Psychology's Replication Crisis and Clinical Psychological Science

Jennifer L. Tackett & Cassandra M. Brandes

Northwestern University

Kevin M. King

University of Washington

Kristian E. Markon

University of Iowa

Manuscript in press, *Annual Review of Clinical Psychology*

Author emails: jennifer.tackett@northwestern.edu;

cassandrabrandes2021@u.northwestern.edu; kingkm@uw.edu; kristian-markon@uiowa.edu

Corresponding author: Jennifer L. Tackett, Department of Psychology, 2029 Sheridan Rd,

Evanston, IL, 60208; jennifer.tackett@northwestern.edu

Abstract

Despite increasing interest in issues of replicability, open science, research transparency, and improving methods and practices in psychological science, the clinical psychology community has been slow to engage. This has been slowly shifting, and the authors of the present article hope to facilitate this emerging dialogue. We begin by examining some potential areas of weakness in clinical psychology in terms of methods, practices, and evidentiary base. We then discuss a select overview of solutions, tools, and current concerns of the reform movement from a clinical psychological science perspective. We examine areas of clinical science expertise (i.e., implementation science) that should be leveraged to inform open science and reform efforts. Finally, we reiterate the call to clinical psychologists to increase engagement with these efforts toward reform that should accelerate a more credible clinical psychological science.

Keywords: replication crisis; open science; clinical science; meta-psychology; clinical psychology; reproducibility

Psychology's Replication Crisis and Clinical Psychological Science

Interest in issues of replicability, open science, research transparency, and improving methods and practices has been at the forefront of scientific discussions in the psychological community in recent years, with renewed interest in these topics emerging around 2011 (e.g., Simmons et al. 2011). Although clinical psychology has been slow to engage in these ongoing efforts, a number of clinical scientists embarked on a paper intended to call clinical psychologists to the replicability conversation (Tackett et al. 2017) and movement within clinical psychology and related disciplines has been growing. For example, a recent call to establish an initial database of clinical psychologists interested in open science and replicability topics received an overwhelming response, with 200+ individuals identifying as clinical scientists signing up within a couple of weeks (<https://goo.gl/forms/tvAijds0as7peCdG2>). These individuals indicated a wide range of topical interests in open science, improving methodology, and replicability, and reflected a broad range of training levels, including graduate students and postdoctoral fellows (38.4%), assistant professors (26.4%), and associate /full professors (25.1%). Thus, despite slow emergence in the conversation thus far, substantial interest exists in the field and is mobilizing toward collaborative and constructive ways forward.

This slow and early growth remains in dire need of more focused and critical examinations of these topics from a clinical psychological lens. In the current review, we aim to contribute to this emerging conversation, hoping that clinical psychologists will continue to increase their involvement as we attempt to broaden this analysis and the potential solutions generated from them. We examine a few areas in clinical psychological science where replicability and credibility concerns are likely. First, we call for a more critical look at those portions of our existing research base (and the methods, practices, and dissemination

culture associated with them) where findings may not be replicable or robust. Then, we briefly review some of the current open science tools, proposals, and solutions that clinical psychological scientists should begin integrating into their own research and workflow. We next discuss some of the challenges in defining operationalization and replication in clinical psychological research, both at the theoretical and methodological levels. We then turn to a consideration of how the broader principles and framework of implementation science might be leveraged to increase uptake and retention of better methods and practices, including open science principles, across the field of clinical psychology. Finally, we end with some basic conclusions and next steps, including, and a plea for other clinical psychological scientists to join us in this work.

Identifying the Weak Spots in the Clinical Science Research Base

The fact that clinical psychology has been less involved in the replicability conversation to date should not be taken to suggest that we have minimal replicability concerns as a field. Early efforts from social, personality, and cognitive psychology have begun highlighting those areas where replicability problems may exist (e.g., Open Science Collaboration 2015). Differences in methodology and statistical approaches (e.g., an emphasis on effect size in clinical research over focal hypotheses examined via an NHST framework in experimental social and cognitive psychology) may point to potential differences in overall replication of findings, but until this is explicitly and comprehensively examined, it remains an empirical question (Tackett et al. 2017). One comprehensive recent effort to replicate comprehensive associations between personality traits and life outcomes demonstrated remarkably high replication rates (Soto 2018), perhaps in part because of the extensive body of research underlying the conceptualization and measurement of personality traits. To the extent that certain areas of clinical psychology similarly rely heavily on

correlational (or observational) estimates of descriptive associations, as is the common level of analysis in personality psychology, it is certainly possible that these areas will prove more robust.

Reardon and colleagues (2018) coded average sample sizes (which are but one aspect informing the statistical power in a study) in studies published in the two top APA clinical journals (*Journal of Abnormal Psychology* and *Journal of Consulting and Clinical Psychology*) in 2000, 2005, 2010, and 2015. This study found that sample sizes in top clinical psychological outlets were larger than those in top social psychological outlets and the field-wide flagship journal *Psychological Science* (Fraley & Vazire 2014; Reardon et al. 2018). Clinical sample sizes were more comparable to those in personality journals, although with higher associated impact factors for the clinical journals (i.e., clinical outlets appear to “reward” higher sample sizes based on correlations between journal impact factor and the clinical *N*-factor). Nonetheless, most samples were not large enough to detect small effect sizes (e.g., correlations in the .10-.20 range, which are likely to be most common for psychological research). Furthermore, from 2000-2015 change in sample size was not evidenced for JCCP (*Journal of Consulting and Clinical Psychology*) and was only beginning to trend up (toward higher sample sizes) for JAP (*Journal of Abnormal Psychology*). This study presents some mixed news for clinical psychology: although sample sizes have been higher than in some psychological subfields (e.g., social psychology) and clinical journals evidence a higher correlation between median sample size and impact factor, an improvement in sample sizes over the last two decades is not clearly evident across these prominent journals and our samples are often too small to detect small effect sizes, which may be a common magnitude for many psychological effects.

Historically, statistical power has been poor in psychology research (Cohen 1962a, Fraley & Vazire 2014), and this is especially true for resource-intensive areas. Two such areas in

clinical psychology include clinical neuroscience and treatment/intervention research. Research on the neuroscientific mechanisms underlying psychopathology have been increasingly prioritized by funding bodies, including the National Institute of Health (NIH), creating a heightened market for this type of research. These studies are costly in time and money, and thus it is likely unsurprising that they often rely on small samples to examine small- to medium-sized effects, resulting in poor statistical power to detect true effects (Button et al. 2013). In addition to the frequent use of small samples, statistical power of this research is often further minimized when researchers simultaneously use a large number of dependent variables and thus need to correct for multiple comparisons (Cremers et al. 2017). Of course, these problems also generalize to clinical neuroscience research as well, presenting a number of methodological considerations in clinical neuroscience research, all of which can impact replicability (Abram & DeYoung 2017). Studies of statistical power in treatment research have been similarly portentous. For example, comparative outcome trials for depression, researchers found that these studies had a maximum of 40% of the participants needed to detect a minimally clinically important effect size ($d=.24$; Cuijpers 2016). These initial investigations indicate that clinical neuroscience and treatment/intervention studies are likely areas for concern when examining clinical psychology's replicability weak spots.

Another potential area for concern in clinical psychology stems from diagnostic unreliability. Interrater reliability for diagnoses has historically been mixed, at best (Frances & Widiger 2012), and with the transition from DSM IV to DSM 5, reliability has only improved for some diagnoses (e.g., posttraumatic stress disorder; PTSD), while major problems remain for many common disorders (e.g., generalized anxiety and major depressive disorders; Freedman et al. 2013, Regier et al. 2013). In addition, categorical diagnoses have

poor predictive power for future diagnoses in comparison to dimensional alternatives (Kim & Eaton 2015), which is likely related (in part) to the loss of power that accompanies dichotomization (MacCallum et al. 2002). Given the primacy of DSM-defined diagnoses in the vast majority of clinical psychological research, such issues pose clear barriers to establishing a replicable and robust literature.

Yet another consideration is the impact of publication bias in clinical science, which has been discussed for decades (e.g., Meehl 1990). Across psychology, significant results are disproportionately published relative to null results, with 96% of papers showing significant effects (Bakker et al. 2012), elevated even in relation to the 80% average across the sciences (Fanelli 2010). In clinical psychology specifically, a consistent bias toward significant results has been documented in intervention studies, resulting in an overestimation of treatment effects for disorders such as depression, generalized anxiety disorder, and ADHD (Cuijpers et al. 2014, Cuijpers et al. 2010, Rapport et al. 2013). This is consistent with evidence that positive findings (via statistical significance) from clinical trials are published more often, and more quickly, than negative/null findings (Hopewell et al. 2009). Descriptive psychopathology research suffers from publication bias as well; one particularly high-profile example suggested that the association between the 5-HTTLPR serotonin transporter gene and unipolar depression was likely a false positive finding initially stoked by publication bias (Culverhouse et al. 2018, Duncan & Keller 2011). Similarly, one recent investigation found substantial publication bias favoring large positive effects to have drastically inflated reported associations between depression and mortality (Miloyan & Fried 2017). Thus, publication bias is likely a problem across various methodological approaches and research topics in clinical psychology.

A related concern is the use of questionable research practices, or researcher behavior that increases the likelihood of achieving desired (typically, statistically significant) results (Simmons et al. 2011). These practices include the selective reporting of desirable results, flexibility in data analytic strategies, and hypothesizing after results are known (HARKing), to name a few. In a study of QRP use among psychologists, John, Loewenstein, & Prelec (2012) found an estimated 91-94% of researchers across subfields reported engaging in at least one QRP, with this rate being lower among clinical researchers, as compared to researchers in other subfields of psychology. Nonetheless, there is evidence of QRPs from within the clinical literature, as well; for example, the selective reporting of results in clinical treatment studies is an established phenomenon often discussed as “allegiance effects” (Leichsenring et al. 2017). Allegiance effects, or the tendency of treatment studies to support the theoretical orientation of the author (e.g., cognitive-behavioral or psychodynamic therapies), have been documented for several decades (Munder et al. 2013). This selective reporting may contribute to null findings for comparative outcome studies, or “the Dodo bird verdict” in treatment research (Cuijpers 2016, Wampold et al. 1997). In a recent critical examination, Leichsenring and colleagues (2017) identified 13 biases they believed to affect replicability of psychotherapy research and, considering them each in turn, offered discussion around how better to counter such biases in this type of research. This serves as an exemplar for future areas of clinical psychological research to begin evaluating their own practices and biases.

To combat biased reporting, QRPs, and other problematic practices, one ameliorative effort has been to create guidelines and norms for registration of intervention trials in psychology (e.g., APA Publications and Communications Board Working Group on Journal Article Reporting Standards 2008). However, in clinical psychology, these guidelines have

not been required by most journals. Therefore, it is unsurprising that research shows that the vast majority of randomized controlled trials (RCTs) in clinical psychology are not pre-registered according to best practice recommendations. Cybulski, Mayo-Wilson, & Wilson (2016) showed that among RCTs published in 2013 in the 10 highest impact clinical psychology journals (N=165), only 25% of trials were registered prospectively (i.e., before data collection) and a mere 1% (N=2) were registered prospectively *and* defined all primary outcomes in their registration. A systematic comparison of registrations to publications has not been conducted in clinical psychology specifically, but extant evidence from health intervention research suggests anywhere between 40-62% of RCTs contain at least one unreported change to the research design following registration (Dwan et al. 2013). It is also worth noting that it is common to not report results of clinical trials; one estimate suggests that 63% of all clinical trials report their findings (<http://fdaaa.trialstracker.net/>). These results collectively suggest that clinical psychologists, like those in other areas, likely engage in practices that increase the likelihood of replicability problems in the long term. Further, this evidence indicates a need to establish better norms and reinforcement for tools that combat bias and increase openness in clinical science.

