From Old School to Open Science: The Implications of New Research Norms for Educational Psychology and Beyond

Hunter Gehlbach

Johns Hopkins University

gehlbach@jhu.edu

Carly D. Robinson

Brown University

carly_robinson@brown.edu

Citation:

Gehlbach, H., & Robinson, C. D. (2021). From old school to open science: The implications of new research norms for educational psychology and beyond. *Educational Psychologist*, 56(2). doi: 10.1080/00461520.2021.1898961

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:

10.1080/00461520.2021.1898961. Once the article is published online, it will be available at the following permanent link: <https://doi.org/10.1080/00461520.2021.1898961>

Abstract

Recently, scholars have noted how several “old school” practices—a host of well-regarded, long-standing scientific norms—in combination, sometimes compromise the credibility of research. In response, other scholarly fields have developed several “open science” norms and practices to address these credibility issues. Against this backdrop, this special issue explores the extent to which and how these norms should be adopted and adapted for educational psychology and education more broadly. Our introductory article contextualizes the special issue’s goals by: overviewing the historical context that led to open science norms (particularly in medicine and psychology); providing a conceptual map to illustrate the interrelationships between various old school as well as open science practices; and then describing educational psychologists’ opportunity to benefit from and contribute to the translation of these norms to novel research contexts. We conclude by previewing the articles in the special issue.

Introduction

Our initiation into open science practices felt a bit like having the rug pulled out from underneath us. It began several years ago, when one of us (Gehlbach) published promising findings on a “birds of a feather” intervention (Gehlbach et al., 2016). By providing feedback from a get-to-know-you survey that highlighted commonalities between teachers and students, the research team found that the intervention resulted in improved classroom relationships and academic achievement. The headline finding showed that the treatment closed the achievement gap for underrepresented minority students by two-thirds. In addition, the authors tried out a relatively new idea at the time that has subsequently become a staple practice of open science: they preregistered six focal hypotheses. In other words, before conducting any data analyses the research team specified which hypotheses were going to be tested and reported regardless of the statistical significance or effect sizes of those results. They reported all other analyses in the final published article (including the headline finding) as “exploratory.”

Around this same time, Robinson joined Gehlbach’s research team, excited to conduct this type of research that might enhance social aspects of adolescents’ schooling experience. Together, we enthusiastically embarked upon a journey to see if the “birds of a feather” intervention might produce similar effects across different contexts and on different populations (including a return to the original school to conduct a direct replication). Because the original study’s findings had been preregistered, we retained a confidence for these replication studies that was atypical of our skeptical dispositions. Unfortunately, as it turned out, skepticism would have been more warranted than confidence.

Years later, having reaped the benefits of hindsight and the hard-earned wisdom that multiple failed replication attempts confer, it seems painfully obvious what went wrong (and

how problematic that first preregistration attempt was). Gehlbach et al. (2016) did not correct for the number of hypotheses tested, articulate specific criteria to adjudicate which participants should be excluded, or anticipate several other pre-analytic decisions that later proved important. Most problematic, the most exciting result had emerged from an exploratory analysis focusing on a subset of the sample. The published paper clearly distinguishes the preregistered findings from the exploratory. However, these exploratory results mirrored many other social psychological findings in educational settings and caused us to remain overly optimistic that we would replicate the original exploratory findings to some degree. Although it took many failed replication efforts, that optimism has now been thoroughly extinguished.

Adopting and Adapting Open Science

Cognitively, we knew that the mere act of preregistering studies was not enough to ensure the level of rigor we sought in our work. But the painful experiential learning gained from these replication attempts proved to be a much more potent catalyst than mere knowledge for changing our research habits. Steadily we began to nudge our research program to adopt increasing numbers of the new research practices that fall under the “open science” umbrella (van der Zee & Reich, 2018), occasionally adapting some of them with our own ideas. We now better understand the processes that lead to “illusory results”—findings that seem real and compelling but that typically arise from researchers taking advantage of multiple analytic choices, arriving at a statistically significant result by chance, and developing a theoretical story to explain the results. We have experienced not only why pre-specifying hypotheses is so important but also why the process is so tricky for applied research (Gehlbach & Robinson, 2018). Since our initial preregistration attempts, we have benefitted immensely from the scores of new resources that have emerged to help others avoid our mistakes. For example, the Center

for Open Science (2020) offers a comprehensive guide to facilitate not only preregistration (<https://osf.io/jea94/>), but also storing and sharing data, code, and other study materials to facilitate more scientific transparency. Hosting services such as [EdArXiv](#) allow authors to post pre-prints and working papers, thus enabling feedback prior to publication. New approaches facilitate the publishing of null results. For instance, a “registered report” entails a peer review process in which reviewers evaluate manuscripts on the merits of the research question and study design, not on their results. One of us (Robinson) even published one of the failed “birds of a feather” intervention replication attempts this way instead of relegating the results of that study to be buried in a file-drawer (Robinson et al., 2019). Finally, we have developed some of our own norms and practices, such as dividing our results sections into pre-specified and exploratory hypotheses (e.g., Robinson et al., 2019).

Beyond our collaborations, however, the discipline of educational psychology faces a strange historical moment. Some worry that the discipline now faces inequities that pose professional problems internally. On the one hand—thanks to knowledge of how these issues have unfolded within psychology—a number of educational psychologists are aware of the full array of these issues (and practices to address them) that fall under the open science umbrella: *p*-hacking versus preregistration; opacity versus open-methods, open-materials, and open-access; publication bias versus replication studies; file-drawer problems versus registered reports; and so forth. On the other hand, many researchers have had no exposure at all. This uneven distribution of knowledge creates major problems: Scholars practice science under radically different conditions. Different peer reviewers may play by different rules. Early career scholars might face changing standards right as they go up for tenure. Certain beloved theories may

prove to be rock solid while others may rest on a foundation of illusory results, leaving researchers with no clear path for adjudicating which is which.

