What Psychology Teachers Should Know about
Open Science and the New Statistics

Beth Morling

University of Delaware

Robert J. Calin-Jageman

Dominican University

In press, *Teaching of Psychology* (Generalists' Corner)

Author Note

Beth Morling, University of Delaware and Robert J. Calin-Jageman, Dominican University.

Correspondence concerning this article should be addressed to Beth Morling, Psychological and Brain Sciences, University of Delaware, Newark, DE, 19716 (email: morling@udel.edu) or to Robert Calin-Jageman, Psychology Department, Dominican University, River Forest, IL, 60305 (email: rcalinjageman@dom.edu)

Abstract

Psychology teachers have likely heard about the “replication crisis” and the “open science movement” in psychology, and they are probably aware that psychologists have proposed new standards for research practice. How should our psychology courses reflect these new practices? We describe several modern practices that have transformed our field and that seem likely to endure: preregistration of studies, transparency of reporting, norms for replication, and the new statistical focus on estimation and precision. We offer suggestions for how to integrate these new practices into psychology courses.

Keywords: College teaching, open science, research transparency, estimation thinking, new statistics

What Psychology Teachers Should Know about

Open Science and the New Statistics

As someone who has been doing research for nearly twenty years, I now can't help but wonder if the topics I chose to study are in fact real and robust. Have I been chasing puffs of smoke for all these years? . . . I'm in a dark place. I feel like the ground is moving from underneath me and I no longer know what is real and what is not (Inzlicht, 2016)

Many of us can relate to Dr. Inzlicht's despair. Psychology's "replication crisis" has been at the forefront of science journalism, psychology conferences, and social media for several years now (Open Science Collaboration, 2015, 2017). Some of psychology's most well-known studies have been marked by a failure to replicate and have left us wondering which studies are "real and robust" and which are "puffs of smoke." For example, the facial feedback hypothesis (Strack, Martin, & Stepper, 1988) was not replicable in a large, multi-lab sample (Wagenmakers, et al., 2016; see Strack, 2016). Similarly, although many studies support ego depletion theory (Baumeister, 2014), at least one large replication attempt failed (Hagger et al., 2016; but see Baumeister, 2019). The effects of "power posing" (Carney, Cuddy, & Yap, 2010) have also been elusive except when using self-report measures (Ranehill, Dreber, Johannesson, Leiberg, & Sul, 2015; Simmons & Simonsohn, 2017; see also Cuddy et al., 2018). Even seemingly uncontroversial examples—such as the hypothesis that people eat more when using larger plates—have not survived replication attempts (Kosıte et al., 2019).

Importantly, a disappointing replication result does not mean that the original finding is wrong. However, a failed replication does suggest that these phenomena are more complex or nuanced than might be reflected in current textbooks.

Psychology's Methodological Upheaval: Problems and Solutions

Stories of failed replications have co-evolved with a multi-dimensional “credibility revolution” (Vazire, 2018). As psychologists inquired why studies failed to replicate, they identified questionable research practices behind the unstable findings (Chambers, 2017). Formerly common research practices, such as HARKing (hypothesizing after the results are known), *p*-hacking, and small samples are now questioned. In their place, new practices are becoming standard,

Publication practices are changing, too. For example, journals are devoting more space to replication studies (e.g., Lindsay, 2015, 2017; Makel, Plucker, & Hegarty, 2012). Similarly, because of the limitations of null hypothesis significance testing (NHST), many journals are now requiring confidence intervals and effect sizes instead of *p* values (Cumming, 2014). Although not all psychologists agree about the existence or severity of the problem (e.g., Gilbert, King, Pettigrew, & Wilson, 2016), one fact seems clear: Psychological science is experiencing a period of rapid methodological change.

Scientific and quantitative reasoning are core learning outcomes of many psychology courses (see, e.g., the APA Introductory Psychology Initiative; APA, 2019). Therefore, psychology teachers should help students practice scientific reasoning skills. Here we discuss two important themes: research transparency and estimation thinking. Figure 1 provides a pictorial introduction to the concepts in each theme.

Blurring the Line Between Confirmatory and Exploratory Research

The scientific method is meant to protect us from our own biases, to see the world as it truly is, not just as we would like to see it (to paraphrase Bacon, 1620/1889; see also Feynman, 1974). Science, though, is never complete--researchers are constantly discovering new ways biases can creep in and develop new protections against these biases (Nuzzo, 2015). The current

period of rapid methodological change reflects how psychologists have identified new sources of error and quickly adopted potential remedies.

These remedies reassert that openness and transparency form the core of scientific practice and help redraw the bright line between confirmatory and exploratory research. Confirmatory research, depicted in psychology textbooks as the “theory-data cycle,” or the hypothetico-deductive model (Chambers, 2017), typically proceeds as follows: Psychologists construct clear tests of their hypotheses, state these hypotheses in advance, collect and analyze their data objectively, and make results public even when the results fail to support the theory. Unfortunately, psychologists have often adopted questionable practices at each stage of this cycle (Chambers, 2017), resulting in a blurred line between confirmatory and exploratory research. These practices include underreporting nonsignificant effects, *p*-hacking, and HARKing. Before addressing how teachers might integrate these issues, we explain why each is problematic.