Meta-analysis and its Relevant History in Clinical Psychology

Historically, clinical psychology as a field has been central to the study of scientific replicability, not only within psychology and the behavioral sciences, but in the sciences more broadly. By pooling published effects, including many effects from many studies (ideally), meta-analysis can—for an unbiased literature—provide a more accurate estimate of both effect size and effect size heterogeneity. Modern meta-analysis, for example, has its origins in the meta-analytic study of clinical phenomena—Glass (1976) in fact coined the term “meta-analysis” in the context of research into psychotherapy efficacy (Glass 2015,

Smith & Glass 1976). Statistical integration of results from multiple studies is much older (Pearson 1904), but work at that time by Glass and colleagues, as well as others (Shadish & Lecy 2015), helped spark modern meta-analysis by leveraging effect sizes to integrate findings across heterogeneous research designs (Gurevitch et al. 2018). Clinical psychology was central to establishing the lingua franca of replicability studies, in terms of effect size, power, and related considerations as they pertain to replicability and meta-scientific concerns.

Parallels with the current replicability crisis are informative. Much like now, concerns—as well as methods for addressing those concerns—were being driven by skepticism about specific scientific claims, in that case, claims of psychotherapy efficacy (Shadish & Lecy 2015). Various lines of evidence have confirmed longstanding concerns that the psychotherapy research literature itself has been affected by publication bias (e.g., Cuijper et al., 2010; Driessen et al., 2015; Sohn 1996). This pattern of findings regarding publication bias in the psychotherapy literature suggests that research on replicability is coming full circle. Whether or not prospective registered replication projects will follow remains to be seen. On the one hand, psychotherapy research has, perhaps, one of the most thorough meta-analytic evidence bases in all of science; the existence of NIH and other agency records has allowed for retrospective registered replication analyses; and registered replication studies of clinical phenomena are expensive in multiple ways that studies of nonclinical phenomena are not. On the other hand, registered replication projects have followed other meta-analytic demonstrations of publication bias in other fields, and there have been counterarguments regarding the magnitude of psychotherapy publication bias (Niemeyer et al. 2013), as well as observations that other types of bias (e.g., measurement error) might counter publication bias *per se* (Staines & Cleland 2007).

Clinical Assessment Gaps Between Research and the Field

Another area of concern for clinical psychology research is in assessment gaps between research and applied settings and extensive heterogeneity in forms of assessment across studies. Problems with defining optimal practices for diagnoses notwithstanding (as discussed previously), ideal, structured assessment practices are rarely, if ever, applied in the same fashion between the laboratory and the clinic due to practical limitations (Reardon et al. 2017). Further, varying assessment approaches (e.g., prospective versus retrospective) have been shown to drastically alter lifetime diagnostic prevalence estimates even between academic studies exclusively (Moffitt et al. 2010). These problems likely contribute to discrepancies between prevalence rates for many disorders between field and academic settings (Garb 2005); findings which imply non-generalizability and non-replication of outcome and inclusion criteria across contexts.

Assessment gaps between clinical and academic research studies are evidenced in differential prevalence estimates between studies conducted in these two settings. The distance of gaps between field and epidemiological or academic studies varies widely by disorder. Some research suggests that primary care practitioners underdiagnose some disorders (e.g., early personality disorder) when comparing field prevalence rates to those obtained from epidemiological surveys (Conway et al. 2017, Garb 2005). Research has also suggested that prevalence rates are inflated for other disorders (e.g., pediatric bipolar disorder) in the clinic relative to the lab due to an overreliance on potentially biased “clinical judgment” over structured assessment, malingering, and pharmaceutical companies creating incentives for practitioners to assign diagnoses (Frances & Widiger 2012, Garb 2005, Jenkins et al. 2011). On the other hand, field trials and epidemiological research have also been criticized for not capturing mid-range diagnoses of the kind that generally present in a

primary care setting, so the potential for underestimation of population prevalence rates in research is also a possible contributor to these gaps (Jones 2012). While many factors contributing to these discrepancies (e.g., malingering) vary by disorder, one primary commonality is the lack of structured assessment use in primary care contexts. As discussed, this can lead to both under- and overdiagnosis, though its full implementation is impeded by sizeable practical constraints. In evaluating the replicability of various effects in clinical psychology, these methodological features of studies warrant close consideration.

The degree of structure in the assessment of psychopathology is not the only measurement factor that may contribute to heterogeneity in disorder prevalence estimates; methodological differences between studies may also factor into assessment discrepancies. One meta-analysis of prevalence rates for a variety of childhood disorders found that estimates varied substantially as a function of sampling frame (e.g., schools or households), sample representativeness, and the type of diagnostic interview (Polanczyk et al. 2015). Other researchers have found vastly different disorder lifetime prevalence rates as a function of retrospective versus prospective study designs, with prospective rates approximately doubling retrospective rates (Moffitt et al. 2010). This substantial method variance has immense implications for the replicability of clinical science, as prospective studies are extremely time- and resource-intensive. Therefore, much descriptive psychopathology research is conducted retrospectively, possibly yielding estimates that may in turn face greater replication problems.

Generalizability of RCTs to the Field/Community Settings

Although RCTs provide the gold standard of evidence for treatment efficacy, there is often little attention paid to how the design of RCTs impacts the external validity or generalizability of findings from RCTs. As Shadish, Cook, & Campbell noted, “*readers of*

experimental results are rarely concerned with what happened in that particular, past, local study. Rather, they usually aim to learn either about theoretical constructs of interest or about a larger policy" (2002, p. 18). The essential irony of many RCTs designed to test efficacy is that maximizing internal validity (the degree to which a treatment is known to cause changes in the outcome) often comes at the cost of external validity (the degree to which effects may be generalized to broader populations and contexts; Shadish et al. 2002). There has long been concern about the differences between populations eligible for RCTs and the general patient population (Rothwell 2005) and mental health intervention research (Stuart et al. 2015). One review of exclusion criteria for studies of substance use disorders (SUDs) estimated that between 63% and 68% of individuals with SUDs would be excluded from RCTs for reasons such as prior treatment, low motivation, other substance use or medical problems (Moberg & Humphreys 2017). On the other hand, evidence for the treatment of many other disorders contradicts this claim. For example, a few studies have shown that most patients seen in outpatient community settings would be eligible for at least 1 published RCT (Stirman et al. 2003; 2005). These and other studies from inpatient settings have suggested that ineligibility of patients in community setting is more often observed due to disorders that are in partial remission, or are more mild compared with those observed in RCTs (Stirman et al. 2003; 2005, von Wolff et al. 2014), and some research has suggested that the most common exclusion criteria seen in RCTs for depression are unlikely to influence the generalizability of their effects (van der Lem et al. 2012).

However, the generalizability of RCTs is not only influenced by client characteristics (i.e. how well they match the population for whom there is evidence of efficacy) but also by the degree to which the characteristics of a treatment delivered in an RCT matches those in applied settings. Here, the gap between research and practice seems to be much wider. In

one study of the efficacy of treatments for children studied in usual care settings, reporting that the typical effect size of treatment was nearly half or less than what was reported in prior meta-analyses of RCTs (Weisz et al. 1995). This study also noted substantial differences between the context of RCT delivery and community settings. RCTs frequently rely on highly trained therapists with homogenous (and intensive) training in the treatment, receiving high levels of supervision from treatment experts, and carry relatively low caseloads of patients with similar disorders, while psychotherapists in the community carry a high caseload of patients with a wide variety of diagnosis and get relatively little supervision (Weisz et al. 1995) that is provided at relatively low ratios (~an hour per week) and rarely focuses on evidence based interventions (Dorsey et al. 2013; 2017).

Merging Current Proposals for Reform with Clinical Psychology Research

Since the most recent instantiation of psychology's replication crisis in 2011, the rate at which tools and incentives for increasing replicability have developed has been staggering. However, the majority of these efforts have been initiated by psychologists within a limited range of subfields, and thus, concerns have been raised as to their applicability to areas outside of social, cognitive, and personality psychology (Finkel et al. 2015, Tackett et al. 2017). In addition to practical tensions in attempts to generalize tools and solutions to other areas of psychological research, more fundamental issues raise additional concerns about generalizability. For example, fundamental questions around defining and operationalizing replication may not translate directly to much clinical psychological research, and similar statistical and analytic issues may necessitate broader discussion to better capture a wide variety of psychological science, including that from the clinical domain. We briefly discuss these three issues in turn: 1) applying proposed tools and solutions to clinical psychological

research, 2) defining and operationalizing replication in clinical psychological research, and 3) extending relevant statistical debates and concerns to clinical psychology.

Proposed Solutions and Tools for Reform

Here we briefly describe five non-exclusive (and non-exhaustive) routes for increasing replicability including (1) open materials, (2) open data, (3) (pre)registration, (4) Registered Reports, and (5) multisite collaborations for both replications and original research. We then provide an overview of various extant incentives for engaging with these practices and how they relate to clinical science. This list is only a very brief and selected overview of current efforts being made toward reform, but will perhaps serve as a starting point for those interested in enhancing openness, transparency, and reproducibility in their research.

Open materials are most typically made available when study authors post all digitally shareable protocols, measures, and other applicable study tools to a publicly available online repository. Open materials allow for more thorough evaluation of a study as well as facilitating standardization and replicability in future efforts by those using the same materials. However, clinical psychology materials are, on average, not as readily shareable as those in some other subfields. First, monetization is a much more central issue to clinical psychological assessments than measures in other subfields (Yates & Taub 2003). Second, public access to assessments for psychopathology is especially ethically problematic due to increased likelihood of malingering and diagnoses by unqualified persons, and test security is explicitly listed under APA Ethical Principles Section 9.11 (American Psychological Association 2017). Third, measure integrity can be compromised if assessment materials are made available to those whom they are designed to assess. Neuropsychological tests, for example, are often tightly controlled in order to protect against coaching and practice effects

(Calamia et al. 2012), and re-norming these materials with sufficient power is enormously resource-intensive. However, many clinical psychology materials are arguably no different from those in other subfields in terms of sharing. Experimental paradigms such as those designed to evoke fear or reward responses, scripts for parent-child interaction tasks, and instructions for daily diary completion are all examples of materials that do not carry the ethical and practical weight of measures such as those described above. As such, open materials in clinical psychology are a nuanced issue, and one that warrants further attention as clinical psychology joins the reproducibility conversation.

Open data involves posting raw study data to an online repository that can be accessed by individuals outside of the study team. The term “open data” includes both fully publicly available data and data posted to trusted third-party repositories that restrict data access to researchers who meet criteria for demonstrating proper handling of sensitive data, or “protected access” (Blohowiak et al. 2018). Open data is intended to increase data availability for verification and confidence in a project’s reported results, especially as data “available on request” is most often practically inaccessible (Wicherts et al. 2006). One major concern with sharing clinical data is identifiability, as clinical psychology data generally include sensitive health and legal information (e.g., endorsement of criminal behavior), low base-rate conditions (e.g., schizophrenia), and rare populations (e.g., individuals aged 89+). Aside from trusted data repositories, another proposed solution is application of the “safe harbor method,” or a de-identification method originally developed for medical data that substantially limits privacy risks by removing 18 features such as precise ages over 89 years, or exact ZIP codes in areas with populations of less than 20,000 (Walsh et al. 2018). In addition, clinical psychology can benefit by learning from efforts from other areas of social

science research that are grappling with similar concerns (e.g., relationships research on sensitive topics; Joel et al. 2018).

