidence doesn't generalize to the replication context. The results is that the effect observed in the replication study has ambiguous meaning. For a concrete example: imagine a dependent variable measured with poor reliability. In the event of a failed replication, it is unclear if the result is due to theory failure or failure to detect a real effect in the face of poor reliability (see Loken & Gelman, 2017 for a detailed discussion). In the event of a replicated effect, the meaning of the numbers is unclear if they do not reliably represent the construct. What does it mean to replicate an effect when the factor structure of the scale in the replication sample conflicts with that observed in the original study, or if for some data collection sites the reliability is much lower than others? Questions like these are what methodologists, open science enthusiasts, and educational psychologists need to grapple with when they plan and interpret replication research – I offer some ideas in the solutions section below.

The potential and promise of large-scale replication research are too great to let measurement challenges impede impact. When scientists work together, they can achieve far more than toiling away, solo in their respective laboratories. The next step is to develop the transparent and rigorous

measurement practices for replication research that will enable the community to realize the potential. Educational psychologists are well suited to adopt and develop such practices because they are measurement savvy and experienced in considering the complications of replication research.

Solutions for Educational Psychologists to Consider

Many areas of science are in the process of reforming research practices to increase transparency with the goal of strengthening the validity of conclusions. Educational psychologists can move forward with integrating open science practices into the research culture as other articles in this special issue describe. The reform effort can be guided by learning from how reforms were implemented in other areas of psychology, from the meta-science conducted in adjacent areas, and from the replication research undertaken thus far. Work thus far has revealed two problems: (1) measurement practices lack transparency in many areas of psychology, which undermines rigor and (2) these practices make replication research more difficult and limit the usefulness of the findings.

Educational psychologists have shaped our notions of what construct validity is with Lee Cronbach's seminal work to define it (Cronbach & Meehl, 1955) and they continue to shape recommendations for the evidence needed to support the use of measures by codeveloping the latest edition of *The Standards for Educational and Psychological Testing* with the American Educational Research Association (AERA, NCME, APA, 2014). Educational psychologists are innovators of methods, pioneering in quantitative modeling and measurement. As educational psychologists incorporate transparency into their research practices and make direct replication a valued output, they can expand the scope of what the open science movement has accomplished so far and help to develop the methods needed that will address measurement challenges head on.

Solution One: Listen to the meta-science and do more of it

A part of improving science is reflecting on how it is done and identifying what needs to change. Such reflections in other areas have rang the alarm bells about warped incentive structures that encourage research practices that threaten the validity of our claims (Flake & Fried, 2020; John et al., 2012; Steegen et al., 2016). Questionable research practices are in the vernacular because scientists took a closer look, identified what was going on, and called it out. It is time for educational psychologists to do the same. Different disciplines will face different challenges and some readers may be thinking of poor research practices not yet mentioned or specific barriers to implementing reforms that haven't been addressed. Educational psychologists will certainly have questions about how open science practices and reform relate to the unique challenges they face studying some of the most vulnerable learners. In what ways does the research need to improve? How to adapt research practices to spur that improvement? These questions get answered through systematic and empirical evaluation of the research.

There is little empirical basis for knowing what can be improved in educational psychology; and no meta-scientific research that examines measurement practices in educational psychology specifically. It is unknown what bad practices are common, what innovations are under utilized, and what is needed next. If the findings from existing meta-science research on measurement practices apply to educational psychology, then transparency of measurement development and use is a problem. Educational psychologists can systematically review the published literature asking questions such as, "how often are the item stems and response options reported?," "what construct validation practices are most frequent for adapted measures?," or, "are measures used without the reporting of any validity evidence?" Meta-science is a step toward improving practices—through it, researchers can see where the cracks in the foundation are and work toward an empirically based improvement strategy. A step forward is to add to the methodological repertoire, meta-scientific research that focuses on the use of

methodologies and how to improve implementation. A part of getting this done is seeing the value in it. The other part of getting it done is supporting the conduct and publication of such research.