More worrisome still, as a discipline within an applied field, the implications for practice are especially concerning. How can we expect practitioners to distinguish between real and illusory findings? Should school leaders make research-based decisions if the research rarely replicates? How should educational policies be constructed if certain fundamental theories rest upon shaky empirical foundations?

Goals of the special issue

Within this context, we thought that a special issue in *Educational Psychologist* could make several important contributions. First, we hope that the articles in the special issue will allow scholars to work from a common understanding of open science. Unless all scholars within the discipline know what the practices are, how they work, and what their limitations are, they will talk past each other. Moreover, practitioners and policymakers will be left with very mixed messages from the research community about what results are credible and which are dubious.

Second, building from this common understanding, we hope the articles open a nuanced discussion about which reforms could benefit the discipline most and where new tensions may arise. Some open science practices may be readily adopted, some adapted, and some inappropriate. We describe this process of figuring out which practices may be useful, in which situations, with which modifications, as a transition from “old school” practices to open science. We use “old school” as an umbrella term that conveys respect for a host of well-regarded scientific norms that often served the discipline well. Yet, many scholars also now realize that these practices may have inadvertently disguised some consequential problems. Old school

practices range from the analytic flexibility researchers allow themselves (such as trying out different covariates, different ways to create composite measures in their study, etc.) that have been well documented elsewhere (Gehlbach & Robinson, 2018; Simmons et al., 2011) to traditions like keeping measures, materials, and analytic code private unless requested by other scholars. Three features unite this broad range of old school practices. First, they are employed with the intention of serving a sensible purpose (e.g., discovery of an important finding that would go undetected without proper covariates or protection of researchers' intellectual property). Thus, many of these research norms are distinct from questionable research practices. Second, without proper guardrails to moderate the practice, they can cause consequential problems such as illusory results or an inability to replicate findings. Third, they tend to decrease transparency, which jeopardizes attempts to understand and replicate findings.

Finally, with this common understanding and the beginnings of a nuanced conversation in place, the articles will explore potential compromises and solutions to these tensions. Although it is premature to think that we can definitively solve all these tensions, the special issue contributes a series of new ideas that should shed light on which open science practices are likely to prove most valuable in which contexts. Authors address questions ranging from whether pre-prints can be posted while still maintaining blind peer review to how, if at all, meta-analyses should differentiate preregistered hypotheses from exploratory (Patall, 2021/this issue) to how researchers should address unexpected events from a field study in their registered reports (Reich, 2021/this issue). Finally, by highlighting the new tensions that will inevitably arise as practices designed for one context are imported into another as well as exploring potential compromises, the special issue aims to accelerate educational psychology's transition to these new scientific norms. In turn, we hope that faster adoption of norms to promote better science

helps bolster practitioners' faith in educational research. Towards these ends, the first half of the special issue aims to educate current and future scholars about the core issues and practices that fall within the open science umbrella. The second half of the special issue anticipates several implications and ripple effects that will occur as these new research practices become the norm.

As the co-editors of this special issue, we claim no unique expertise regarding the facets of open science. We wrestle with these issues anew with each additional study we conduct; we continue to make mistakes and learn from them. That said, we will credit ourselves for having assembled a team of scholars with tremendous expertise. We invite readers to join us in learning from this panel of early-career to senior colleagues who are pushing our thinking on seminal open science issues.

To frame the special issue, we use our introductory article to first look beyond educational psychology to provide some key historical context driving this push to change scientific norms. Second, we illustrate the ways in which some old school scientific practices can reinforce persistent problems, with distinct consequences for researchers and for practitioners. Mirroring this illustration, we show how open science practices could catalyze a new set of scientific norms for the benefit of researchers, practitioners, and policymakers. Next, we zero in on educational psychology to describe how the discipline occupies an ideal position for wrestling with the next generation of open science challenges and identify what some of those challenges are likely to be. Specifically, we propose that educational psychologists could reap substantial benefits from open science practices and simultaneously could advance the sophistication of open science practices within the discipline of educational psychology as well as for the field of education more broadly. We conclude by previewing the specific articles in the special issue.

Key Historical Context for Open Science

Our failed replication story is hardly the first such episode. However, scholars' collective awareness regarding the pervasiveness of the challenges in developing reproducible scientific findings has developed relatively recently. In medicine, Ioannidis (2005) forcefully brought the issue to the fore with the provocatively titled article, "Why most published research findings are false." Through simulations, he showed why most research claims are more likely to be illusory than accurate. He went on to identify multiple study characteristics—small sample sizes, modest effect sizes, execution of a large number of statistical tests, financial interest, and methodological flexibility—that frequently combine to boost the odds that researchers may be reporting illusory results. This work contributed to a number of reforms to how medicine conducts randomized clinical trials. Although detailing the effects of these reforms falls beyond our focus, Blattman's (2016) humorous blog title speaks volumes about the impact of these changes: "Preregistration of clinical trials causes medicines to stop working."

In addition to sparking reforms in medicine, these ideas migrated into other disciplines—including psychology. In a comparably influential article, Simmons, Nelson, and Simonsohn (2011) took a different approach to demonstrating the likely pervasiveness of illusory results in psychology. They devised a study whose characteristics closely resembled many laboratory-based studies in psychology with one key difference: they strategically combined a sufficient number of old school practices to essentially ensure that they would find a statistically significant, but impossible, result. In their demonstration experiment, they asked undergraduates to listen to different types of music to "prove" that listening to the Beatles *When I'm Sixty-Four* makes people younger—note the causal language. For the bulk of the article, the authors then transparently described the numerous "researcher degrees of freedom" they took advantage of by

iteratively conducting post hoc analyses until they found a statistically significant, yet illusory, result. Of particular interest, they carefully calculated how combining many long accepted (and important) old school practices—including flexibility in: terminating data collection, determining a study’s minimum sample size, fully reporting the variables collected during the study, reporting of all experimental conditions, identifying which data points constitute outliers, and using covariates as statistical controls—in *combination* bolstered their chances of identifying a statistically significant result. Trying out so many analyses enabled them to find just the right sample size with just the right covariates on just the right outcome measure to capitalize on chance to produce the statistically significant result. The genius of this demonstration study was to produce an illusory result—that the Beatles’ music can turn back the hands of time—that all readers would easily recognize as impossible.