A second major concern about open data among clinical psychology researchers is the incentive for conducting time- and resource-intensive data collection when open data is normative. While this problem (i.e., the “Tragedy of the Commons”) is applicable to all of psychology, it is even more salient in the context of data that require years, even decades, of effort and funding on the part of investigators, while incentives generally favor researchers who publish papers over those who collect data (or perhaps, favor those who get grants over those who publish papers over those who collect data). Some proposed solutions to this problem include data embargoes, where researchers can publish data only after a prespecified period of time in which they plan to publish on the dataset themselves, and venues allowing for the publication of open datasets so that researchers may benefit from citations when the data are used (Munafò et al. 2017).

Preregistration, in a strict sense, entails public posting of time-stamped records of study design, hypotheses, outcomes, and data analysis pipelines prior to the execution of a study, including prior to the start of data analysis. Preregistration is intended to clearly delineate which findings in a study were predicted a priori, and thus clarify the diagnosticity of inferential statistical tests (Nosek et al. 2018). Preregistration of primary and secondary outcomes is now mandated for clinical trials by the United States government, as well as many medical journals (see also International Committee of Medical Journal Editors; De Angelis et al. 2004). Despite this push within the broader medical field, clinical psychology has been slow to adopt preregistration. This may be due to a focus on “strict” preregistration in psychology’s replicability conversation – or one in which all aspects of a study can be neatly delineated prior to study commencement. Strict pre-registration is a standard that is

not achievable for much research outside of lab-based studies with undergraduates, which may result in researchers opting out completely (Finkel et al. 2015, Tackett et al. 2018). The lack of uptake in clinical psychology specifically may also be influenced by the more common usage of archival datasets, protracted nature of data collection and hypothesis development in larger and longer studies, and unpredictable sampling due to base rates, difficulties with recruitment, and additional research challenges encountered in clinical research (see Tackett et al. 2017 for a review of these obstacles). There is a growing recognition of these problems within the replicability conversation, however, and a call to avoid throwing out the proverbial baby with the bathwater (Tackett et al. 2018). Indeed, we urge clinical psychological researchers to consider various options along the registration continuum, as many options exist beyond strict pre-registration. Recent recommendations can be summed up by this – be as transparent as is feasible, as some basis for evaluating credibility is better than none at all.

Registered Reports (RRs) are a type of empirical article that undergoes peer review at the design phase of the project, with in-principle acceptance awarded before results are known by authors and reviewers to ensure publication of a project regardless of whether findings support the hypotheses or not. This initiative is intended to reduce publication bias and the selective reporting of significant results. Registered Reports are essentially detailed preregistrations of studies that are then evaluated by peer reviewers, updated based on reviewer feedback, and in-principle accepted such that, barring deviations from the agreed-upon study design and analytic plan or author withdrawal, the project will be published once the study is completed and written up. First implemented in 2013, Registered Reports are currently offered as an article type in 108 journals across psychology, neuroscience, and medicine. At the time of writing of this review, very few clinical psychology journals have yet

confirmed RRs as an accepted article type, however. Barriers to RRs for clinical psychologists overlap with those for preregistration, as well as those for open materials and data, as these are often requirements for RRs. The use of extant datasets has also been raised as a concern for the interface between RRs in clinical psychology, but over 50% of journals that offer RRs accept submissions based on existing data (<https://cos.io/rr/>). RRs have been cited as ideal solutions to many problems with replicability, but this is an initiative that requires further reconciliation with the needs of clinical science in the future.

Multisite collaborations are projects initiated and maintained by researchers at multiple institutions, both for the purposes of conducting large-scale replications as well as original research. Multisite collaborations are intended to increase diversity and statistical power in research – especially among populations that are hard to recruit – by pooling resources between labs. The last several years has seen the rise of many impactful multisite collaborations, including the Human Connectome Project (Van Essen et al. 2013), the Psychological Science Accelerator (Moshontz et al. 2018), and StudySwap (Chartier & McCarthy 2018). Still other efforts exist for the purpose of conducting large-scale replications, including the Many Labs (Klein et al. 2014) and ManyBabies (Frank et al. 2017) projects, as well as Registered Replication Reports, with an article type similar to RRs and focused on multisite replications

(<https://www.psychologicalscience.org/publications/replication>). Multisite collaboration is especially important for resource intensive research, and it is therefore optimally applicable to clinical science. While multisite collaborations have become commonplace for original research in clinical psychology (for example, in neuroimaging research), large-scale replication efforts for clinical research findings have not yet taken hold. Given the obstacles that clinical psychology faces in doing exact replications (see Tackett et al. 2017),

collaborations present rich opportunities to “bake in” replication in an efficient manner. However, without structural incentives for engaging in these efforts, their uptake is likely to be slow.

All researchers are motivated by structural incentives and norms, and changes in these factors are likely to have an effect on clinical psychology researchers’ engagement with reproducibility and open science going forward (Lilienfeld 2017). In perhaps one of the most powerful shifts in structural incentives, the National Institutes of Health (NIH) announced that a discussion of “Rigor and Reproducibility” would be required in the Significance/Approach sections of NIH grant applications submitted after May 25, 2016, with plans to apply these standards to institutional training grants, institutional career development grants, and individual fellowships imminently. Although this change is promising, it is unclear how it will be operationalized and what direct effect it will have on scientific practices.

Beyond funding structures, journals – including several high impact clinical science journals – have also begun adopting standards of openness including the Transparency and Openness Promotion Guidelines (TOP; <https://cos.io/top/>). The TOP Guidelines are gradients of enforcement of eight open science standards/practices (e.g., preregistration, replication, code transparency), with four tiers of stringency. Some journals have also begun implementing a “badge” system, in which individual articles are displayed with acknowledgements of engagement with open science practices including open data, open materials, and preregistration. At least two clinical psychology journals (*Clinical Psychological Science* and *Collabra: Clinical Psychology*) have adopted badges promoting open science practices, at the time of this writing. Since implementing badges in psychology journals, data sharing has risen from approximately 3% to 40% of articles in participating journals (with no change

in non-participants), and open materials have likewise increased, though to a lesser extent (Kidwell et al. 2016). Beyond this, open access to clinical psychological research has expanded in the last decade, with NIH announcing a mandate on public availability of manuscripts published on NIH-funded research. Clinical psychology researchers have been relatively slow to adopt reproducibility and open science, however, structural incentives and norms enforced by funding agencies and clinical psychology journals may advance this engagement substantially. As such, it is imperative that clinical scientists join the conversation on developing tools, incentives, and norms around these issues to optimize their applicability and overcome obstacles to implementation.

Defining and Operationalizing Replication in Clinical Psychological Research

An ongoing challenge in the replication crisis is how best to define replicability. For example, what aspects of a design must be replicated in order to consider something a “direct” versus “conceptual” replication (Zwaan et al. 2018)? Even researchers at the forefront of the reform movement have struggled to answer this question when working primarily from an experimental cognitive/social psychological framework, where constraining direct replications are likely much more feasible. Multiple barriers emerge that prevent direct replication in clinical psychology (Tackett et al. 2017), and a truly direct replication may never be possible (Tackett & McShane 2018). Some critical examples of barriers in clinical psychological research are the variable nature of clinical samples across sites and labs, the changing nature of diagnostic constructs which result in evolving construct definitions and measurement, and statistical approaches (i.e., prioritization of effect size estimation over NHST) that require more nuanced approaches to defining replication than dichotomous success/failure decisions allow (Tackett et al. 2017, Tackett & McShane 2018). Highly relevant and generalizable conversations are occurring within

epidemiology, with parallels in terms of primary concerns and challenges for solutions (Lash et al. 2018).

As the broader reform movement considers the best way to operationalize replication success and failure in certain types of psychological research, the clinical psychology community must directly address the question ourselves: by what criteria should we define a replication attempt? And how do we operationalize the outcomes of these attempts? Given the nature of much clinical psychological research, we may be well served by considering replication as a latent construct, and attempt to work toward construct validation efforts—a space in which many clinical psychology researchers are actually quite comfortable (Cronbach & Meehl 1955). This will require active and evolving attempts to define the construct, and an iterative process at measurement in which measurement attempts and construct refinement occur in an ongoing feedback loop. Such an approach would also likely be best served by more complex and nuanced measurement approaches, including, for example, the use of intervals or reaction ranges in place of dichotomous (e.g., success/failure) or point estimates, as well as the use of hierarchical/meta-analytic models to account for effects across studies, sites, and samples (Tackett & McShane 2018).

One potential starting point is a comprehensive review of previous attempts in the clinical psychology literature, which certainly exist (e.g., see the thoughtful analysis by Sher et al. 2011, to recover the “cat’s cradle” pattern in latent trajectory analysis of substance use while varying aspects of study design and features). A thorough review of such examples already in the clinical literature would be helpful to highlight those ways clinical researchers have already been tackling this underlying question, and allow us to move toward a more systematic and structured approach that might be adopted across the field. In general, we see

this as an exciting and incredibly important area for further consideration as clinical psychology becomes more and more engaged with replication and other reform efforts.

Extending Statistical Concerns to Clinical Psychological Research

Contrasts between replication paradigms in other areas and how those paradigms might or not apply to clinical psychology extend to statistical methods that are used.

Differences in the types of questions and designs that constitute clinical psychological science lead to different replicability concerns, which in turn lead to implications for how replication is conceived of and operationalized in clinical psychology versus other areas.

For instance, in contemporary discussion of replicability, a great deal of focus has been placed on p-values, especially in terms of two general concerns: how they are interpreted vis-a-vis an observed sample or result, and their vulnerability to manipulation in the form of p-hacking. The former concern, about p-value interpretation, derives from the insufficiency of p-values as an index of the replicability or veridicality of a given result (Colquhoun 2017)—p-values represent $P(X|h)$, the probability of the observed dataset or result X given a hypothesis of interest h , when what is needed is an estimate of the probability of a hypothesis, given an observed dataset or result, $P(h|X)$. The latter concern about p-hacking arguably ultimately derives in part from a focus in the literature on null hypothesis significance testing (NHST) in its various forms.

These two general concerns about p-values apply to nearly, if not all, areas of science, including clinical psychology. However, certain characteristics of clinical psychological research relative to other areas of psychology lend additional emphasis to these concerns. For instance, in clinical psychology the ubiquity of observational designs and focus on applied utility lends additional weight to the arguments for moving away from NHST and toward effect size estimation; when faced with large, highly powered epidemiological

samples, or comparisons of novel versus existing treatments, significance per se might be less important than effect size and cost-benefit considerations (McShane et al. 2018).

Similarly, the heterogeneity of clinical populations and extremely low base rates of many clinical phenomena amplify concerns about use of p-values in making inferences about findings, as $P(h | X)$ depends on the probabilities of the hypotheses being examined in the population at hand.