Blurring the Line by Underreporting

One questionable practice involves selective reporting of significant effects. For example, a researcher may include multiple dependent variables (DVs) in a study. When making results public, the researcher might report only the DVs that supported the hypothesis. There is nothing wrong with administering multiple dependent measures, but this practice becomes misleading when the researcher never reports outcomes that did not support predictions. The researcher has misrepresented an *exploratory* study as *confirmatory*.

An alleged example of this practice comes from Rohrer, Pashler, and Harris (2015), who attempted to replicate a study on money and cognition. The original researchers found that exposure to money caused people to endorse attitudes that justify inequality (Caruso, Vohs,

Baxter, & Waytz, 2013). While corresponding with Rohrer and colleagues, Caruso et al. revealed that they had reported only nine of the 28 dependent variables they had measured. Rohrer et al. argued that by reporting only some of the DVs (only the ones that “worked”), the original authors misled readers about the strength of their evidence. By underreporting, Caruso et al. present their work as confirmatory when they had actually been *exploring* the effects of money on a wide variety of possible outcomes. (In response to Rohrer et al.’s failure to replicate, Vohs [2015] reviewed 10 years of studies on money primes and maintained their importance.)

Blurring the Line by *p*-hacking

A second practice that blurs the line between confirmatory and exploratory science is “*p*-hacking,” or “exploiting researcher degrees of freedom” (Chambers, 2017; Simmons, Nelson, & Simonsohn, 2018). During data analysis, researchers make dozens of decisions: whether to include outliers, include covariates in a statistical test, and so on. Researchers may consciously or unconsciously make choices that lead to significance. The term “*p*-hacking” reflects that researchers’ choices are driven by the desire to obtain a significant *p* value rather than by the goal of determining how well the data support the hypothesis (Simmons, Nelson, & Simonsohn, 2011; Simonsohn, Nelson, & Simmons, 2014). Many researchers argue—and rightfully so—that there is nothing inherently wrong with exploratory data analysis. However, such explorations become misleading when researchers fail to disclose all of the data permutations and statistical dead-ends they pursued.

Blurring the Line by HARKing

A third practice carries the disparaging name, HARKing, for *hypothesizing after the results are known* (Kerr, 1998). HARKing is the practice of presenting the process as if the

results were expected all along. To illustrate how widely accepted this practice once was, consider Bem's (2004) advice:

There are two possible articles you can write: (a) the article you planned to write when you designed your study or (b) the article that makes the most sense now that you have seen the results. They are rarely the same, and the correct answer is (b). (pp.171-172)

Bem argued that HARKing makes an empirical article easier to read. He is probably correct; however, as with underreporting and *p*-hacking, HARKed effects are less likely to be robust. HARKing, most clearly of all these practices, presents an exploratory process as if it were a confirmatory one.

The Transparency Solution

The solution to all three problems has been a call for radical transparency and openness in reporting of science (Table 1). This push for transparency is one reason people use the label Open Science Movement. Transparency applies to three areas: Method sections, hypotheses, and raw data. Journals have begun awarding "badges" to published work that uses these new standards (see Figure 1).

Transparent Methods Sections: Open Materials

Researchers are now expected to disclose every study detail. Some journals have eliminated word limits on Method sections so they can be as long as necessary for full disclosure. Many journals also enable the use of online supplementary materials, where researchers provide the complete list of conditions and variables used in their studies. An additional approach is the "21 Word Solution," with which researchers announce in their Method section that "We report how we determined our sample size, all data exclusions (if any), all manipulations, and all measures in the study" (Simmons, Nelson, Simonsohn, 2012). When

researchers disclose all of the conditions and variables they included, readers are better able to evaluate the strength of the evidence and science makes more progress.

Transparent Results: Open Data

Another form of transparency is known as “open data,” in which researchers publicly share the data files from published research. Open data also means that researchers share how they prepared the data and how they computed any composite scores. Before sharing data, researchers remove any identifying information and when data are impossible to anonymize, they might not be shared in their entirety. Some journals publish the data files with the published manuscript; other scholars publish their data separately on the Open Science Framework (osf.io).

There are multiple benefits to sharing data openly (Chambers, 2017). Open data can be re-analyzed to confirm the results. Open data helps address underreporting and *p*-hacking. It is easier to detect fraud because independent scientists can look for patterns in the data that are consistent with data fabrication. Data are stored for posterity, and other scientists can use an old dataset to test novel hypotheses. Finally, shared data seem ethically reasonable for studies funded by taxpayers, who arguably deserve to have access to the knowledge gained by the investment they have made.

Transparent Hypotheses: Preregistration and Registered Reports

Preregistration—the antidote to HARKing—involves answering specific, structured questions about study design, hypotheses, and data analysis. Researchers document these decisions in advance, often in great detail, even down to drafting the code and posting the syntax files needed to analyze the data. When filed through public repositories such as aspredicted.org or the Open Science Framework, registrations carry a timestamp. Therefore, it can be clearer that researchers collected and analyzed their data only after preregistering.