In psychology the ensuing debate sparked intense discussion within the field about a replication crisis (Maxwell et al., 2015). At times, the discussion became directed at individual scholars and spilled out into the popular press (e.g., Dominus, 2017). Controversy persists regarding the replicability of psychological experiments (Gilbert et al., 2016; Open Science Collaboration, 2015). Despite the challenges and controversy, psychology has developed numerous innovations that not only address researcher degrees of freedom, but have improved upon a number of related scientific issues and have pushed to change long-standing norms in the field (Nosek et al., 2015). For example, the flagship journal of the *Association of Psychological Science* offers badges for preregistered studies, for articles in which researchers make their study materials openly available, and for open-data (Association for Psychological Science, 2017). One of the Association’s new publications even explicitly lists metascience—using scientific methods to study how science is conducted—as a core topic within the journal’s scope (Simons,

2018). Multi-site collaborative replication projects have also helped clarify how robust particular findings have been. For instance, the “Many Labs” Replication Projects examine the variation in replicability of classic and contemporary psychological effects across new samples and settings (Klein et al., 2014; Klein et al., 2018). Efforts like these identify psychological studies with effects that may have been underestimated (e.g., anchoring; Jacowitz & Kahneman, 1995), overestimated (e.g., sex differences in implicit math attitudes; Nosek et al., 2002), possibly illusory, or reliable and robust. Furthermore, researchers are engaging in new collaborations to replicate their previously published research findings and updating the scientific record when they do not replicate (e.g., Kristal et al., 2020). In short, much progress has been made that might benefit educational psychologists.

How Key Facets of Open Science Interrelate

Before thinking through how these innovations might need to be tailored for educational psychology, another key piece of context is important. Old school and open science practices emerge out of a historical context; however, they also exist in a current research ecosystem where different scientific norms affect one another in important ways (see Figures 1a and 1b).

Old school science in context

Scientific practices proliferate and endure presumably because scholars engage in the type of research they were trained to conduct within a context that expects and rewards those practices. Many foundational tenets of the educational psychology literature and many other social science disciplines have been developed by dedicated, thoughtful scholars who thoroughly analyzed their data. As a result of carefully examining their data, researchers have often been better able to present cleaner data (e.g., removing the “noise”—or random variation—from data

Figure 1a: A Hypothesized Pathway through which Old School Science Research Practices could affect Research and Practice Processes

Figure 1b: A Hypothesized Pathway through which Open Science Research Practices could affect Research and Practice Processes

by disregarding outliers or adding additional covariates) that may facilitate the discovery of new, interesting nuances to established theories. Over time, increasingly nuanced theories may have come to rest upon an empirical foundation of data that appear exceptionally “well-behaved.” The well-behaved data and clean findings presumably shaped the expectations of other researchers, including journal reviewers. Consequently, scholars may have felt new pressure to produce additional novel and exciting caveats to those theories with even less noisy data. At first, careful examinations, cleaner data, theoretical nuances, etc. may sound like positive developments. However, these efforts and pressures often lead researchers to engage in research practices that have become increasingly problematic for the scientific process. In Figure 1a, we present a hypothesized model of how decades of these unrealistic expectations, combined with editors’ and journals’ appetite to publish novel and exciting findings, might impact the scientific process. Old school research practices, on the left-hand side of the figure, have led peer reviewers to expect studies with tight narratives that lead to clear, novel findings.

Consequently, researchers face increased pressures to produce exciting and statistically significant—and often illusory—results. These pressures engender a system that reinforces academics for engaging in old school practices. The implications of illusory results—as illustrated on the far right-hand side of the figure—is that practitioners and policymakers drift increasingly further away from realizing value from educational psychology research. To be clear, the hypothesized model is designed as an extreme case to illustrate one end of the continuum and one, of many, interrelationships between these practices.

The figure’s left-hand column illustrates several common old school research practices. Researchers have described “*p*-hacking” in many different ways—as researchers walking down a “garden of forking paths” to find the most appealing, statistically significant results (Gelman &

Loken, 2014); as scholars who take advantage of “researcher degrees of freedom” (Simmons et al., 2011); or as data-dredging, snooping, fishing, and significance chasing (Nuzzo, 2014).

Regardless of one’s preference in terminology, the core problem remains: By running enough analyses, researchers can easily cross over the fine line from thoroughly analyzing their data into torturing their data until it tells them what they want to hear. In short, *p*-hacking is the practice of running a multitude of analyses in search of statistically significant, publishable results.

Although we presume that most researchers engage in this practice inadvertently—in the spirit of thorough analysis rather than torturing data—illusory results follow just the same, regardless of researcher intent. Furthermore, because humans are inveterate story tellers (Cron, 2012), academics can readily make sense of their findings for their audience of choice.