At the danger of overgeneralizing, it might be argued that clinical research designs are often more expensive than the types of designs that have dominated psychological replicability discussions in other areas, due to participant recruitment costs, large sample sizes, and other considerations (Tackett et al. 2017). As such, some approaches to studying replicability, such as collaborative pre-registered replications, might be more difficult or expensive in clinical psychology than in other areas of psychology. This may increase the weight placed on more meta-observational approaches to replicability, especially meta-analysis, which has always concerned itself fundamentally with questions about replicability and generalizability of results. The field of clinical psychology and psychiatry is likely to benefit substantially from improvements in meta-analysis, including a better understanding of how inferences derived from traditional meta-analytic approaches coincide with, or can be better calibrated with, those of collaborative preregistration approaches, which might not always be feasible.

The expense of clinical designs also has implications for understanding how publication bias might occur in clinical psychology versus other fields. Although participants might be relatively expensive either in number or in recruitment costs, measures are sometimes still relatively cheap to administer, which might give rise to "measure-hacking," whereby results are manipulated by selectively reporting results for specific measures or

variants of measures. The possibility of this phenomenon gives weight to the importance of developing and using well-validated and justified measures (Marshall et al. 2000). It also underscores the importance of understanding the psychometric structures of clinical constructs, as these structures point to potential similarities and points of convergent replication in designs. Even if findings cannot be replicated across participants, they might be replicated across different measures of the same constructs, or across different constructs that operate equivalently within a theory (McShane et al. 2018). Traditional psychometric considerations, such as of reliability, convergent and discriminant validity, and psychometric structure, help define the contours of replicability expectations, in terms of what measures a finding might or might not be expected to replicate across (Markon 2015). The next stages require more specific efforts to outline quantitative criteria—more likely in the form of likely intervals rather than precise point estimates—for replicability and psychometric standards from a construct validity perspective, which will more accurately capture much research in clinical psychology (Tackett et al. 2018).

The types of topics pursued in clinical research shape questions about replicability in other ways. Structural modeling is an important focus of research in clinical science, for example, but it is sometimes unclear how to define replicability of structural findings. Schieber and colleagues (2017) have shown that a given definition of structural equivalence can result in causally distinct structures being identified as equivalent, and that the set of structures identified as equivalent can vary depending on the definition used. This ambiguity in whether two structures are deemed equivalent derives from how different patterns in causal paths can be weighted; to the extent that different weighting methods are employed, different conclusions can be drawn about the extent to which one structure is replicating another. Moreover, a given well fitting structural model only represents one of many

possible well fitting models, and should not be taken as confirmation of any one model as empirical truth (Tomarken & Waller 2003).

Identifying replicable causal structures can also be difficult regardless of how replicability is defined, however. One emerging problem is in identifying replicable or generalizable details in causal structures, where a number of variables might be correlated due to multiple mechanisms simultaneously, including direct effects, indirect effects, dynamic processes, or unmeasured third variables. In much the same way that multicollinearity in regression creates challenges in distinguishing unique from shared components of correlated predictive effects, in causal networks or systems, it becomes difficult to isolate direct causal relationships between variables controlling for other direct, indirect, and latent effects in the broader causal system. Although some evidence suggests that broad features of causal networks might generalize across samples (Fried et al. 2018), other research has suggested that specific causal pathways, isolated from the broader structural features, are often poorly replicable (Bulteel et al. 2018, Forbes et al. 2017). These difficulties in identification of generalizable structural pathways create challenges in interpretation of causal systems above and beyond broad superordinate features. Research is needed to identify mechanisms for increasing the statistical power with which unique causal pathways can be identified.

Leveraging Intervention (Implementation) Science to Improve Credibility in Clinical Psychological Research (and Beyond)

In addition to considering areas where the clinical psychological evidence base may be weak, and ways in which we can bring the open science/replicability conversation to clinical psychology, it is time to think more broadly about the ways in which bringing clinical psychology to the conversation can help shape and facilitate it. One highly salient example is

the area of intervention/implementation science, which provides a broad framework aimed at better understanding how to get people to change – ways in which individuals (and groups/organizations) initially adopt new strategies, and how these adaptive changes can be maintained. In the following section, we discuss in greater depth the ways in which the open science/replicability conversation may benefit from leveraging the existing evidence base on implementation science.

Diffusion and dissemination of best practices are likely to be ineffective

There are several lines of evidence that suggest that there is a substantial gap between the optimal and actual practice of research methods in psychological science. First, although some tutorials on research methods are among the most highly cited manuscripts in the field (Sharpe 2013, Sternberg 1992), most articles introducing new statistical methods are not highly cited even years after their publication (Altman & Goodman 1994). The majority of articles in applied psychology cite one or zero quantitative articles (Mills et al. 2010). Despite the popularity of tutorials, applied researchers exhibit fundamental misunderstandings of methodological principles and statistical ideas, such as p-values and null-hypothesis significance tests (Nickerson 2000), Cronbach’s alpha (Schmitt 1996, Sijtsma 2009), confidence intervals (Belia et al. 2005, Cumming et al. 2004), or misunderstand the assumptions and interpretation of regression (Williams et al. 2013) and analysis of covariance (Miller & Chapman 2001, Westfall & Yarkoni 2016). Indeed, statistical “myths” are common enough to have books dedicated to correcting them (Lance & Vandenberg 2008).

Instead of applying statistical thinking to data analyses (Chance 2002), researchers rely on heuristics and rules of thumb to guide many decisions about data analysis (“mindless rituals”; Gigerenzer 2004, Gigerenzer & Gaissmaier 2011), regardless of the evidence in support of such heuristics or the degree to which they have been shown to be misleading

(Greenland et al. 2016). For example, researchers have frequently used rules of thumb to determine sample size (Cohen 1990), model fit (Marsh et al. 2004), or test reliability (Sijsma 2009). Gigerenzer (2004) described the “mindless ritual” of null-hypothesis significance testing, where researchers begin with a null hypotheses (such as no correlation or no mean differences), fail to specify the predictions of an alternative hypothesis, and use $p < .05$ to reject the null and accept whatever version of their hypothesis they favor. Recent survey data indicated that, when psychology researchers used rules of thumb to estimate necessary sample size to examine a typically sized effect, they routinely overestimated power and underestimated the required sample size (Bakker et al. 2016). Rituals, rules and magic numbers can be helpful checks against common errors in human judgment (Kahneman & Tversky 1974), but when they are over-used in a research context where developments in the quantitative literature are largely ignored, they can become quickly outdated.

There are frequent errors in the application and reporting of statistical models. Several studies have estimated that between 10% and 20% of the psychological literature reports incorrect significance tests (Bakker & Wicherts 2011, Berle & Starcevic 2007, García-Berthou & Alcaraz 2004, Nuijten et al. 2016), frequently in favor of the researcher’s hypotheses (Bakker & Wicherts 2011). Another recent review estimated that nearly half of articles in social psychology report means or standard deviations that are inconsistent with the sample size (Brown & Heathers 2017). The methodological literature is rife with reviews of common errors or misapplications of research methods. Researchers struggle to properly adjust for covariates (Lord 1967, Miller & Chapman 2001, Westfall & Yarkoni 2016), use power analysis to inform study design (Cohen 1962b, Sedlmeier & Gigerenzer 1989), establish construct validity (Campbell & Fiske 1959, Fiske & Campbell 1992), test for mediation (Shrout & Bolger 2002, Zhao et al. 2010) test exploratory factor models (Fabrigar

et al. 1999), or judge the fit of structural equation models (Jackson et al. 2009, McNeish et al. 2018). In short, if a method of analysis is popular, it is likely that one can find a review of frequent reporting and analysis errors of its application.

The gap between what is optimal scientific practice and what is standard practice likely arises in part because there is no systematic model for the implementation of best practices in research methods. To date, efforts to improve scientific practice have focused on diffusion, dissemination, and to a lesser extent, implementation. Diffusion, or the passive delivery of information (such as publishing journal articles in methodology journals), is likely to be the least effective method, given how few methodological papers are cited in substantive articles (Mills et al. 2010). Dissemination, or targeted knowledge delivery, may be represented in methodological tutorials for applied audiences (i.e. Grimm 2007, King et al. 2017), which are highly cited (Sharpe 2013), or by workshops and pre-conferences. These relatively passive methods of increasing knowledge may inadvertently contribute to the problems with the application or interpretation of statistical models described above. Because innovation is highly incentivized, it may be that new statistical methods are rapidly adopted with little attention to the marginal improvement in inferences they might provide compared to existing methods, or the degree to which novel methods increase the risk of error. For example, Bauer (2007) described the rapid adoption of growth mixture models despite the numerous problems with their application and interpretation (such as mis-identification of classes, dependency of class solutions on study design; see also King et al., 2017; Sher et al., 2011).

Finally, some structural efforts at implementation (or purposeful and targeted efforts at changing practices) have been attempted, where funding agencies, practice organizations (such as APA or APS), or journals provide recommendations or requirements around the

reporting of scientific studies. For example, the APA Publication manual provides standards for reporting many basic statistics (American Psychological Association 2010), editorial boards have attempted to mandate or incentivize statistical reporting practices (Finch et al. 2004, Ioannidis 2018, Kidwell et al. 2016). However, without systematic change at the organizational level, structural efforts will only last as long as the reformers are in charge (Finch et al. 2004).

Perhaps the expected solution would be to improve graduate and post-graduate training in research methods. However, evidence suggests that there is an increasing gap between the training that doctoral students receive and the research methods that are commonly applied in current research. The typical graduate statistical training sequence has been repeatedly noted to cover the basics of experimental and correlational analysis, requiring a mean of 1.2 semesters of statistics training, and doing a poor job of covering more advanced topics (such as latent variables, measurement, or quasi-experimental design (Aiken et al. 1990; 2008). Although there is a proliferation of workshops available for research methods, there is no evidence that they actually improve participants' methodological skill in a way that substantively improves their research practices. Adult and professional education research suggests that the impact of workshops is limited (Lyon et al. 2011). Passive learning is well known to be ineffective at producing changes in practice or skill mastery (Beidas & Kendall 2010, Herschell et al. 2010). Evidence suggests that learning is influenced not only by the specific practices of teaching, but also by the multiple contexts in which learning occurs (Bransford et al. 2000), such that optimal learning in a new method will occur by building on existing knowledge and practice of new skills in the context in which they will be applied (Lovett & Greenhouse 2000). Coupled with the evidence for a shortage of quantitative psychologists (Aiken et al. 2008, Golinski & Cribbie 2009), it is

increasingly likely that applied researchers will be called upon to use or evaluate research methods for which they lack expertise.

Leveraging The Advent of Implementation Science in Healthcare For Open Science Reform

The notion of a gap between typical practice and optimal practice has been widely studied in other fields, perhaps most notably in the area of healthcare (Chambers et al. 2013). Studies of medical interventions suggest it can take decades for evidence based practices (EBPs) to appear in typical practice (Contopoulos-Ioannidis et al. 2008, Morris et al. 2011). For example, one study estimated that around 10% of therapists used manuals for EBPs in community mental health settings (Becker et al. 2013). Other work has suggested that although many therapists apply elements of EBPs in standard practice, they are used in doses too low to have therapeutic effect (Garland et al. 2010). This suggests that when a typical individual obtains treatment, they are either unlikely to receive an evidence based treatment (Zima et al. 2005), or what they do receive will be ineffective. This is at least in part because while the evidence base is developed using therapists with low caseloads who are closely supervised and treating homogenous caseloads of patients, clinicians in the community have heterogenous and high caseloads and receive low levels of supervision (Weisz et al. 1995). It is difficult for therapists to choose the appropriate EBP for each case from among the potential hundreds available, and it is likely impossible that any single therapist may be able to master a variety of EBPs for all clients (Chorpita et al. 2007).

