Published as Christopherson, C.D. (2016) Replication and open science: Tools for progress in psychotherapy research. *Psychotherapy Bulletin*, 51 (3), 19-23.

Please Note: This manuscript is posted for noncommercial use. It is the pre-publication version and, as such, may contain errors that are not present in the final publication.

PSYCHOTHERAPY RESEARCH

Replication and Open Science: Tools for Progress in Psychotherapy Research

Cody D. Christopherson
Department of Psychology
Southern Oregon University

Replication has been a recent hot topic in Psychology research. With all of the concerns that have been raised, many of us may wonder how replication problems will impact practitioners and psychotherapy researchers. The purpose of this article is to review some recent research on publication and replication. I will make suggestions and argue that open science principles and replication will lead to a healthier and more progressive psychotherapy literature.

In August 2015, *Science* published “Estimating the Reproducibility of Psychological Science,” the first fruit of the Replication Project: Psychology or RP:P (Open Science Collaboration, 2015). Surprisingly, fewer than half of the 100 attempted studies obtained the same results as the original. Though the RP:P met its own criticisms (Gilbert, King, Pettigrew, & Wilson, 2016), the findings have led to questions about the reliability of research. In this, psychology is not alone. Similar conversations about reliability of research are happening in medicine, biology, and economics, among others, and the causes of low reproducibility seem to be endemic to many fields of scientific research (Ioannidis, 2005).

Problems in Psychotherapy Research

The fundamental problem is challenging to address, but easy to identify: Incentives for researchers may have undermined the reliability of research. Specifically, a premium has been put on publishing first and publishing in quantity. This gives rise to at least three sub-problems: publication bias, p-hacking, and hypothesizing after results are known (or HARKing). I will briefly explain these and apply them to psychotherapy research below, after which I will describe four solutions.

Problem 1: Publication bias. Non-significant findings are common but unlikely to be published. Therefore, the published literature becomes biased in favor of studies with significant results. Attempts to synthesize the literature (for example, with meta-analysis) are only as valid as the literature being summarized. Imagine we want to know whether diaphragmatic breathing reduces symptoms of schizophrenia. Five studies show a positive effect and 15 studies have failed to show such an effect, so we would have little hope for this intervention. However, the five

positive studies are more likely to be published than the 15 negative studies. Imagine all five positive studies were published and just one of the negative studies. After responsibly reviewing the peer-reviewed literature, we might conclude the use of diaphragmatic breathing as a clinical intervention will reduce symptoms of schizophrenia. This problem has important implications for psychotherapy practice. At that point, based on the *published* research, one might conclude it would be unethical NOT to use this proven approach; however, based on the *entire* body of research, the treatment is questionable at best. This illustration is fanciful, but there is evidence of publication bias in real psychotherapy research; to take one example, our best estimates for the effects of psychotherapy on depression appear substantially influenced by publication bias (Flint, Cuijpers, Horder, Koole, & Munafò, 2015).

Problem 2: P-hacking. P-hacking is the manipulation of data or analyses until reaching a p-value of less than .05. The motivation: p generally must be less than .05 to publish. P-hacking can be accomplished through many different methods, including collecting data until $p < .05$, eliminating outliers to decrease p , selectively choosing outcomes to demonstrate significance, and so on.

P-hacking is facilitated by a researcher having many options in post-hoc analysis. These options, or *researcher degrees of freedom*, can lead to outright incorrect results (Simmons, Nelson, & Simonsohn, 2011). Consider a real measure of aggression called the “Competitive Reaction Time Task,” in which participants “blast” one another with sound in multiple trials. According to a researcher tracking this instrument, it has been used in 122 publications but, amazingly, has been quantified in 150 different ways. Quantifications of the sound blast include volume in the first trial, duration in the first trial, volume times duration in the first trial, the mean of volume times duration in the first 25 trials, the average volume in trials 14 to 19, and so on (Elson, 2016). For a single study of this sort, there are at least 150 different ways to try to achieve $p < .05$. There is some evidence this is a concern for psychotherapy research. For example, psychotherapy outcome studies published over a three-year period averaged nearly four different outcomes measures each, with some using up to 14 (Ogles, 2013). The use of multiple outcomes measures is not itself wrong, but without preregistration (see below), it does greatly increase researcher degrees of freedom and risk of p-hacking.