Because illusory results emerge from researchers exploring so many forking pathways in their data, reading fully transparent manuscripts that detail *all* data analysis would be excruciating. Not unreasonably, researchers typically write scientific manuscripts to detail the data analysis pathway that lead to the successful, statistically significant results of interest. Failed pathways are rarely documented, and the measures that did not contribute to the significant finding often go unreported.¹ Naturally, without full access to all measures and analyses, replication attempts are often doomed from the outset. Even if researchers have access to the same measures and similar participants, if the original analyses capitalized on the aforementioned flexibility in “researcher degrees of freedom,” the odds of a successful replication are vanishingly small (see Plucker & Makel, 2021/this issue for more on different types of replication studies). Additionally, knowing that journals prefer to publish statistically

¹ Simmons et al. (2011) offer a powerful illustration of what transparent reporting of results should look like by fully reporting the key results in their Table 3, but then **bolding** only those parts that would typically be reported in an old school approach.

significant results creates a “file drawer” problem (Rosenthal, 1979) in which researchers are incentivized to stick their failed studies, including failed replications, into a metaphorical file drawer rather than trying to publish them. As a result, individual researchers soon learn that pursuing original, rather than replication, research may leave them better off. At the collective level, educational psychology then faces major challenges with respect to the aggregation of knowledge. Those replications that do exist may use different measures on different populations (see Flake, 2021/this issue for more on these challenges). Perhaps more problematic, the sample of studies to draw from for research syntheses is likely to be biased towards those that find statistically significant results (Patall, 2021/this issue).

These challenges around aggregated knowledge may lead to a plethora of competing, nuanced, idiosyncratic theories that attempt to cope with the mixed empirical results that emerge in a particular area of study. In other words, as different researchers take up different measures or navigate different pathways through their data in the pursuit of a statistically significant finding, underlying explanatory theories may need to undergo creative contortions to accommodate the array of different findings. As an example of measurement challenges, Flake (2021/this issue) shows how jingle-jangle fallacies can easily inflict additional theoretical complexity. If grit and conscientiousness are the same construct, scholars should be able to simplify theories of how goals and personality traits affect schooling outcomes (e.g., Muenks et al., 2017). Wentzel (2021/this issue) provides a similar example regarding how self-efficacy/competence/confidence/expectancies for future success/academic self-concept are described, measured, and theoretically distinguished.

The additional theoretical complexity that multiple analytic pathways can potentially cause creates headaches for policymakers and practitioners. The values affirmation intervention

that draws from self-affirmation theory offers an intriguing case. Over a host of studies, scholars have found that when students remind themselves of some of their most important personal values, a number of positive academic outcomes have emerged. Benefits have been documented for African-American seventh graders (Cohen et al., 2006), Latinx middle schoolers (Sherman et al., 2013), first-generation college biology students (Harackiewicz et al., 2014), and collegiate female physics students (Miyake et al., 2010). However, null effects have also been found for minority middle school students (Dee, 2015; Hanselman et al., 2016), under-represented minorities in the same aforementioned undergraduate biology class (Harackiewicz et al., 2014), as well as the same collegiate female physics students when those data were analyzed by different researchers (Serra-Garcia et al., 2020).² Numerous moderators have emerged to create nuanced theoretical explanations for these heterogeneous results including: the timing of the intervention (e.g., before a perceived threat), the framing of the writing exercise that comprises the crux of the intervention (e.g., as a normal part of class or as beneficial), the type of control group activity, degree of identification with the negatively stereotyped group, knowledge of self-relevant negative stereotypes, caring about school performance, classroom composition (e.g., demographic heterogeneity), and the extent to which a threat is “in the air” within the classroom contexts (see Hanselman et al., 2016).

The key challenge is this: Until there is a norm of transparency around which (and how many) analytic pathways are tested and which are reported (e.g., through preregistration), it becomes nearly impossible to know which pathways and moderators reflect existing real-world

² It is worth noting that part of the reason educational psychologists know about all the conflicting results in this line of research is because key individuals were willing to share materials and data sets. This level of openness has helped this line of research mature rapidly. The additional step of preregistered future studies may help provide the definitive evidence practitioners need to ascertain the extent to which these interventions work for different populations of students.

complexities and which represent analytic pathways that capitalized on chance. As frustrating as this uncertainty is for researchers, it creates an untenable situation for practitioners and policymakers hoping to put the intervention to use in classrooms.

With these complex, nuanced theories in hand, researchers face expectations from peer reviewers to articulate how their studies contribute new nuances to already complex theories. Reviewers may disproportionately expect authors to make these contributions through unrealistically tidy, statistically significant results and may recommend publication (or not) based on those expectations.

By contrast, practitioners and policymakers may not have much of a sense of what to do with the complicated, conflicting results. Moreover, if the promise of a particular line of research rests on a foundation of illusory results and the odds of replication are low, then practitioners' attempts to use research will likely prove disappointing. This confusion and frustration may serve as a disincentive for using research in their daily decision-making if no clear guidance is offered. Having had mixed experiences with research, practitioners' and policymakers' appetite for future research may diminish. Perhaps these challenges help to explain why practitioners and policymakers do not currently embrace educational psychological scholarship as much as researchers might hope (Phillips et al., 2009).

Open science

In Figure 1b (a hypothesized model designed to illustrate the other end of the continuum), we portray a very different scenario under a hypothetical world in which educational psychologists adopt open science norms and practices widely. The left-hand column of the figure highlights many interrelated practices. Preregistration is a practice in which scholars publicly post their exact study design and analytic approach before the study is conducted or

before the data are analyzed (Gehlbach & Robinson, 2018). Because the focal analyses are pre-specified, results can be written up in a fully transparent way—neither fishing for statistical significance nor *p*-hacking one’s way to a compelling set of findings are possible with preregistered hypotheses. Exploratory analyses would still be an important contribution in many articles, but they would have to be clearly labeled as such. Because the measures of interest are also pre-specified, it becomes much easier and less risky for researchers to openly share their materials and data sets; in other words, which measures were included in a study is already public information. This transparency suddenly allows replication studies to become much more viable—materials, methodologies, and even statistical code are already available. As replication studies become more viable, some findings will inevitably replicate while others will not. However, the need for the scholarly community to see successful *and* failed replications will become clearer. “Registered reports” offer a new publishing approach that meets this need. In this process, authors first submit an introduction, methods section, and a detailed analysis plan (ideally in the form of a publicly available preregistration). At this point, the submission undergoes peer review. With no results present, reviewers must evaluate the importance of the research question and can suggest revisions or improvements to the research design and analysis plan. If and when the editors are satisfied, authors receive an *in principle acceptance*, meaning that their article will be accepted provided that they adhere to the research plan as described or offer compelling explanations for deviations from the original plan (Reich et al., 2020). Because registered reports are evaluated without respect to the magnitude and direction of the research results, they address the file-drawer problem and help build an empirical knowledge base of much higher fidelity. As publications of null findings emerge more frequently, the support for a

number of current theories may atrophy and educational psychology could be left with fewer, stronger theories.