First, researchers need funding streams that directly support meta-science research. For example, the non-profit foundation the Campbell Collaboration funds systematic reviews of methodology and methodological research, but more opportunities are needed. Further, explicit statements of aims and scopes in journals to support meta-science, are another way to bolster this research. For example, the aims and scope for the journal *Advances in Methods and practices in Psychological Science* states, “AMPPS welcomes meta-science contributions that examine research practices in the field,” meta-science is also explicitly discussed in the editor’s statement about the launch of the journal (Simons, 2018).

Solution Two: Use Transparency to Eliminate QRPs AND QMPs

A part of the methodological reform movement is opening researcher’s eyes to questionable research practices that threaten the validity of their work, acknowledging that they are common and even encouraged. Taking improvement seriously means reading research with renewed scrutiny and asking critical questions about the information presented. Simmons et al. (2011) showed that a number of flexible analytic practices increase Type 1 error rates and these practices are often used to hack p -values down to the less than .05 size. Just like statistical practices need a transparency overhaul, so do measurement practices. Increasing the transparency regarding statistical analyses is only one part of strengthening the validity of our claims.

To expand the scope of transparent practices, researchers, reviewers, and editors can ensure there are no lingering questions about the measures by explicitly acknowledging QMPs and working to prevent them. The simple solution: make it easy for someone else to replicate the study and continue the process of on-going construct validation. Guidelines exist for getting this done, reviewers and editors

just need to implement them (e.g., Transparent and Open Practices Guidelines; Nosek et al., 2015; CONSORT guidelines; Schulz et al., 2010). Focusing specifically on measurement practices and QMPs, Flake and Fried (2020) provide suggest making sure the following questions are answered:

1. What is the construct and how is it defined?
2. How is it operationalized?
3. Why was the measure selected?
4. How was the measure quantified?
5. Was the measure modified? If so, how?
6. Was the measure created on-the-fly?

For some measures the answers to these questions are straightforward, and for others they are not. For certain constructs, populations, and research questions, these answers might be supported with more evidence, but less evidence for others. The first step to rigorously answering these and using measures with strong validity evidence is transparently reporting what decisions were made and why. When these questions are answered the community can evaluate the construct validity of the measure use and the validity of the study conclusions, it facilitates the replication of that study, and it supports the appropriate calibration of claims from the study.

Solution Three: Plan Replication Research with Measurement in Mind

Replication research is complex and presents logistical and methodological challenges. Educational psychologists can anticipate these challenges because of the large-scale efforts undertaken in adjacent areas of psychology. Learning from those insights, researchers can support, plan, and conduct the replication research that will make a lasting contribution to the field. To do this, researchers must consider validity holistically, evaluate the original research from the foundation up, and mend cracks in that foundation through comprehensive replication planning that includes construct validation.

If the meta-science on measure use and the reviews I discussed earlier are any insight, it is likely that original studies will lack rigorous validity evidence for the measures used. Replicators may not even be able to extract key information about the measures, such as their source, how many items there were, what the response format was, or how they were scored. If researchers are considering an original study to replicate and can't determine basic validity information for the measures, a concrete solution is to conduct a validation study first. Adding in a validation study to a replication study is a significant investment, but this is a situation in which the spirit of open and team science can support. In addition to recruiting data collection labs, replicators can recruit teams to focus on the measures and conduct the corresponding measurement research. In fact, this is an aspiration of the Psychological Science Accelerator: a methods and measurement team works alongside primary studies to conduct methodological research (Moshontz et al., 2018). Even still, resources are limited, and it is important to consider which studies are ready to be replicated. Ideally, the foundation of a study is strong before it is replicated. When replication studies use measures with poor construct validity evidence, that lack of evidence introduces confounds into the result. Conducting on-going construct validation as a part of the replication process eliminates the threats to validity that measures pose from the outset.