Problem 3: Hypothesizing after results are known (HARK). HARKing distorts both statistical methods and creates a problem of infinite theoretical support (Kerr, 1998). A hypothetical example: A researcher studies whether mindfulness training influences daily stress. Without making a prediction before data collection, the data become difficult to profitably interpret. If mindfulness training decreases stress, then we might say it is a useful therapeutic tool. If it increases stress instead, then we can show how mindfulness training helps people become more sensitive to their daily stress, which could also be considered to be therapeutic. But what if the real effect of mindfulness training, overall, is nothing at all? For some people it reduces stress, for others it increases it in equal measure. Many studies may be published on both sides and it may be years before we can collectively know the real average effect of mindfulness training. If researchers are HARKing, it is difficult to know for sure because any finding can be supported by some theory post-hoc. If the true effect is zero, many studies may be published on both sides. HARKing and the other two problems above can lead to debates that may never be settled. Practitioners and clients are not well-served by this needless confusion.

Potential Solutions

The above are examples of possible problems with our psychotherapy literature. They are blind spots caused by our own traditions and common practices. As intractable as these problems may seem, there are proposed solutions which would be effective if adopted broadly. In this, individual labs as well as journals, granting agencies, and universities have an obligation to act to improve our field.

Solution 1: Pre-register. Pre-registration can involve creating a time-stamped record of your variables, target N for adequate power, and planned analyses. It solves at least two problems. First, it creates a record of all studies so unpublished studies can be incorporated into a meta-analysis. It also eliminates p-hacking. A pre-registration analysis plan includes a dated document with a target N for adequate power, stopping rules, and planned analyses. Exploratory analyses can be identified as such.

This is a powerful tool. In one survey of published antidepressant trials, 94% of studies showed the drug effective. Because drug trials require registry with the U.S. Food and Drug Administration (FDA), it was possible to discover that only 51% of studies had a positive result (Turner, Matthews, Linardatos, Tell, & Rosenthal, 2008). However, even with the required registry, drug trials commonly switch outcomes between proposal and publication (Goldacre et al., 2016.) To create a strong literature, we need a way to ensure outcomes are identified in advance. When they are changed, the rationale should be in the report. We do not yet know what elements of our literature might be subject to a bias similar to those in drug trials.

Though powerful, pre-registration is rarely used. Of all randomized controlled trials (RCTs) in the top 25 journals in clinical psychology published in 2013, only 15% were pre-registered and 44% registered overall (Cybulski, Mayo-Wilson, & Grant, 2016). Unlike treatments regulated by the FDA, there is not a regulatory body to enforce the registration of psychological treatments, though the organizations and guidelines exist. Consolidated Standards of Reporting Trials (CONSORT) is one such set of voluntary guidelines (Schulz, Altman, & Moher, 2010). Over 600 journals have agreed to these guidelines, including about 12 psychology journals. However, CONSORT remains ideal, not enforceable. For example, in one study, fewer than 50% of articles in journals endorsing CONSORT defined their primary outcome in advance (Turner, Shamseer, Altman, Schulz, & Moher, 2012). To be effective field-wide, journals and granting agencies must eventually adopt and enforce a policy of pre-registration.

Badges are a related tool. Journals may opt to display badge symbols next to studies with pre-registration, open data, and open materials. This simple policy does meaningfully influence researcher behavior (Kidwell et al., 2016).

Solution 2: Share. Science is considered progressive and collaborative (“verify”), rather than reliant on authority or tradition (“trust”). The American Psychological Association (APA) Ethics Code requires sharing of data with other professionals to verify claims (APA, 2002); however, in a recent study, only 38% of authors responded to a request and reminder to send their data for re-analysis, including only 25% of authors in the Journal of Abnormal Psychology (Vanpaemel, Vermorgen, Deriemaeker, & Storms, 2015). More than half the time, researchers were

implicitly asked to trust rather than verify. Collaborative science does not exist without independent verification.

I recently asked a colleague whether he had considered incorporating any elements of open science in a new project. Alarmed, he asked, “You mean give away all our data before we publish it?” Open data may seem scary or unnecessary. Why do it? In economics, researchers attempted to replicate 60 studies using identical data and analysis code. Even with help from the original authors, the researchers got the same results only half the time (Chang & Li, 2015). Without open data, we do not know the extent of this problem in our literature.