As these open science practices are adopted, several research-related consequences of these new practices would follow, as illustrated in Figure 1b. First, new norms would become instantiated during the peer review process (Mellor, 2021/*this issue*; Nosek et al., 2015). Authors and reviewers would need to recalibrate to become more appreciative of preregistered hypotheses, more tolerant of non-statistically significant and noisy findings, more discriminating about exploratory results, more demanding about which research questions might be important to ask in the first place, less expectant of nuanced stories where the data work out perfectly, and even open to reviewing research designs with no results at all (in the case of registered reports). To the extent that authors write for these types of reviewers, the incentives will be to produce credible results; the pressure for statistically significant findings will abate. As a result, educational psychology should have fewer problems with publication bias, which, in turn, should reinforce the adoption of all of these open science practices.

On the right-hand side of Figure 1b we illustrate the implications for practitioners and policymakers. The hypothesized winnowing to a smaller number of much stronger results should clarify how research might help them make smarter, data-driven decisions. As they have positive experiences using research, it should boost their appetite for and faith in the scholarship from educational psychology.

To be sure, the contrasts we have drawn between the processes of old school and open science are oversimplified and speculative. Multiple other connections exist between elements in both models and the open science practices are somewhat separable. For instance, it may not make sense to preregister a certain study (Reich, 2021/*this issue*), but the researchers could still

openly share their data and materials (Fleming et al., 2021/this issue). It seems immensely likely that theories will continue to face challenges in navigating the context-dependent nature of findings. However, it is clear that scientific processes will shift. These shifts will cause major changes in educational psychology which, in turn, will inevitably impact practitioners and policymakers. Thus, our hope is that this article provides a starting point for the conversation that guides the most sensible way to make these changes.

Open Science in Educational Psychology: Adoption Challenges and Future Tensions

Given the historical context that includes so much progress from other disciplines—particularly psychology— and seeing the benefits of how open science practices can support each other, it might seem straight-forward for educational psychology to adopt these new norms and the infrastructure that supports them. However, several key challenges will complicate adoption.

Winning over skeptical hearts and minds

We anticipate colleagues' skepticism will be far and away the biggest barrier to the adoption of open science practices. Many educational psychologists will very reasonably question how big an issue illusory results are in the first place. A key ingredient to the progress made by medicine and psychology was the publication of a high impact article that spoke to the prevalence of illusory results in the field. Ioannidis' (2005) simulation and Simmons et al.'s, (2011) faux study are powerful illustrations in part because many scholars in their fields utilize experimental research designs. Thus, large numbers of researchers within medicine and psychology can readily resonate with the magnitude of the problem even if particular scholars typically engage in other types of research. While the red flags for illusory results identified in these other disciplines may worry many educational psychologists, others need to see formal

documentation of the problem before they are convinced. Furthermore, given the context-specific nature of teaching and learning, how much replication of results should be expected in the first place?

Few would argue that educational research suffers from an excessive number of replication studies. Makel and Plucker (2014) found that 0.13% of published educational research studies comprised replication attempts. This paucity of replication studies is a major contributor to why it is so hard to establish a baseline frequency of false findings. Furthermore, establishing the prevalence of illusory results within a discipline that embraces inevitably messy, context-dependent, methodologically diverse studies is a daunting task to say the least. Although it seems unlikely that a single article might have the same impact as occurred in medicine and psychology, some research approaches could shed light on the pervasiveness of the problem.

Researchers might examine how problematic illusory results may be in educational psychology through a number of approaches. For instance, scrutinizing the distribution of statistically significant *p*-values—or the “*p*-curve”—for studies published in educational psychology journals would allow scholars to identify, in the aggregate, whether effects are likely to be true or a reflection of selective reporting. As Simonsohn, Nelson, and Simmons (2014) explained, a left-skewed *p*-curve (i.e., finding *p*-values of .03 and .04s disproportionately more than *p*-values of .01 and .001) suggest the presence of selective reporting and, therefore, illusory results. On the other hand, a right-skewed *p*-curve would indicate that, generally, we can have greater faith in the published results.

Another approach could involve documenting the prevalence of phenomena that suggest publication bias and/or that researchers are taking advantage of small samples to obtain statistically significant results. For example, some researchers have examined the correlations

between effect sizes and sample size (e.g., Slavin & Smith, 2009). Scholars who review and synthesize the effectiveness of reading and mathematics programs have found that studies with smaller sample sizes tend to have much larger positive effect sizes than those with larger sample sizes (Slavin & Lake, 2008; Slavin et al., 2009). Educational psychologists could similarly take advantage of the fact that effect size and sample size should be uncorrelated to explore this association within the discipline's flagship journals as a step toward identifying the presence of illusory results.

Finally, the mere act of documenting the prevalence of study characteristics that can produce illusory results in published educational psychology studies could provide a useful baseline. For example, Gehlbach and Robinson (2018) identified study features such as small sample sizes, key results occurring in sub-group analyses, multiple dependent variables, liberal testing of covariates, and other common study characteristics that are often a part of rigorous science. However, without preregistering, as a single study employs increasing numbers of these practices, it boosts the probabilities that illusory findings may result. Thus, as the field moves towards open science practices, scholars might compare the prevalence rates of these study features over time (particularly in journals that explicitly adapt to embrace open science practices) to see whether the rates of these potentially problematic characteristics begin to decline.