Implementation science is a relatively new field in healthcare that offers frameworks, methods, and outcomes that could be usefully applied to enhance the quality of research methods. An implementation science perspective is grounded in the evidence that diffusion and dissemination of research evidence and best practices does little to change behaviors or

systems that maintain the status quo, at least in part because they only attempt to increase knowledge about better practices, which only exert one influence on behavior change (Joyce & Showers 2002). Indeed, the implementation science perspective recognizes that implementation is multiply determined by characteristics of the intervention, the individual, and the interpersonal, organizational, and macro-level context of an individual (Damschroder et al. 2009). Successful implementation is characterized by changes in not only the final outcome (such as more reliable and valid research findings, or improved patient health outcomes), but by the outcomes of implementation strategies, such as increases in fidelity, widespread adoption, and increases in perceived acceptability, feasibility and appropriateness (Lewis et al. 2015, Proctor et al. 2011). Successful implementation is achieved by using strategies to target a variety of implementation outcomes at multiple levels (Powell et al. 2015).

The principles of implementation science may be readily applied to the improvement of research methods in psychology. First, the multilevel determinants perspective recognizes that there are multiple systems and stakeholders at many levels that influence the implementation of effective research practices. Implementation will be more effective if applied researchers understand and believe in the new method, believe they can apply and interpret the method themselves, that the method meets their needs and resources, and that their peers (other researchers, editors, reviewers, funding agencies) expect that method to be applied. Implementation will be more widespread in contexts (e.g. research groups and programs, journals, professional organizations) that expect, support and reward the implementation of a method, and have policies and incentives that promote better practices.

It will be critical to consider the degree to which the characteristics of a new method, such as its adaptability, complexity and design, relevance for given problems, and cost,

influence implementation. The widespread adoption of macros, spreadsheets and online tools for common statistical procedures (Dawson 2014, Preacher et al. 2006) highlights the importance of reducing complexity and cost as a way of increasing implementation. Kirk (1996) noted that some statistical information was more common (such as R^2 as a measure of effect size in regression) than others because it was commonly included in statistical output. Thus, implementation will be improved if software packages can present output in a way that conforms to best reporting practices for psychologists (and discourage poor practices).

It will be important to measure implementation specific outcomes to understand successes and failures of efforts to implement new methods (Lewis et al. 2015, Proctor et al. 2011). Research on the effectiveness of methods (i.e. whether or not the method solves a specific problem well) has long been the central focus of methodological research. However, effective interventions may not be implemented if they are not viewed as acceptable to stakeholders, seen as appropriate solutions to a given problem, or perceived as actually feasible to use (Crale et al. 2018). Implementation success should also consider the degree to which methods are adopted, penetrate research in a given field, and are used sustainably with fidelity. Theoretically, implementation outcomes are the precursor to the effectiveness of interventions in practice (Proctor et al. 2011). Without achieving success in implementation outcomes, a quantitative intervention may be effective, in that it produces appropriate and accurate answers, but it may not be viewed as acceptable, feasible or sustainable, leading to low adoption, or it may be widespread but conducted with poor fidelity, so it is not very effective. In other words, successful implementation is defined by widespread, appropriate and accurate use of a method.

Finally, there are a wide variety of strategies that might be used to improve the implementation of methods that can be borrowed from successful implementation efforts in healthcare (Powell et al. 2015). Formal implementation blueprints can be used to guide implementation efforts, identify change targets, timelines and measures of progress or success. Organizational efforts can work to incentivize best practices, such as the badge system that is used to recognize research that conforms to open science practices (Kidwell et al. 2016), or develop best-practice guidelines that can be used to guide the work of applied researchers (Proctor et al. 2009). Some work has suggested, for example, that providing guidelines for reporting results RCTs was related to improvements in statistical reporting (Turner et al. 2012). An emerging area of interest is in how the design and usability of interventions and implementation strategies may impede or facilitate widespread implementation (Lyon & Koerner 2016). Finally, it will be critical to design and test effective methods of training, supervision and consultation that are based in current best-practices (Dorsey et al. 2013, Lyon et al. 2011).

Looking Toward a More Transparent and Reproducible Clinical Psychological Science: Conclusion and Future Directions

In the short period of time—several years—since we initiated a paper questioning the absence of clinical psychology from the replicability conversation (Tackett et al. 2017), substantial strides have already been made toward engagement in reform efforts by the clinical psychology community. Clinical psychology has now seen the emergence of badges promoting open science and transparency in two field-specific journals, a special issue on the topic is forthcoming in the *Journal of Abnormal Psychology*, discussion panels and symposia have emerged at field-specific and field-wide (e.g., the *Association for Psychological Science* convention) conferences, and initial efforts at establishing a database of interested clinical psychologists

was met with resounding success. It's clear that there is a sector of the clinical psychology community ready and willing to tackle these problems head-on, but it is simultaneously clear that clinical psychology, as a subfield, is lagging behind other areas of psychological science in these efforts.

This current status places us in a critical position at the current time, with a need to have more engagement from the clinical psychology community across many levels – from examining and charting the potential weak points in our literature base, to defining and operationalization replication on both theoretical and statistical levels, to applying (and modifying, as necessary) existing tools to promote a more open, transparent, and reproducible science. There are also great opportunities to consider the many ways we, as a field, have already been tackling these issues, even if not specifying them under an explicitly open science framework, and leveraging our existing knowledge and expertise to accelerate and facilitate the open science and reform movement writ large.

As clinical psychology continues to engage with this conversation, and we hope increasingly so, we will undoubtedly encounter new challenges along the way. There are concerns discussed somewhat informally, that we have encountered, including the sense that those advancing the open science and reform movement may not be open-minded, judicious, or considerate. This is a larger challenge than we can address here—certainly, a movement that is largely predicated on criticizing an established structure will (probably necessarily) come across as antagonistic at times. Nonetheless, that does not mean that high levels of antagonism are necessary or even particularly helpful in advancing these efforts. Perhaps this is a domain where the increased involvement of individuals with clinical training might be of some utility to the broader movement, overall.

There are other tensions currently, for example in approaches to the peer review and editorial process. If editors or reviewers are overly critical of theories or results that contradict their own, while being overly accommodating to those that they favor, this might point to an existing bias in the system. A transparent author may inadvertently provide extra ammunition for reviewers to go beyond the usual in identifying potential flaws in the methods or results (regardless of their verisimilitude). One potential solution to this issue is RRs, as after the initial “results-free” peer review of the study introduction and methods, reviewers and editors may not rescind their decision to publish the paper on the basis of the results obtained, effectively eliminating that source of bias. Other forms of preregistration may also protect against reviewer bias. When methods and analytical strategies are prespecified and timestamped in a registration document, potential adhoc reviewer criticisms are considered exploratory and bear the burden of (registered) empirical investigation to present a strong defense against final results. Beyond registration, other solutions may be employed, such as signed reviews, open (but anonymous) reviews, the use of preprints for feedback from a wider scientific audience, and monitoring for this sort of reviewer behavior by editors. None of these approaches are panaceas; each has benefits and drawbacks, and some of these are discussed in detail elsewhere (Ross-Hellauer et al. 2017). These, and other, issues are actively being discussed in the open science community. There is a place for clinical psychologists to weigh in, and be part of working through some of these potentially thornier issues.

In sum, we hope that we have provided some acceleration to the ongoing conversations about open science, transparency, and replicability within the field of clinical psychological science. It is our assessment that clinical psychology has many strengths to leverage, including aspects of our previous methods and practices that have potentially

served us well, as well as expertise and knowledge that will benefit the open science and reform movement more broadly. However, there remains much work to be done, and with yet a small minority of clinical psychologists actively engaging these efforts, it is clear that this movement has not yet taken hold in our subfield. We call out to our clinical colleagues to join us in these challenging efforts, and facilitate a movement that will result in better science and a more trustworthy evidence base.

Summary Points List

1. Although clinical psychological scientists have been slower to engage in the replicability conversation, recent years have shown increased awareness and involvement from clinical psychologists.
2. Identifying weak points in the clinical psychology evidence base is a critical need, and likely includes areas reliant on small sample sizes (e.g., clinical neuroscience and intervention research), gaps between diagnostic assessment in the field and in research settings, and generalizability of randomized control trials to field and community settings.
3. A number of problematic methods and practices likely impact the replicability of clinical psychological findings, including low sample size and power, diagnostic unreliability, publication bias, the use of questionable research practices, and variable adherence to registration guidelines (e.g., for intervention trials).
4. Several areas for clinical psychology to consider in a replication context includes the use of meta-analysis (which has its historical roots in the field of clinical psychology), the best way to define and operationalize replication in clinical psychology, and the extent to which current statistical concerns (e.g., the problematic use of null hypothesis significance testing) generalize to clinical psychological research.
5. Many current proposals for improvement and reform can be applied and adapted for clinical psychological research, including open materials, open data, (pre)registration, Registered Reports, and multisite collaborations for both replications and original studies.
6. The clinical psychological expertise and knowledge base can also be leveraged to facilitate advances in the open science movement; one such example is considering existing work from intervention/implementation science as a mechanism for enhancing adoption and maintenance of open science methods.

Future Issues List

1. Explicit efforts are needed at identifying weak spots in the clinical science research base. This mapping may occur by examining research topics with evidence for limited replicability or generalizability or identifying those fields most associated with known problems that limit replicability (e.g., small sample sizes, reduced power, problematic measurement, publication bias, etc.).
2. A better understanding is needed of known lab-to-field gaps, for example in diagnostic discrepancies between research and field settings and problematic generalizability of RCTs to the field and to community settings.
3. As clinical psychologists increasingly engage with current open science and reform tools, these solutions will most likely be in need of adaptation to better suit the types of research designs, samples, and methodological/analytic approaches that clinical psychologists typically use. This tailoring will most expediently follow from early attempts at adapting existing tools and solutions in current research efforts.
4. Clinical psychologists need to directly engage with persistent questions around how to define a replication attempt and a replication outcome (e.g., success/failure), tackling this question at both theoretical and methodological levels.
5. Clinical psychologists should leverage our existing knowledge base to enhance open science and reform efforts, for example, by applying principles from implementation science to these domains.

6. The field of clinical psychology is in need of greater involvement and engagement across the field. These reforms will take effort, and will benefit most by receiving input from researchers across domains of clinical psychological science.