Even when measures have a wealth of validity evidence for their use in the original study or from previous research, large scale replication can still be challenging. When collecting data across borders, translation is needed, and procedures may need to be modified for a protocol that many labs can easily follow. Ensuring modified and translated measures hold the same interpretation across heterogeneous samples in large scale replication research is necessary for meaningful replication results. The specific methods, criteria, and thresholds for measures used in replication studies hasn't been developed yet. Should researchers conduct a full test of measurement invariance of translated measures and how should they proceed if measures are not invariant? These are the next wave of questions that need to be answered. A concrete initial solution is to gather validity evidence for measures as a first

phase of a replication study and include a construct validation section in the replication report. If the construct validity evidence is not consistent with the original evidence or doesn't support using the measure in the replication study, important knowledge about the measure has been generated and the larger goal of understanding what is and isn't replicable has been met. Generating evidence in support of measure use further solidifies the foundation on which final conclusions rest.

Conclusion

Reforming research practices to increase transparency and validity will take some trial and error before researchers get it right. The larger field of psychology, which educational psychology is part of, has undertaken a period of self-reflection to identify which reforms are needed and tested them out. This process is on-going. Though measurement practices were not a focus of these efforts at first, meta-scientific research has revealed the need for improvement. Questionable measurement practices are common, threaten the validity of original research, and complicate and limit the usefulness of replication research. As educational psychologists move forward with open science reform, they can plan for this. They can undertake meta-science research on measurement practices to pinpoint areas for improvement, scrutinize research and increase transparency with an eye to both questionable research and measurement practices, and plan replication studies that tend to the importance of construct validity. These are broad suggestions and the devil is in the details: educational psychologists can contribute to the development of the specific research practices needed to increase the transparency of measurement and impact of replication research. Through this, researchers will strengthen the foundation of science, eliminate threats to validity that occur throughout the research process, and ultimately do a better job of supporting learners.

References

- American Educational Research Association, National Council on Measurement in Education, & American Psychological Association. (2014). *Standards for educational and psychological testing*. American Psychological Association.
- Baker, M. (2016). Is there a reproducibility crisis? *Nature*, 533, 452–454.
<https://doi.org/10.1038/533452A>
- 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>
- Chambers, C. D., Feredoes, E., Muthukumaraswamy, S. D., & Etchells, P. J. (2014). Instead of “playing the game” it is time to change the rules: Registered reports at AIMS Neuroscience and beyond. *AIMS Neuroscience*, 1(1), 4–17. <https://doi.org/10.3934/Neuroscience.2014.1.4>
- Christenson, S. L., Wylie, C., & Reschly, A. L. (2012). Handbook of research on student engagement. In *Handbook of Research on Student Engagement* (Issue November). <https://doi.org/10.1007/978-1-4614-2018-7>
- 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>
- Ebersole, C. R., Mathur, M. B., Baranski, E., Bart-Plange, D.-J., Buttrick, N. R., Chartier, C. R., Corker, K. S., Corley, M., Hartshorne, J. K., IJzerman, H., Lazarević, L. B., Rabagliati, H., Ropovik, I., Aczel, B., Aeschbach, L. F., Andrighetto, L., Arnal, J. D., Arrow, H., Babincak, P., ... Nosek, B. A. (2020). Many

Labs 5: Testing Pre-Data-Collection Peer Review as an Intervention to Increase Replicability.

Advances in Methods and Practices in Psychological Science, 3(3), 309–331.

<https://doi.org/10.1177/2515245920958687>

Fabrigar, L. R., & Wegener, D. T. (2016). Conceptualizing and evaluating the replication of research results. *Journal of Experimental Social Psychology*, 66, 68–80.

<https://doi.org/10.1016/j.jesp.2015.07.009>

Flake, J. K., Barron, K. E., Hulleman, C., McCoach, B. D., & Welsh, M. E. (2015). Measuring cost: The forgotten component of expectancy-value theory. *Contemporary Educational Psychology*, 41.

<https://doi.org/10.1016/j.cedpsych.2015.03.002>

Flake, J. K., Davidson, I. J., Wong, O., & Pek, J. (2020). *Construct validity and the validity of replication studies: A systematic review*. [Unpublished manuscript]. Department of Psychology, McGill University, Montreal, Canada.