Regardless of whether scholars can precisely document the prevalence of illusory results, it is hard to argue against the goal of improving scientific practices. Furthermore, educational psychology seems particularly well-suited to take on the challenge of adapting and adopting these new norms for multiple reasons. First, as a branch of psychology, it should be easy to benefit from much of the progress that other subdisciplines of psychology have made while also

tackling the challenges that are unique to educational psychology. Second, given the prominent role of educational psychology within the broader educational research community, developing pragmatic norms and resolving key tensions could provide valuable models for other disciplines within the field of education. For instance, how should authors and journals balance the transparency concerns around open data with the need to maintain participants' privacy (Fleming et al., 2021/this issue)? What types of research methodologies are viable for registered reports (Reich, 2021/this issue)? What expectations are reasonable for replication studies across distinct methodological approaches (Plucker & Makel, 2021/this issue)? Answers to these types of questions could be helpful for sociologists of education, educational historians, educational policy scholars, and so forth. Third, because of the applied nature of the discipline, we owe it to practitioners and policymakers to provide them with the highest integrity, most robust knowledge to guide their decision-making.

Additional challenges and future tensions

Convincing sage, skeptical colleagues that they should undertake a slate of new practices—despite their substantial success with old school practices—is not the only adoption challenge. For open science practices to be thoughtfully infused into educational psychology, multiple other issues warrant consideration. We focus on three broad challenges that we perceive to be grist for some of the richest discussions moving forward—beginning with, but extending beyond, this special issue.

First, the context-specific nature of so many of the studies in educational psychology makes replication a hazy proposition. To what degree should we expect measures and materials used successfully in an urban school to work comparably well in a rural district? If they use different textbooks, should the same intervention designed to bolster students' interest in the

subject matter have identical effects in two different college classrooms? How educational psychologists create guidelines for expectations around replication attempts will have major implications for everyone from journal editors deciding the fate of articles to policymakers deciding what research results seem sufficiently robust to influence policies.

Second, educational psychology draws from a particularly rich diversity of methodological approaches. How should researchers adapt these new norms and practices for measurement studies, secondary data analysis on large data sets, correlational/longitudinal studies, qualitative inquiry, and so forth? In one scenario, the field might dismiss these open science practices as only relevant for controlled experiments. To us, doing so would represent a missed opportunity to not only improve scientific processes within educational psychology but also to have educational psychology advance the open science practices of the future.

Third, educational psychologists' tradition of careful attention to measurement raises numerous intriguing tensions with respect to open science practices. Researchers focused on measure development are attuned to the fact that there is no such thing as a "valid measure." Rather, measures show greater or lesser evidence of validity for a certain use on a certain population within a certain context (Gehlbach, 2015). So if validity can only properly be assessed after data describing the participants and their context are collected and analyzed, how will open science practices like preregistration and registered reports that require *a priori* assumptions work (Flake, 2021/this issue)? If measures function differently with different participant populations, how can we evaluate whether a replication attempt has succeeded (Plucker & Makel, 2021/this issue)? Many researchers view their measures as intellectual property and charge fees for their use. This approach creates substantial tension with practices like open materials (Fleming et al., 2021/this issue). Again, educational psychologists have

much to figure out in terms of how open science practices might be adapted for the discipline and much to contribute to the continued development of open science norms.

By no means are these the only challenges for educational psychologists in adopting and adapting open science practices. Other challenges arise around specific practices. For instance, the practice of preregistration forces researchers to make *a priori* choices about the precise steps they will take. This norm may be less realistic for educational psychology research that occurs through field studies. The fidelity with which researchers can execute their study can be attenuated by a nearly infinite array of unplanned events: busy teachers stop participating, students get sick on the day of a follow-up assessment or observation, not to mention that fire-alarms and snow days seem to occur at the worst possible times for researchers. Starting to raise and address these adoption and adaptation challenges represents a major part of the impetus to begin a discussion with this special issue. Researchers, practitioners, and policymakers all stand to benefit the sooner educational psychology researchers begin to thoughtfully consider adoption and adaptation challenges and begin experimenting with ways to incorporate open science practices into their research processes.

Overview of the special issue

The first half of the special issue lays out several core challenges that educational psychologists have faced and how open science practices can help address them. In the first article, Plucker and Makel (2021/this issue) address the challenge of replication with special attention to the different types of replication studies that might be pursued and what factors might need to be considered for non-experimental replication studies. In the second article, Reich (2021/this issue) addresses the practices of preregistration and registered reports. Drawing on his experience as an editor of registered reports, he examines the range of approaches to these

two practices across different research methodologies. Next, Fleming et al. (2021/this issue) discuss open practices in educational psychology. Themes of transparency and openness are being applied broadly to address issues such as open materials, open data, open statistical code, and open access to research reports (including preprints). They discuss these practices with particular attention towards democratizing science. Because adoption of all of these practices will require a discipline that shifts its respective valuing of transparency and novelty, Mellor (2021/this issue) addresses the topic of changing norms, which serves as a transition to the second half of the special issue. He describes how cultural change might begin through the actions of grassroots groups, communities of practice, the leadership of policymakers, and the adoption of specific incentives such as badging.