There was a time when data archives were unwieldy punch cards and expensive paper journal space could not be used to detail explicit procedures adequate for replication. Today, the Internet makes data sharing and storage extremely convenient. Online articles and supplements make space issues disappear. It is simple to share data, protocol, and registration with tools such as the Open Science Framework (which may be accessed online at: www.osf.io).

Sharing can also extend to articles themselves. The public cannot access most research without paying high prices for articles. A problem specific to clinical researchers is that the divide between practitioner and researcher grows ever wider if we cannot freely communicate. Though it is difficult to estimate, one common claim is that takes 17 years for research to influence practice (Morris, Wooding, & Grant, 2011). Thousands of studies have been designed, funded, and executed in order to shape clinical practice. What if clinicians could easily access them? Open access research may finally help close the perennial gap.

Solution 3: Determine the extent and the location of the problem. In RP:P, not all subfields were equally replicable. Cognitive psychology results matched the original approximately half the time, but for social psychology it was about 25% (OSC, 2015). We do not know which areas of psychotherapy research are prone to the problems described above, but there are ways to probe certain areas without conducting replications. Statistical examples of this probing include p-curve analysis (Simonsohn, Nelson, & Simmons, 2014) and looking at the relationship of effect size and sample size in a meta-analysis via funnel plot and sensitivity analysis (Copas & Shi, 2000). Other ways of probing might include a multi-method, meta-science approach. For example, we can survey or interview researchers regarding research practices. This can generate self-awareness with thick understandings of research practices. As an aside, we might also consider taking qualitative and mixed methods research more seriously in general. Qualitative research is not subject to the problems associated with significance testing.

Solution 4: Replicate directly. Direct replication means recreating the original experiment as precisely as possible. This differs from a conceptual replication which shares a theory with the original study but changes a variable. For example, it may use novel means of inducing the experimental condition. These two types of replications answer important, but different, questions. Direct replication answers, “Is this finding reliable?” Conceptual replication answers, “Is this finding generalizable?” Conceptual replication is common in psychology, independent pre-registered direct replication is rare.

Direct replication avoids endless theory protection. If a conceptual replication fails, the theory is not hurt. There may be some other reason for failure. If an experiment is powerfully and directly replicated and it fails, then that represents a true threat to the theory. From a positivist point of view, the only way for science to advance is to put our theories at grave danger (Meehl, 1978). Without direct replication, we risk building impenetrable belts around our theories.

Barriers to Replication

One reason for focusing on original research rather than replication is that psychotherapy research is difficult and expensive. Specific challenges include recruiting a clinical sample; assuring adequate treatment training and fidelity for clinicians; spending the time necessary to administer the treatment and measure follow-up; addressing high drop-out rates; dealing with the inherent risk in using a vulnerable population; and identifying a true placebo.

As an alternative to many early psychotherapy direct replications, one researcher proposes adopting the methods of the FDA in drug trials (Coyne & Kok, 2014). This approach contains three phases of research. The first two phases filter out treatments that are unsafe, unfeasible, unreliable, or unacceptable to patients. Following Phase 3, an adequately powered RCT, independent pre-registered replications can be conducted.

Another way to progress without conducting a full-scale clinical trial replication is to use specific ingredients of a therapeutic approach, particularly those that should show an effect in a nonclinical population as well as a clinical population. Then, run the study, including multi-site independent direct replications, using a nonclinical population. Once a replicable effect has been identified, build up to the full therapy in a clinical population. Whether any of these ideas are eventually adopted, a direct replication should be considered a gold standard that is merited after research has passed hurdles such as successful pre-registration, demonstrating adequate power, and providing open data (for verification purposes at a minimum.)

Impact and the Future

The delivery of psychotherapy has been broadly impacted by the empirically based treatment movement. Research on empirically based relationships and therapeutic processes has also been influential on the field of clinical training, and continues to grow. How supported is our empirical support itself? We do not know the degree to which each aspect of clinical research, from common factors and therapeutic alliance to assessment and diagnosis to outcome, has been influenced by the above considerations. But we also cannot dismiss these as potential concerns. The practices that have given rise to low replicability are present in clinical research. We have tools to help establish the veracity of our more important claims. The onus is on us to police ourselves. To improve, we need collaboration with journals and granting agencies willing to commit to publishing and funding replications to accompany new studies.