References

- Abram SV, DeYoung CG. 2017. Using personality neuroscience to study personality disorder. *Personal. Disord. Theory Res. Treat.* 8: 2-13
- Aiken LS, West SG, Millsap, RE. 2008. Doctoral training in statistics, measurement, and methodology in psychology: Replication and extension of Aiken, West, Sechrest, and Reno's (1990) survey of PhD programs in North America. *Am. Psych.* 63(1): 32–50. <https://doi.org/10.1037/0003-066X.63.1.32>
- Aiken LS, West SG, Reno RR, Kazdin AE, Sherman SJ. 1990. Graduate training in statistics, methodology, and measurement in psychology a survey of PhD programs in North America. *Am. Psych.* 45(6): 721–34.
- Altman DG, & Goodman SN. 1994. Transfer of technology from statistical journals to the biomedical literature. *JAMA.* 272(2): 129. <https://doi.org/10.1001/jama.1994.03520020055015>
- American Psychological Association. 2010. *Publication Manual of the American Psychological Association, Sixth Edition. Publication Manual of the American Psychological Association, Sixth Edition.* Washington, DC: American Psychological Association. <https://doi.org/10.1006/mgme.2001.3260>
- American Psychological Association. 2017. Ethical principles of psychologists and code of conduct.
- APA Publications and Communications Board Working Group on Journal Article Reporting Standards. 2008. Reporting standards for research in psychology: Why do we need them? What might they be? *Am. Psych.* 63(9): 839–51. <https://doi.org/10.1037/0003-066X.63.9.839>

- Bakker M, Hartgerink CHJ, Wicherts JM, van der Maas HLJ. 2016. Researchers' intuitions about power in psychological research. *Psych. Sci.* 27: 1069-77.
- Bakker M, van Dijk A, Wicherts JM. 2012. The rules of the game called psychological science. *Persp. Psych. Sci.* 7(6): 543-54. <https://doi.org/10.1177/1745691612459060>
- Bakker M, Wicherts JM. 2011. The (mis)reporting of statistical results in psychology journals. *Beh. Res. Meth.* 43(3): 666-78. <https://doi.org/10.3758/s13428-011-0089-5>
- Bauer, D. J. (2007). Observations on the use of growth mixture models in psychological research. *Mult. Beh. Res.*, 42(4): 757-786.
- Becker EM, Smith AM, Jensen-Doss A. 2013. Who's using treatment manuals? A national survey of practicing therapists. *Beh. Res. Ther.* 51(10): 706-10. <https://doi.org/10.1016/j.brat.2013.07.008>
- Beidas R, Kendall P. 2010. Training therapists in evidence-based practice: a critical review of studies from a systems contextual perspective. *Clin. Psychol. Sci. Pract.* 17: 1-30. <https://doi.org/10.1111/j.1468-2850.2009.01187.x>
- Belia S, Fidler F, Williams J, Cumming G. 2005. Researchers misunderstand confidence intervals and standard error bars. *Psychol. Meth.* 10(4): 389-96. <https://doi.org/10.1037/1082-989X.10.4.389>
- Berle D, Starcevic V. (2007). Inconsistencies between reported test statistics and p-values in two psychiatry journals. *Int. Journ. Meth. Psych. Res.* 16(4): 202-7. <https://doi.org/10.1002/mpr.225>
- Blohowiak BB, Cohoon J, de-Wit L, Eich E, Farach FJ, et al. 2018. Badges to acknowledge open practices. Retrieved from: <https://osf.io/tvyxz/>
- Bransford JD, Brown AL, Cocking RR. (2000). *How people learn: Brain, mind, experience, and school: Expanded Edition*. [https://doi.org/10.1016/0885-2014\(91\)90049-J](https://doi.org/10.1016/0885-2014(91)90049-J)

- Brown NJL, Heathers JAJ. 2017. The GRIM test: A simple technique detects numerous anomalies in the reporting of results in psychology. *Soc. Psychol. Pers. Sci.* 8(4): 363–9. <https://doi.org/10.1177/1948550616673876>
- Button KS, Ioannidis JPA, Mokrysz C, Nosek BA, Flint J, Robinson ESJ, Munafò MR. 2013. Power failure: why small sample size undermines the reliability of neuroscience. *Nat. Rev. Neuro.* 14(5): 365–76. <https://doi.org/10.1038/nrn3475>
- Calamia M, Markon K, Tranel D. 2012. Scoring higher the second time around: Meta-analyses of practice effects in neuropsychological assessment. *The Clin. Neuropsychol.* 26(4): 543–70. <https://doi.org/10.1080/13854046.2012.680913>
- Campbell DT, Fiske DW. 1959. Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychol. Bull.* 56(2): 81–105. <https://doi.org/10.1037/h0046016>
- Chambers DA, Glasgow RE, Stange KC. 2013. The dynamic sustainability framework: Addressing the paradox of sustainment amid ongoing change. *Implement. Sci.* 8(1): 1–11. <https://doi.org/10.1186/1748-5908-8-117>
- Chance BL. 2002. Components of statistical thinking and implications for instruction and assessment. *J. Stat. Educ.* 10(3). <https://doi.org/10.1080/10691898.2002.11910677>
- Chartier CR, McCarthy RJ. 2018. StudySwap: A platform for inter-lab replication, collaboration, and research resource exchange. Retrieved from: <https://osf.io/rd37b/>
- Chorpita BF, Becker KD, Daleiden EL. 2007. Understanding the common elements of evidence-based practice: Misconceptions and clinical examples. *J. Am. Acad. Child Adol. Psychiat.* 46(5): 647–52. <https://doi.org/10.1097/chi.0b013e318033ff71>
- Cronbach LJ, Meehl PE. 1955. Construct validity in psychological tests. *Psychol. Bull.* 52: 281–302.

- Cohen J. 1962. The statistical power of abnormal-social psychological research: A review. *J. Abnorm. Soc. Psychol.* 65(3): 145.
- Cohen J. 1990. Things I have learned (so far). *Am. Psychol.* 45(12): 1304–12.
- Contopoulos-Ioannidis DG, Alexiou G, Gouvas TC, Ioannidis JP. 2008. Life cycle of translational research for medical interventions. *Science* 321: 1298–9.
<https://doi.org/10.1126/science.1160622>
- Conway CC, Tackett JL, Skodol AE. 2017. Are personality disorders assessed in young people? *Am. J. Psychiat.* 174(10): 1000–1.
<https://doi.org/10.1176/appi.ajp.2017.17040454>
- Crable EL, Biancarelli D, Walkey AJ, Allen CG, Proctor EK, Drainoni ML. 2018. Standardizing an approach to the evaluation of implementation science proposals. *Implement. Sci.* 13(1): 71. <https://doi.org/10.1186/s13012-018-0770-5>
- Cremers HR, Wager TD, Yarkoni T. 2017. The relation between statistical power and inference in fMRI. *PLoS ONE* 12: e0184923.
- Cuijpers P. 2016. Are all psychotherapies equally effective in the treatment of adult depression? The lack of statistical power of comparative outcome studies. *Evid. Based Ment. Health* 19(2): 39–42. <https://doi.org/10.1136/eb-2016-102341>
- Cuijpers P, Sijbrandij M, Koole S, Huibers M, Berking M, Andersson G. 2014. Psychological treatment of generalized anxiety disorder: A meta-analysis. *Clin. Psychol. Rev.* 34(2): 130–40. <https://doi.org/10.1016/j.cpr.2014.01.002>
- Cuijpers P, Smit F, Bohlmeijer E, Hollon SD, Andersson G. 2010. Efficacy of cognitive–behavioural therapy and other psychological treatments for adult depression: meta-analytic study of publication bias. *Brit. J. Psychiat.* 196(03): 173–8.
<https://doi.org/10.1192/bjp.bp.109.066001>

- Culverhouse RC, Saccone NL, Horton AC, Ma Y, Anstey KJ, Banaschewski T, et al. 2018. Collaborative meta-analysis finds no evidence of a strong interaction between stress and 5-HTTLPR genotype contributing to the development of depression. *Mol. Psychiat.* 23(1): 133–42. <https://doi.org/10.1038/mp.2017.44>
- Cumming G, Williams J, Fidler F. 2004. Replication and researchers' understanding of confidence intervals and standard error bars. *Und. Stat.* 3(4): 299–311. https://doi.org/10.1207/s15328031us0304_5
- Damschroder LJ, Aron DC, Keith RE, Kirsh SR, Alexander JA, Lowery JC. 2009. Fostering implementation of health services research findings into practice : a consolidated framework for advancing implementation science. *Implement. Sci.* 4(1): 50. <https://doi.org/10.1186/1748-5908-4-50>
- Dawson JF. (2014). Moderation in management research: What, why, when, and how. *J. Bus. Psychol.* 29(1): 1–19. <https://doi.org/10.1007/s10869-013-9308-7>
- De Angelis C, Drazen JM, Frizelle FA, Haug C, Hoey J, Horton R, et al. 2004. Clinical trial registration: A statement from the International Committee of Medical Journal Editors. *N. Eng. J. Med.* 351(12): 1250–1. <https://doi.org/10.1056/NEJMe048183>
- Dorsey S, Pullmann MD, Deblinger E, Berliner L, Kerns SE, Thompson K., et al. 2013. Improving practice in community-based settings: a randomized trial of supervision—study protocol. *Implement. Sci.* 8(89): 1–12. <https://doi.org/10.1186/1748-5908-8-89>
- Dorsey S, Pullmann MD, Kerns SEU, Jungbluth N, Meza R, Thompson K, & Berliner L. 2017. The juggling act of supervision in community mental health: Implications for supporting evidence-based treatment. *Adm. Policy Ment. Health* 44(6): 838–52. <https://doi.org/10.1007/s10488-017-0796-z>

- Duncan LE & Keller MC. 2011. A critical review of the first 10 years of candidate gene-by-environment interaction research in psychiatry. *Am. J. Psychiat.* 168(10): 1041–9.
<https://doi.org/10.1176/appi.ajp.2011.11020191>
- Dwan K, Gamble C, Williamson PR, Kirkham JJ, & the Reporting Bias Group. 2013. Systematic review of the empirical evidence of study publication bias and outcome reporting bias — An updated review. *PLoS ONE* 8(7): e66844.
<https://doi.org/10.1371/journal.pone.0066844>
- Fabrigar LR, Wegener DT, MacCallum RC, & Strahan EJ. 1999. Evaluating the use of exploratory factor analysis in psychological research. *Psychol. Methods* 79(3): 272–99.
<https://doi.org/doi:10.1037/1082-989X.4.3.272>
- Fanelli D. 2010. Do pressures to publish increase scientists' bias? An empirical support from US states data. *PLoS ONE* 5(4): e10271.
<https://doi.org/10.1371/journal.pone.0010271>
- Finch S, Cumming G, Williams J, Palmer L, Griffith E, Alders C, et al. 2004. Reform of statistical inference in psychology: the case of memory & cognition. *Behav. Res. Methods Instrum. Compute.* 36(2): 312–24. <https://doi.org/10.3758/BF03195577>
- Finkel EJ, Eastwick PW, & Reis HT. 2015. Best research practices in psychology: Illustrating epistemological and pragmatic considerations with the case of relationship science. *J. Pers. Soc. Psychol.* 108(2): 275–97. <https://doi.org/10.1037/pspi0000007>
- Fiske DW & Campbell DT. 1992. Citations do not solve problems. *Psychol. Bull.* 112(3): 393–5. <https://doi.org/10.1037/0033-2909.112.3.393>
- Fraley RC & Vazire S. 2014. The N-pact factor: Evaluating the quality of empirical journals with respect to sample size and statistical power. *PLoS ONE* 9(10), e109019.
<https://doi.org/10.1371/journal.pone.0109019>