Flake, J. K., & Fried, E. I. (2020). Measurement schmeasurement: Questionable measurement practices and how to avoid them. *Advances in Methods and Practices in Psychological Science*, 1–22.

<https://doi.org/https://doi.org/10.31234/osf.io/hs7wm>

Flake, J. K., Pek, J., & Hehman, E. (2017). Construct validation in social and personality research: Current practice and recommendations. *Social Psychological and Personality Science*, 8(4), 370–378.

<https://doi.org/10.1177/1948550617693063>

Fleming, J. I., Wilson, S. E., Hart, S. A., Therrien, W. J., & Cook, B. G. (this issue). Open accessibility in education research: Enhancing the credibility, equity, impact, and efficiency of research. *Educational Psychologist*.

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

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

Gehlbach, H., & Brinkworth, M. E. (2011). Measure Twice, Cut Down Error: A Process for Enhancing the Validity of Survey Scales. *Review of General Psychology*, 15(4), 380–387.

<https://doi.org/10.1037/a0025704>

Gelman, A., & Loken, E. (2014). 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. *Psychological Bulletin*, 140(5), 1272–1280.

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

Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on “Estimating the reproducibility of psychological science.” *Science*, 351(6277), 1037-b. <https://doi.org/10.1126/science.aac4716>

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>

Kelley, T. L. (1927). Interpretation of educational measurements. In *Interpretation of educational measurements*. World Book Co.

Kidwell, M. C., Lazarević, L. B., Baranski, E., Hardwicke, T. E., Piechowski, S., Falkenberg, L. S., Kennett, C., Slowik, A., Sonnleitner, C., Hess-Holden, C., Errington, T. M., Fiedler, S., & Nosek, B. A. (2016). Badges to Acknowledge Open Practices: A Simple, Low-Cost, Effective Method for Increasing Transparency. *PLoS Biology*, 14(5), 1–15. <https://doi.org/10.1371/journal.pbio.1002456>

Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr., R. B., Bahník, Š., Bernstein, M. J., Bocian, K., Brandt, M. J., Brooks, B., Brumbaugh, C. C., Cemalcilar, Z., Chandler, J., Cheong, W., Davis, W. E., Devos, T., Eisner, M., Frankowska, N., Furrow, D., Galliani, E. M., ... Nosek, B. A. (2014). Investigating variation

in replicability: A “many labs” replication project. *Social Psychology*, 45(3), 142–152.

<https://doi.org/10.1027/1864-9335/a000178>

Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science*, 355(6325), 584–585. <https://doi.org/10.1126/science.aal3618>

Marsh, H. W., Pekrun, R., Parker, P. D., Murayama, K., Guo, J., Dicke, T., & Arens, A. K. (2019). The murky distinction between self-concept and self-efficacy: Beware of lurking jingle-jangle fallacies. *Journal of Educational Psychology*, 111(2), 331–353. <https://doi.org/10.1037/edu0000281>

Mellor, D. T. (2021/this issue). Changing norms in research culture to value transparency over novelty. *Educational Psychologist*, 56(2), ###-###.

Messick, S. (1995). Validity of psychological assessment: Validation of inferences from persons' responses and performances as scientific inquiry into score meaning. *American Psychologist*, 50(9), 741–749.

Moshontz, H., Campbell, L., Ebersole, C. R., Ijzerman, H., Urry, H. L., Forscher, P. S., McCarthy, R. J., Musser, E. D., Antfolk, J., Castille, C. M., Evans, T. R., Fiedler, S., Flake, J. K., & Chartier, C. R. (2018). The Psychological Science Accelerator: Advancing psychology through a distributed collaborative network. *Advances in Methods and Practices in Psychological Science*.