The second half of the special issue looks to the future to describe the implications and ripple effects of these practices within educational psychology and across educational research more broadly. In other words, as these more rigorous norms become instantiated, what will change for researchers, policymakers, and practitioners? Although these articles are necessarily more speculative, they build from the progress made in other disciplines. Flake (2021/this issue) looks within the discipline at how open science practices will affect educational psychology's rich tradition of measurement and validity. How will measurement practices and approaches to establishing evidence of validity change under new norms of transparency? Patall (2021/this issue) continues to explore issues of validity in her article, but at the level of aggregated knowledge across the field. If meta-analyses are built upon studies that are riddled with problematic practices, do the biases from these practices cancel each other out or magnify one another? More importantly, how might research syntheses need to change in an era when open science practices are the norm? Finally, Wentzel (2021/this issue) examines open science

practices within a broader set of challenges that could be addressed by these reforms. She focuses particularly on the likely impact of preregistration and registered report processes if they require researchers to incorporate a more explicit focus on theory. Greater attention to the conceptual underpinnings of research could lead to improvements in theoretical specificity and by implication, to sampling and measurement issues. In turn, this focus has the potential for long-term positive consequences such as enhancing the theoretical foundations of educational psychology. Using achievement motivation as a case example, she entertains how core understandings may change as theory becomes incorporated into the research process to a greater extent. She also explores the challenges of open science practices for informing educational policies and practice.

Our initial efforts to incorporate open science practices into our own research were riddled with trial and error. We have since found many likeminded people—many of whom are authors in this special issue—working to increase the transparency and rigor of educational research. They too are borrowing best practices from other fields, adapting them, and engaging in trial and error. By assembling these ideas in one place, we hope that others can learn from past mistakes and start testing new ideas. Shifting from old school science to open science will undoubtedly face challenges and roadblocks. At the risk of stating the obvious, open science is not a panacea that will solve all the problems educational psychologists and educational researchers face. However, the ideas presented within these pages represent important starting points for a discussion on how to improve scientific processes; we hope early adopters and benevolent skeptics will innovate upon them. This collection of articles provides foundational knowledge about open science practices and a foundation for educational psychology to lead the open science movement in the field of education. Instead of scattered individuals advocating for

discrete adjustments to the research process, we hope to be part of a unified generation of scholars that profoundly improves the scientific process in educational research for the benefit of practitioners and policymakers—not to mention the youth who they serve.

References

- Association for Psychological Science. (2017). *Open Practice Badges*. Retrieved July, 31 from <https://www.psychologicalscience.org/publications/badges>
- Blattman, C. (2016, March 27). Preregistration of clinical trials causes medicines to stop working. <https://chrisblattman.com/2016/03/01/13719/>
- Center for Open Science. (2020). *OSF Pre-registration template*. Retrieved February 18 from <https://osf.io/jea94/>
- Cohen, G. L., Garcia, J., Apfel, N., & Master, A. (2006). Reducing the racial achievement gap: A social-psychological intervention. *Science*, 313(5791), 1307-1310.
- Cron, L. (2012). *Wired for story: The writer's guide to using brain science to hook readers from the very first sentence* (1st ed.). Ten Speed Press.
- Dee, T. S. (2015, 2015 2016-03-15). Social identity and achievement gaps: Evidence from an affirmation intervention. *Journal of Research on Educational Effectiveness*, 8(2), 149-168.
- Dominus, S. (2017, Oct. 18). When the revolution came for Amy Cuddy. *The New York Times Magazine*. <https://www.nytimes.com/2017/10/18/magazine/when-the-revolution-came-for-amy-cuddy.html>
- Flake, J. K. (2021/this issue). Strengthening the foundation of educational psychology by integrating construct validation into open science reform. *Educational Psychologist*, 56(2), ###-###.
- Fleming, J. I., Wilson, S. E., Hart, S. A., Therrien, W. J., & Cook, B. G. (2021/this issue). Open accessibility in education research: Enhancing the credibility, equity, impact, and efficiency of research. *Educational Psychologist*, 56(2), ###-###.
- Gehlbach, H. (2015). Seven survey sins. *The Journal of Early Adolescence*, 35, 883-897. <https://doi.org/10.1177/0272431615578276>
- Gehlbach, H., Brinkworth, M. E., King, A. M., Hsu, L. M., McIntyre, J., & Rogers, T. (2016, 2016-02-17). Creating birds of similar feathers: Leveraging similarity to improve teacher–student relationships and academic achievement. *Journal of Educational Psychology*, 108(3), 342-352. <https://doi.org/dx.doi.org/10.1037/edu0000042>
- Gehlbach, H., & Robinson, C. D. (2018). Mitigating illusory results through preregistration in education. *Journal of Research on Educational Effectiveness*, 11(2), 296-315. <https://doi.org/10.1080/19345747.2017.1387950>
- Gelman, A., & Loken, E. (2014, November 2014 2016-03-17). The statistical crisis in science. *American Scientist*, 102(6), 460-465.
- Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on “Estimating the reproducibility of psychological science”. *Science*, 351(6277), 1037-1037. <https://doi.org/10.1126/science.aad7243>
- Hanselman, P., Rozek, C. S., Grigg, J., & Borman, G. D. (2016, 2016 Aug 08