- Frances AJ & Widiger T. 2012. Psychiatric diagnosis: Lessons from the DSM-IV past and cautions for the DSM-5 future. *Annu. Rev. Clin. Psychol.* 8(1): 109–30.
<https://doi.org/10.1146/annurev-clinpsy-032511-143102>
- Frank MC, Bergelson E, Bergmann C, Cristia A, Floccia C, Gervain J, et al. 2017. A collaborative approach to infant research: Promoting reproducibility, best practices, and theory-building. *Infancy* 22(4): 421–35. <https://doi.org/10.1111/infa.12182>
- Freedman R, Lewis DA, Michels R, Pine DS, Schultz SK, Tamminga CA, et al. 2013. Editorial: The initial field trials of DSM-5: New blooms and old thorns. *Am. J. Psychiat.* 170: 1-5.
- Garb HN. 2005. Clinical judgment and decision making. *Annu. Rev. Clin. Psychol.* 1(1): 67–89.
<https://doi.org/10.1146/annurev.clinpsy.1.102803.143810>
- García-Berthou E & Alcaraz C. 2004. Incongruence between test statistics and P values in medical papers. *BMC Med. Res. Methodol.* 4(13). <https://doi.org/10.1186/1471-2288-4-13>
- Garland AF, Brookman-Frazee L, Hurlburt MS, Accurso EC, Zoffness RJ, Haine-Schlagel, R, & Ganger W. 2010. Mental health care for children with disruptive behavior problems: A view inside therapists' offices. *Psychiatr. Serv.* 61(8): 788–95.
<https://doi.org/10.1176/appi.ps.61.8.788>
- Gigerenzer G. 2004. Mindless statistics. *J. Socio. Econ.* 33(5): 587–606.
<https://doi.org/10.1016/j.socec.2004.09.033>
- Gigerenzer G & Gaissmaier W. 2011. Heuristic decision making. *Annu. Rev. Psychol.* 62: 451–82. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Golinski C & Cribbie RA. 2009. The expanding role of quantitative methodologists in advancing psychology. *Can. Psychol.* 50(2): 83–90. <https://doi.org/10.1037/a0015180>

- Greenland S, Senn SJ, Rothman KJ, Carlin JB, Poole C, Goodman SN, & Altman DG. 2016. Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations. *Eur. J. Epidemiol.* 31(4): 337–50. <https://doi.org/10.1007/s10654-016-0149-3>
- Grimm KJ. 2007. Multivariate longitudinal methods for studying developmental relationships between depression and academic achievement. *Int. J. Behav. Dev.* 31(4): 328–39. <https://doi.org/10.1177/0165025407077754>
- Herschell AD, Kolko DJ, Baumann BL, & Davis AC. 2010. The role of therapist training in the implementation of psychosocial treatments: A review and critique with recommendations. *Clin. Psychol. Rev.* 30(4): 448–66. <https://doi.org/10.1016/j.cpr.2010.02.005>
- Hopewell S, Loudon K, Clarke MJ, Oxman AD, & Dickersin K. 2009. Publication bias in clinical trials due to significance of trial results. *Cochrane Database Syst. Rev.* 2009(1).
- Huys QJM, Maia TV, & Frank MJ. 2016. Computational psychiatry as a bridge from neuroscience to clinical applications. *Nat. Neurosci.* 19(3): 404–13. <https://doi.org/10.1038/nn.4238>
- Insel T, Cuthbert B, Garvey M, Heinssen R, Pine DS, Quinn K, et al. 2010. Research domain criteria (RDoC): toward a new classification framework for research on mental disorders. *Am. J. Psychiatr.* 167(7): 748–51.
- Ioannidis JPA. 2005. Why most published research findings are false. *PLoS Med.* 2(8): e124. <https://doi.org/10.1371/journal.pmed.0020124>
- Ioannidis JPA. 2018. The proposal to lower P value thresholds to .005. *JAMA.* <https://doi.org/10.1001/jama.2018.1536>

- Jackson DL, Gillaspay JA, & Purc-Stephenson R. 2009. Reporting practices in confirmatory factor analysis: an overview and some recommendations. *Psychol. Methods* 14(1): 6–23.
<https://doi.org/10.1037/a0014694>
- Jenkins MM, Youngstrom EA, Washburn JJ, & Youngstrom JK. 2011. Evidence-based strategies improve assessment of pediatric bipolar disorder by community practitioners. *Prof. Psychol. Res. Pr.* 42(2): 121–9. <https://doi.org/10.1037/a0022506>
- Joel S, Eastwick PW, & Finkel EJ. 2018. Open sharing of data on close relationships and other sensitive social psychological topic: Challenges, tools, and future directions. *Adv. Methods Pr. Psychol. Sci.* 1: 86-94.
- John LK, Loewenstein G, & Prelec D. 2012. Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychol. Sci.* 23(5): 524–32.
<https://doi.org/10.1177/0956797611430953>
- Jones KD. 2012. A critique of the DSM-5 field trials. *J. Nerv. Ment. Dis.* 200(6): 517–9.
<https://doi.org/10.1097/NMD.0b013e318257c699>
- Joyce B & Showers B. 2002. Student achievement through staff development. *Designing Training and Peer Coaching: Our Needs for Learning.*
- Kahneman D & Tversky A. 1974. Judgment under uncertainty : Heuristics and biases. *Science* 185(4157): 1124–31.
- Kidwell MC, Lazarević LB, Baranski E, Hardwicke TE, Piechowski S, Falkenberg LS, et al. 2016. Badges to acknowledge open practices: A simple, low-cost, effective method for increasing transparency. *PLoS Biol.* 14(5): e1002456.
<https://doi.org/10.1371/journal.pbio.1002456>

- Kim H & Eaton NR. 2015. The hierarchical structure of common mental disorders: Connecting multiple levels of comorbidity, bifactor models, and predictive validity. *J. Abnorm. Psychol.* 124(4): 1064–78. <https://doi.org/10.1037/abn0000113>
- King KM, Littlefield AK, McCabe CJ, Mills KL, Flournoy J, & Chassin L. 2017. Longitudinal modeling in developmental neuroimaging research: Common challenges, and solutions from developmental psychology. Manuscript in press, *Dev. Cog. Neurosci.* <https://doi.org/10.1016/j.dcn.2017.11.009>
- Kirk RE. 1996. Practical significance : A concept whose time has come. *Educ. Psychol. Meas.* 56(5): 746–59.
- Klein RA., Ratliff KA, Vianello M, Adams RB, Bahník Š, Bernstein MJ, et al. 2014. Investigating variation in replicability: A “Many Labs” replication project. *Soc. Psychol.* 45(3): 142–52. <https://doi.org/10.1027/1864-9335/a000178>
- Lance CE & Vandenberg RJ. 2008. *Statistical and methodological myths and urban legends: Doctrine, verity and fable in the organizational and social sciences*. New York: Taylor & Francis <https://doi.org/10.4324/9780203867266>
- Lash TL, Collin LJ, & Van Dyke ME. 2018. The replication crisis in epidemiology: Snowball, snow job, or winter solstice? Manuscript in press, *Curr. Epidemiol. Rep.*
- Leichsenring F, Abbass A, Hilsenroth MJ, Leweke F, Luyten P, Keefe JR, et al. 2017. Biases in research: risk factors for non-replicability in psychotherapy and pharmacotherapy research. *Psychol. Med.* 47(06): 1000–11. <https://doi.org/10.1017/S003329171600324X>
- Lewis CC, Fischer S, Weiner BJ, Stanick C, Kim M, & Martinez RG. 2015. Outcomes for implementation science: An enhanced systematic review of instruments using

- evidence-based rating criteria. *Implement. Sci.* 10(1): 1–17.
<https://doi.org/10.1186/s13012-015-0342-x>
- Lilienfeld SO. 2017. Psychology’s replication crisis and the grant culture: Righting the ship. *Perspect. Psychol. Sci.* 12: 660–4.
- Lord FM. 1967. A paradox in the interpretation of group comparisons. *Psychol. Bull.* 68(3): 304–5. <https://doi.org/10.1037/h0025105>
- Lovett MC & Greenhouse JB. 2000. Applying cognitive theory to statistics instruction. *Am. Stat.* 54(3): 196–206. <https://doi.org/10.1080/00031305.2000.10474545>
- Lyon AR & Koerner K. 2016. User-centered design for psychosocial intervention development and implementation. *Clin. Psychol. Sci. Pract.* 23(2): 180–200.
<https://doi.org/10.1111/cpsp.12154>
- Lyon AR, Stirman SW, Kerns SEU, & Bruns EJ. 2011. Developing the mental health workforce: Review and application of training approaches from multiple disciplines. *Adm. Pol. Ment. Health* 38(4): 238–53. <https://doi.org/10.1007/s10488-010-0331-y>
- MacCallum RC, Zhang S, Preacher KJ, & Rucker DD. 2002. On the practice of dichotomization of quantitative variables. *Psychol. Methods* 7: 19–40.
- Marsh HW, Hau KT, & Wen Z. 2004. In search of golden rules: Comment on hypothesis-testing approaches to setting cutoff values for fit indexes and dangers in overgeneralizing Hu and Bentler’s (1999) findings. *Struct. Equ. Modeling* 11(3): 320–41. https://doi.org/10.1207/s15328007sem1103_2
- McNeish D, An J, & Hancock GR. 2018. The thorny relation between measurement quality and fit index cutoffs in latent variable models. *J. Pers. Assess.* 100(1): 43–52.
<https://doi.org/10.1080/00223891.2017.1281286>

- McShane BB, Gal D, Gelman A, Robert C, & Tackett JL. 2017. Abandon statistical significance. Manuscript in press, *Am. Stat.*
- McShane BB, Tackett JL, Bockenholt U, & Gelman A. 2017. Large scale replication projects in contemporary psychological research. Manuscript in press, *Am. Stat.*
- Miller GA & Chapman JP. (2001). Misunderstanding analysis of covariance ANCOVA. *J. Abnorm. Psychol.* 110(1): 40–8.
- Mills L, Abdulla E, & Cribbie RA. 2010. Quantitative methodology research: Is it on psychologists' reading lists? *Tutor. Quant. Methods Psychol.* 6(2): 52–60.
<https://doi.org/10.20982/tqmp.06.2.p052>
- Miloyan B & Fried EI. 2017. A reassessment of the relationship between depression and all-cause mortality in 3,604,005 participants from 293 studies. *World Psychiat.* 16: 219-20.
- Moberg CA & Humphreys K. 2017. Exclusion criteria in treatment research on alcohol, tobacco and illicit drug use disorders: A review and critical analysis. *Drug Alcohol Rev.* 36(3): 378–88. <https://doi.org/10.1111/dar.12438>
- Moffitt TE, Caspi A, Taylor A, Kokaua J, Milne BJ, Polanczyk G, & Poulton R. 2010. How common are common mental disorders? Evidence that lifetime prevalence rates are doubled by prospective versus retrospective ascertainment. *Psychol. Med.* 40(6): 899–909. <https://doi.org/10.1017/S0033291709991036>
- Morris ZS, Wooding S, & Grant J. 2011. The answer is 17 years, what is the question: Understanding time lags in translational research. *J. R. Soc. Med.* 104(12): 510–20.
<https://doi.org/10.1258/jrsm.2011.110180>
- Moshontz H, Campbell L, Ebersole CR, IJzerman H, Urry HL, Forscher PS, et al. 2018. The Psychological Science Accelerator: Advancing psychology through a distributed