This is a crossroads and a wonderful time to be a psychological scientist. The internet has not yet changed everything about research methods, but it will. Memory will continue to get cheaper and it will be trivial to store and share mountains of data. I predict that within one generation, some version of open science practices will be the standard in psychology. This will certainly benefit our discipline because it will certainly, ultimately, benefit clients.

References

- American Psychological Association. (2010). *Ethical principles of psychologists and code of conduct: 2010 amendments*. Retrieved from <http://www.apa.org/ethics/code/index.aspx>
- Chang, A. C., & Li, P. (2015). Is economics research replicable? Sixty published papers from thirteen journals say 'usually not'. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2669564
- Copas, J., & Shi, J. Q. (2000). Meta-analysis, funnel plots and sensitivity analysis. *Biostatistics*, 1, 247–262. doi: 10.1093/biostatistics/1.3.247
- Coyne, J. C., & Kok, R. N. (2014). Salvaging psychotherapy research: a manifesto. *Journal of Evidence-Based Psychotherapies*, 14(2), 105-124.
- Cybulski, L., Mayo-Wilson, E., & Grant, S. (2016). Improving transparency and reproducibility through registration: The status of intervention trials published in clinical psychology journals. *Journal of Consulting and Clinical Psychology*. doi: 10.1037/ccp0000115
- Elson, M. (2016). FlexibleMeasures.com: Competitive Reaction Time Task (CRTT). Retrieved July 7, 2016, from <http://crtt.flexiblemeasures.com/>
- Flint, J., Cuijpers, P., Horder, J., Koole, S. L., & Munafò, M. R. (2015). Is there an excess of significant findings in published studies of psychotherapy for depression? *Psychological Medicine*, 45(2), 439–446.
- Gilbert, D., King, G., Pettigrew, S., & Wilson, T. (2016) Comment on 'Estimating the reproducibility of psychological science', *Science*, 351(6277), 1037a-1037b.
- Goldacre, B., Drysdale, H., Powell-Smith, A., Dale, A., Milosevic, I., Slade, E., ... Heneghan, C., (2016). The COMPare Trials Project. Retrieved July 19, 2016, from <http://compare-trials.org>
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124. Retrieved from <http://www.plosmedicine.org/article/info:doi/10.1371/journal.pmed.0020124>
- Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2, 196–217.
- Kidwell, M. C., Lazarević, L. B., Baranski, E., Hardwicke, T. E., Piechowski, S., Falkenberg, L.-S., ... Nosek, B. A. (2016). Badges to acknowledge open practices: A simple, low-cost, effective method for increasing transparency. *PLOS Biology*, 14, e1002456. doi: 10.1371/journal.pbio.1002456

- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806-834.
- Morris, Z. S., Wooding, S., & Grant, J. (2011). The answer is 17 years, what is the question: Understanding time lags in translational research. *Journal of the Royal Society of Medicine*, 104, 510–520. doi: 10.1258/jrsm.2011.110180
- Ogles, B. M. (2013). Measuring change in psychotherapy research. In Lambert, M. J. (Ed). *Bergin and Garfield's handbook of psychotherapy and behavior change*. 134–166 John Wiley & Sons. Hoboken, NJ.
- Open Science Collaboration (2015). Estimating the reproducibility of psychological. *Science*, 349, aac4716. doi: 10.1126/science.aac4716
- Schulz, K. F., Altman, D. G., & Moher, D (2010). CONSORT 2010 Statement: Updated Guidelines for Reporting Parallel Group Randomised Trials. *PLOS Medicine*, 7(3), e1000251. doi: 10.1371/journal.pmed.1000251
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359–1366. doi: 10.1177/0956797611417632
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143, 534–547. doi: 10.1037/a0033242
- Turner, E. H., Matthews, A. M., Linardatos, E., Tell, R. A., & Rosenthal, R. (2008). Selective publication of antidepressant trials and its influence on apparent efficacy. *New England Journal of Medicine*, 358, 252–260.
- Turner, L., Shamseer, L., Altman, D. G., Schulz, K. F., & 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. *Systematic Reviews*, 1, 60. doi: 10.1186/2046-4053-1-60
- Vanpaemel, W., Vermorgen, M., Deriemaeker, L., & Storms, G. (2015). Are we wasting a good crisis? The availability of psychological research data after the storm. *Collabra* 1:1–5 doi: <http://doi.org/10.1525/collabra.13>