- 2016-11-19). New evidence on self-affirmation effects and theorized sources of heterogeneity from large-scale replications. *Journal of Educational Psychology*.
<https://doi.org/http://dx.doi.org/10.1037/edu0000141>
- Harackiewicz, J. M., Canning, E. A., Tibbetts, Y., Giffen, C. J., Blair, S. S., Rouse, D. I., & Hyde, J. S. (2014, 2015-12-09). Closing the social class achievement gap for first-generation students in undergraduate biology. *Journal of Educational Psychology*, 106(2), 375-389.
<https://doi.org/http://dx.doi.org/10.1037/a0034679>
- Ioannidis, J. P. A. (2005, August 2005
 2016-01-12). Why most published research findings are false. *PLoS medicine*, 2(8), 1.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21(11), 1161-1166.
- 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. (2014). Investigating variation in replicability. *Social psychology*.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams Jr, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., & Bahník, Š. (2018). Many Labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443-490.
- Kristal, A. S., Whillans, A. V., Bazerman, M. H., Gino, F., Shu, L. L., Mazar, N., & Ariely, D. (2020). Signing at the beginning versus at the end does not decrease dishonesty. *Proceedings of the National Academy of Sciences*, 117(13), 7103-7107. <https://doi.org/10.1073/pnas.1911695117>
- Makel, M. C., & Plucker, J. A. (2014). Facts are more important than novelty: Replications in the education sciences. *Educational Researcher*, 43(6), 304-316.
- Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does “failure to replicate” really mean? *American Psychologist*, 70(6), 487-498.
<https://doi.org/10.1037/a0039400>
- Mellor, D. T. (2021/this issue). Changing norms in research culture to value transparency over novelty. *Educational Psychologist*, 56(2), ###-###.
- Miyake, A., Kost-Smith, L. E., Finkelstein, N. D., Pollock, S. J., Cohen, G. L., & Ito, T. A. (2010). Reducing the gender achievement gap in college science: A classroom study of values affirmation [Article]. *Science*, 330(6008), 1234-1237. <https://doi.org/10.1126/science.1195996>
- Muenks, K., Wigfield, A., Yang, J. S., & O'Neal, C. R. (2017). How true is grit? Assessing its relations to high school and college students' personality characteristics, self-regulation, engagement, and achievement. *Journal of Educational Psychology*, 109(5), 599-620.
<https://doi.org/10.1037/edu0000153>
- 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., Ishiyama, J., Karlan, D., Kraut, A., Lupia, A., Mabry, P., Madon, T. A., Malhotra, N., Mayo-Wilson, E., McNutt, M., Miguel, E., Paluck, E.

- L., Simonsohn, U., Soderberg, C., Spellman, B. A., Turitto, J., VandenBos, G., Vazire, S., Wagenmakers, E.-J., Wilson, R., & Yarkoni, T. (2015, 2015-07-29). Promoting an open research culture: Author guidelines for journals could help to promote transparency, openness, and reproducibility. *Science*, 348(6242), 1422-1425.
<https://doi.org/http://dx.doi.org/10.1126/science.aab2374>
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice*, 6(1), 101.
- Nuzzo, R. (2014, 2014 Feb 13). Statistical errors. *Nature*, 506(7487), 150-152.
<https://doi.org/http://dx.doi.org/10.1038/506150a>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science [Article]. *Science*, 349(6251), 716-711-716-718. <https://doi.org/10.1126/science.aac4716>
- Patall, E. A. (2021/this issue). Implications of the open science era for educational psychology research syntheses. *Educational Psychologist*, 56(2), ###-###.
- Phillips, D. C., Floden, R., Gehlbach, H., Lee, C., Warren Little, J., Maynard, R. A., Metz, M. H., Moje, E. B., Sandoval, W. A., Silverman, S., & Soudien, C. (2009). *The preparation of aspiring educational researchers in the empirical qualitative and quantitative traditions of social science: Methodological rigor, social and theoretical relevance, and more*
- Plucker, J. A., & Makel, M. C. (2021/this issue). Replication is important for educational psychology: Recent developments and key issues. *Educational Psychologist*, 56(2), ###-###.
- Reich, J. (2021/this issue). Preregistration and Registered Reports. *Educational Psychologist*, 56(2), ###-###.
- Reich, J., Gehlbach, H., & Albers, C. (2020). “Like Upgrading from a Typewriter to a Computer”: Registered reports in education research *AERA Open*. <https://doi.org/10.1177/2332858420917640>
- Robinson, C. D., Scott, W., & Gottfried, M. A. (2019, 2019/01/01). Taking it to the next level: A field experiment to improve instructor-student relationships in college. *AERA Open*, 5(1), 2332858419839707. <https://doi.org/10.1177/2332858419839707>
- Rosenthal, R. (1979, 2016-08-22). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86(3), 638-641. <https://doi.org/http://dx.doi.org/10.1037/0033-2909.86.3.638>
- Serra-Garcia, M., Hansen, K. T., & Gneezy, U. (2020). Can short psychological interventions affect educational performance? Revisiting the effect of self-affirmation interventions. *Psychological Science*, 31(7), 865-872. <https://doi.org/10.1177/0956797620923587>
- Sherman, D. K., Hartson, K. A., Binning, K. R., Purdie-Vaughns, V., Garcia, J., Taborsky-Barba, S., Tomassetti, S., Nussbaum, A. D., & Cohen, G. L. (2013). Deflecting the trajectory and changing the narrative: How self-affirmation affects academic performance and motivation under identity threat. *Journal of Personality and Social Psychology*, 104(4), 591-618.
<https://doi.org/10.1037/a0031495>

- 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. (2018, 2018/03/01). Introducing Advances in Methods and Practices in Psychological Science. *Advances in Methods and Practices in Psychological Science*, 1(1), 3-6. <https://doi.org/10.1177/2515245918757424>
- Slavin, R. E., & Lake, C. (2008). Effective programs in elementary mathematics: A best-evidence synthesis. *Review of Educational Research*, 78(3), 427-515.
- Slavin, R. E., Lake, C., Chambers, B., Cheung, A., & Davis, S. (2009). Effective reading programs for the elementary grades: A best-evidence synthesis. *Review of Educational Research*, 79(4), 1391-1466.
- Slavin, R. E., & Smith, D. (2009). The relationship between sample sizes and effect sizes in systematic reviews in education. *Educational Evaluation and Policy Analysis*, 31(4), 500-506. <https://doi.org/10.3102/0162373709352369>
- van der Zee, T., & Reich, J. (2018). Open education science. *AERA Open*, 4(3), 2332858418787466. <https://doi.org/10.1177/2332858418787466>
- Wentzel, K. R. (2021/this issue). Open science reforms: Strengths, challenges and future directions. *Educational Psychologist*, 56(2), ###-###.