- collaborative network. Manuscript in press, *Adv. Meth. Prac. Psychol. Sci.* Retrieved from: <https://doi.org/10.31234/osf.io/785qu>
- Munafò MR, Nosek BA, Bishop DVM, Button KS, Chambers CD, Percie du Sert N, et al. (2017). A manifesto for reproducible science. *Nat. Hum. Behav.* 1(1): 21. <https://doi.org/10.1038/s41562-016-0021>
- Munder T, Brüttsch O, Leonhart R, Gerger H, & Barth J. 2013. Researcher allegiance in psychotherapy outcome research: An overview of reviews. *Clin. Psychol. Rev.* 33(4): 501–11. <https://doi.org/10.1016/j.cpr.2013.02.002>
- Nickerson RS. 2000. Null hypothesis significance testing: A review of an old and continuing controversy. *Psychol. Methods* 5(2): 241–301. <https://doi.org/10.1037/1082-989X.5.2.241>
- Nosek BA, Ebersole CR, DeHaven AC, & Mellor DT. 2018. The preregistration revolution. *Proc. Natl. Acad. Sci.* 115(11): 2600–6. <https://doi.org/10.1073/pnas.1708274114>
- Nuijten MB, Hartgerink CHJ, van Assen MALM, Epskamp S, & Wicherts JM. 2016. The prevalence of statistical reporting errors in psychology (1985–2013). *Behav. Res. Methods* 48(4): 1205–26. <https://doi.org/10.3758/s13428-015-0664-2>
- Polanczyk GV, Salum GA, Sugaya LS, Caye A, & Rohde LA. 2015. Annual research review: A meta-analysis of the worldwide prevalence of mental disorders in children and adolescents. *J. Child Psychol. Psychiat.* 56(3): 345–65. <https://doi.org/10.1111/jcpp.12381>
- Powell BJ, Waltz TJ, Chinman MJ, Damschroder LJ, Smith JL, Matthieu MM, et al. 2015. A refined compilation of implementation strategies: Results from the Expert Recommendations for Implementing Change (ERIC) project. *Implement. Sci.* 10(1): 1–14. <https://doi.org/10.1186/s13012-015-0209-1>

- Preacher KJ, Curran PJ, & Bauer DJ. 2006. Computational tools for probing interactions in multiple linear regression, multilevel modeling, and latent curve analysis. *J. Educ. Behav. Stat.* 31(4): 437–48. <https://doi.org/10.3102/10769986031004437>
- Proctor EK, Landsverk J, Aarons G, Chambers D, Glisson C, & Mittman B. 2009. Implementation research in mental health services: An emerging science with conceptual, methodological, and training challenges. *Adm. Pol. Ment. Health* 36(1): 24–34. <https://doi.org/10.1007/s10488-008-0197-4>
- Proctor EK, Silmere H, Raghavan R, Hovmand P, Aarons G, Bunger A, et al. 2011. Outcomes for implementation research: Conceptual distinctions, measurement challenges, and research agenda. *Adm. Pol. Ment. Health* 38(2): 65–76. <https://doi.org/10.1007/s10488-010-0319-7>
- Rapport MD, Orban SA, Kofler MJ, & Friedman LM. 2013. Do programs designed to train working memory, other executive functions, and attention benefit children with ADHD? A meta-analytic review of cognitive, academic, and behavioral outcomes. *Clin. Psychol. Rev.* 33(8): 1237–52. <https://doi.org/10.1016/j.cpr.2013.08.005>
- Reardon KW, Mercadante EJ, & Tackett JL. 2017. The assessment of personality disorder: methodological, developmental, and contextual considerations. *Curr. Opin. Psychol.* 21: 39–43.
- Reardon KW, Smack AJ, Herzhoff K, & Tackett JL. 2018. An *N*-pact factor for clinical psychological research. Manuscript under review.
- Regier DA, Narrow WE, Clarke DE, Kraemer HC, Kuramoto SJ, Kuhl EA, & Kupfer DJ. 2013. DSM-5 field trials in the United States and Canada, part II: Test-retest reliability of selected categorical diagnoses. *Am. J. Psychiat.* 170(1): 59–70. <https://doi.org/10.1176/appi.ajp.2012.12070999>

- Ross-Hellauer T, Deppe A, & Schmidt B. 2017. Survey on open peer review: Attitudes and experience amongst editors, authors, and reviewers. *PLoS ONE* 12(12): e0189311. <https://doi.org/10.1371/journal.pone.0189311>
- Rothwell PM. 2005. External validity of randomised controlled trials: “To whom do the results of this trial apply?” *Lancet* 365(9453): 82–93. [https://doi.org/10.1016/S0140-6736\(04\)17670-8](https://doi.org/10.1016/S0140-6736(04)17670-8)
- Schmitt N. 1996. Uses and abuses of coefficient alpha. *Psychol. Assess.* 8(4): 350–3. <https://doi.org/10.1037/1040-3590.8.4.350>
- Sedlmeier P & Gigerenzer G. 1989. Do studies of statistical power have an effect on the power of studies? *Psychol. Bull.* 105(2): 309–16. <https://doi.org/10.1037/0033-2909.105.2.309>
- Shadish W, Cook T, & Campbell D. 2002. Experimental and quasi-experimental designs. In *Experimental and Quasi-Experimental Designs for Generalized Causal Inference.* (171–206). <https://doi.org/10.1093/obo/9780195389678-0053>
- Sharpe D. 2013. Why the resistance to statistical innovations? Bridging the communication gap. *Psychol. Methods* 18(4): 572–82. <https://doi.org/10.1037/a0034177>
- Sher KJ, Jackson KM, & Steinley D. 2011. Alcohol use trajectories and the ubiquitous cat's cradle: Cause for concern?. *J. Abnorm. Psychol.* 120: 322–35.
- Shrout PE & Bolger N. 2002. Mediation in experimental and nonexperimental studies: New procedures and recommendations. *Psychol. Methods* 7(4): 422–45. <https://doi.org/10.1037/1082-989X.7.4.422>
- Sijtsma K. 2009. On the use, the misuse, and the very limited usefulness of cronbach's alpha. *Psychometrika* 74(1): 107–20. <https://doi.org/10.1007/s11336-008-9101-0>

- Simmons JP, Nelson LD, & Simonsohn U. 2011. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychol. Sci.* 22(11): 1359–66. <https://doi.org/10.1177/0956797611417632>
- Soto CJ. 2018. How replicable are links between personality traits and consequential life outcomes? The Life Outcomes of Personality Replication Project. Retrieved from <https://osf.io/9dta7/>
- Sternberg RJ. 1992. Psychological Bulletin's top 10 "hit parade." *Psychol. Bull.* 112(3): 387–8. <https://doi.org/10.1037/0033-2909.112.3.387>
- Stirman SW, DeRubeis RJ, Crits-Christoph P, & Brody PE. 2003. Are samples in randomized controlled trials of psychotherapy representative of community Outpatients? A new methodology and initial findings. *J. Consult. Clin. Psychol.* 71(6): 963–72. <https://doi.org/10.1037/0022-006X.71.6.963>
- Stirman SW, DeRubeis RJ, Crits-Christoph P, & Rothman A. 2005. Can the randomized controlled trial literature generalize to nonrandomized patients? *J. Consult. Clin. Psychol.* 73(1): 127–35. <https://doi.org/10.1037/0022-006X.73.1.127>
- Stuart EA, Bradshaw CP, & Leaf PJ. 2015. Assessing the generalizability of randomized trial results to target populations. *Prev. Sci.* 16(3): 475–85. <https://doi.org/10.1007/s11121-014-0513-z>
- Tackett JL, Brandes CM, & Reardon KW. 2018. Leveraging the Open Science Framework in clinical psychological assessment research. Manuscript in press, *Psychol. Assess.*
- Tackett JL, Lilienfeld SO, Patrick CJ, Johnson SL, Krueger RF, Miller JD, et al. 2017. It's time to broaden the replicability conversation: Thoughts for and from clinical psychological science. *Perspect. Psychol. Sci.* 12(5): 742–56. <https://doi.org/10.1177/1745691617690042>

- Tackett JL & McShane BB. 2018. Conceptualizing and evaluating replication across domains of behavioral research. Manuscript in press, *Behav. Brain Sci.*
- Tomarken AJ & Waller NG. 2003. Potential problems with "well fitting" models. *J. Abnorm. Psychol.* 112: 578.
- Turner L, Shamseer L, Altman DG, Schulz KF, & Moher D. 2012. Does use of the CONSORT Statement impact the completeness of reporting of randomised controlled trials published in medical journals? A Cochrane review. *Syst. Rev.* 1(1).
<https://doi.org/10.1186/2046-4053-1-60>
- van der Lem R, de Wever WWH, van der Wee NJA, van Veen T, Cuijpers P, & Zitman F. G. 2012. The generalizability of psychotherapy efficacy trials in major depressive disorder: An analysis of the influence of patient selection in efficacy trials on symptom outcome in daily practice. *BMC Psychiat.* 12(192).
<https://doi.org/10.1186/1471-244X-12-192>
- Van Essen DC, Smith SM, Barch DM, Behrens TEJ, Yacoub E, & Ugurbil K. 2013. The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80: 62–79.
<https://doi.org/10.1016/j.neuroimage.2013.05.041>
- von Wolff A, Jansen M, Hölzel LP, Westphal A, Härter M, & Kriston L. 2014. Generalizability of findings from efficacy trials for chronic depression: An analysis of eligibility criteria. *Psychiat. Serv.* 65(7): 897–904.
<https://doi.org/10.1176/appi.ps.201300309>
- Walsh CG, Xia W, Li M, Denny JC, Harris PA, & Malin BA. 2018. Enabling open-science initiatives in clinical psychology and psychiatry without sacrificing patients' privacy: Current practices and future challenges. *Adv. Methods Pract. Psychol. Sci.* 1(1): 104–14.
<https://doi.org/10.1177/2515245917749652>

- Wampold BE, Mondin GW, Moody M, Stich F, Benson K, & Ahn H. 1997. A meta-analysis of outcome studies comparing bona fide psychotherapies: Empirically, “all must have prizes.” *Psychol. Bull.* 122(3): 203–15. <https://doi.org/10.1037//0033-2909.122.3.203>
- Weisz JR, Donenberg GR, Han SS, & Weiss B. 1995. Bridging the gap between laboratory and clinic in child and adolescent psychotherapy. *J. Consult. Clin. Psychol.* 63(5): 688–701. <https://doi.org/10.1037//0022-006X.63.5.688>
- Westfall J & Yarkoni T. 2016. Statistically controlling for confounding constructs is harder than you think. *PLoS ONE* 11(3): 1–22. <https://doi.org/10.1371/journal.pone.0152719>
- Wicherts JM, Borsboom D, Kats J, & Molenaar D. 2006. The poor availability of psychological research data for reanalysis. *Am. Psychol.* 61(7): 726–8. <https://doi.org/10.1037/0003-066X.61.7.726>
- Williams M, Grajales CAG, & Kurkiewicz D. 2013. Assumptions of multiple regression: Correcting two misconceptions. *Prac. Assess. Res. Eval.* 18(11): 1–14.
- Yates BT & Taub J. 2003. Assessing the costs, benefits, cost-effectiveness, and cost-benefit of psychological assessment: We should, we can, and here’s how. *Psychol. Assess.* 15(4): 478–95. <https://doi.org/10.1037/1040-3590.15.4.478>
- Zhao X, Lynch Jr. JG, & Chen Q. 2010. Reconsidering Baron and Kenny: Myths and truths about mediation analysis. *J. Consum. Res.* 37(2): 197–206. <https://doi.org/10.1086/651257>
- Zima BT, Hurlburt MS, Knapp P, Ladd H, Tang L, Duan N, et al. 2005. Quality of publicly-funded outpatient specialty mental health care for common childhood psychiatric

disorders in California. *J. Am. Acad. Child Adol. Psychiat.* 44(2): 130-144.

<https://doi.org/10.1097/00004583-200502000-00005>

Zwaan RA, Etz A, Lucas RE, & Donnellan MB. 2017. Making replication mainstream. *Behav.*

Brain Sci. 41: e120. <https://doi.org/10.1017/S0140525X17001972>.

.