Murphy, P. K., & Alexander, P. A. (2000). *A Motivated Exploration of Motivation Terminology*. 53, 3–53. <https://doi.org/10.1006/ceps.1999.1019>

Nosek, B. A., Alter, G., Banks, G. C., Borsboom, D., Bowman, S. D., Breckler, S. J., Buck, S., Chambers, C. D., Chin, G., Christensen, G., Contestabile, M., Dafoe, A., Eich, E., Freese, J., Glennerster, R., Goroff, D., Green, D. P., Hesse, B., Humphreys, M., ... Yarkoni, T. (2015). Promoting an open research culture: Guidelines for journal could help to promote transparency, openness, and reproducibility.

Science, 348(6242), 1422–1425.

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>

Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, 2017(15), 201708274. <https://doi.org/10.1073/pnas.1708274114>

O'Donnell, C. L. (2008). Defining, Conceptualizing, and Measuring Fidelity of Implementation and Its Relationship to Outcomes in K-12 Curriculum Intervention Research. *Review of Educational Research*, 78(1), 33–84. <https://doi.org/10.3102/0034654307313793>

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716–aac4716. <https://doi.org/10.1126/science.aac4716>

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>

Patall, E. A. (2021/this issue). Implications of the open science era for educational psychology research syntheses. *Educational Psychologist*, 56(2), ###-###.

Plucker, J. A., & Makel, M. C. (2021/this issue). Replication is important for educational psychology: Recent developments and key issues. *Educational Psychologist*, 56(2), ###-###.

Qualls, A. L., & Moss, A. D. (1996). The degree of congruence between test standards and test

documentation within journal publications. *Educational and Psychological Measurement*, 209–214.

Reich, J. (2021/this issue). Preregistration and Registered Reports. *Educational Psychologist*, 56(2), ###-###.

Rhoads, C. H. (2011). The implications of “contamination” for experimental design in education. *Journal of Educational and Behavioral Statistics*, 36(1), 76–104.

<https://doi.org/10.3102/1076998610379133>

Schulz, K. F., Altman, D. G., Moher, D., & Group, the C. (2010). CONSORT 2010 Statement: updated guidelines for reporting parallel group randomised trials. *Trials*, 11(1), 32.

<https://doi.org/10.1186/1745-6215-11-32>

Shaw, M., Cloos, L. J. R., Luong, R., Elbaz, S., & Flake, J. K. (2020). Measurement practices in large-scale replications : Insights from Many Labs 2. *Canadian Psychology*, 61(4), 289–298.

<https://doi.org/10.1037/cap0000220>

Shrout, P. E., & Rodgers, J. L. (2018). Psychology, science, and knowledge construction: Broadening perspectives from the replication crisis. *Annual Review of Psychology*, 69(1), 487–510.

<https://doi.org/10.1146/annurev-psych-122216-011845>

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>

Simons, D. J. (2014). The Value of Direct Replication. *Perspectives on Psychological Science*, 9(1), 76–80.

<https://doi.org/10.1177/1745691613514755>

Simons, D. J. (2018). Introducing Advances in Methods and Practices in Psychological Science. *Advances in Methods and Practices in Psychological Science*, 1–4.

<https://doi.org/10.1177/2515245918757424>

- Slaney, K. (2017). *Validating psychological constructs: Historical, philosophical, and practical dimensions* (J. Martin (ed.)). Palgrave Macmillan. <https://doi.org/10.1057/978-1-137-38523-9>
- 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>
- Thorndike, E. L. (1904). *An introduction to the theory of mental and social measurements*. Teachers College, Columbia University.
- 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*.
- Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., van Aert, R. C. M., & van Assen, M. A. L. M. (2016). Degrees of freedom in planning, running, analyzing, and reporting psychological studies: A checklist to avoid P-hacking. *Frontiers in Psychology*, 7, 1–12.
<https://doi.org/10.3389/fpsyg.2016.01832>
- Zumbo, B. D., & Chan, E. K. H. (2014). *Reflections on Validation Practices in the Social , Behavioral , and Health Sciences*. i, 321–327. <https://doi.org/10.1007/978-3-319-07794-9>