Registered reports take preregistration a step closer by integrating it with peer review. Researchers write a plan (Introduction, Method, and data analysis) and submit it to a journal before commencing data collection. Peer reviewers can recommend “conditional acceptance” if the study is deemed important, rigorous, and feasible. This process not only controls against HARKing and *p*-hacking; it also changes the incentives. Previously, a journal’s evaluation of a paper might have been based, in part, on the strength of its results. Indeed, journals have long been accused of preferentially publishing only significant results—a bias that can lead to the “file drawer problem” (Rosenthal, 1979). In contrast, registered reports are evaluated on the importance of the research question and the quality of the proposed methods for testing that question (e.g., Benjamin, 2019; Lindsay, 2017).

Teaching About Transparency

The practices we have explained here—underreporting, *p*-hacking, HARKing, transparency, registered reports, and data sharing—all belong in the curriculum for the modern psychology major. The details may overwhelm students in the introductory course, but in general, teachers should add the theme of transparent science.

Teach Preregistration

Teaching the theory-data cycle in psychology courses is essential: It reinforces that psychology is a science because researchers test their ideas with data. Teachers can modernize their coverage of the theory-data cycle by discussing the practice of preregistration (see Figure 2). Links to preregistration are usually available when an article is published.

Prompt Students to Predict

The modernized theory-data cycle merges nicely with the pedagogical technique of prediction. Prediction (rebranded as preregistration) is a simple and powerful teaching tool; in

order to make a prediction, students have to activate what they already know, consolidate it, and apply it. Predictions create positive emotions, such as interest and anticipation, that can facilitate learning (Kornell, Jenson, Hayes, & Bjork, 2009; Ogan, Aleven, Jones, 2009). Instructors can describe a theory, explaining the study procedure, and then asking students to predict the results. Students can preregister their predictions in a notebook, turn to a classmate, use a clicker, or even sketch a graph to document their predictions.

Discuss Transparency

Another suggestion for teaching about transparency is through brief discussions, perhaps jump-started by clicker questions. Specifically, after discussing the scientific method, teachers can ask students,

What should scientists do if they conduct a study and the results that do not support their hypothesis?

- a) ignore the result
- b) make the result publicly available
- c) revise the hypothesis as if it was expected all along
- d) something else

Most students will *not* select a) or c), even though these were common practices in psychology before the dawn of the credibility revolution. The discussion that follows is likely to be engaging and informative about the processes of science.

Ideas for Upper-Level Students

For statistics and methods students, the website fivethirtyeight.com produces an online tool called “[*p*-]hack your way to scientific glory.” As students investigate the correlation

between political party and economic growth, they can try different combinations of variable operationalizations until they get a p value under .05.

Upper-level statistics and methods students could replicate published data analyses using open data. For example, one study that has simple methods and open data found that people were less likely to cheat when manipulated into the emotional states of gratitude (DeSteno, Duong, Lim, & Kates, 2019). The article has an open data badge, so students can recreate the tables reported in the original article (open data link: <https://osf.io/pzvue/>)

Supplemental Materials

We offer three suggestions for supplemental content assignments to introduce the replication crisis. One is artist Miki Naro's graphic depiction of the key moments, called "Repeat after Me" (Naro, 2016). The second is a video from the HBO comedy show *Last Week Tonight* (search "Scientific Studies" but warn students of adult language). Finally, several journalists have covered the replication crisis. One excellent example is Daniel Engber (2017), who described how Bem and colleagues' (2010) ESP study arguably started it all.

Estimation Thinking

The second area of rapid change is in the area of statistical inference. Here, the key development is an increasing emphasis on psychology as a quantitative endeavor that accumulates knowledge through replication and synthesis. Often described as the "New Statistics" or the "Estimation Approach" (Cumming, 2012, 2014), it involves reporting and interpreting effect sizes ("How much?"), countenancing uncertainty in all statistical conclusions ("How wrong?"), and synthesizing results through meta-analysis ("What else is known?").

The estimation approach stands in contrast to the decision-making approach, which has dominated psychology (see Table 2). In the decision-making approach, researchers ask a

qualitative, binary (Yes/No) question (“Do antidepressants treat depression?”) and use a p value to make a qualitative conclusion (“Yes, antidepressants reduce symptoms of depression,” $p < .0001$). Researchers often treat their results as definitive, so there is little motivation to conduct replications. Moreover, there is often scant attention to practical significance. To be clear, this is *not* the way p values were meant to be used (e.g., Fisher, 1926), but this “one-and-done” approach to statistical inference has, until recently, been pervasive.

Estimation offers a different lens for interpreting data, one that should nudge researchers toward more thoughtful and nuanced conclusions. Rather than a binary question, estimation focuses on quantifying effects (“To what extent do antidepressants treat depression?”). Results focus not on a p value but rather on an effect size and an expression of uncertainty (“The average improvement with antidepressants was 10%, 95% CI [7%, 13%]”; Horder, Matthews, & Waldmann, 2010). Estimation—which emphasizes uncertainty—leads to meta-analysis, which combines data to reduce the uncertainty. Estimation thinking is a broad statistical tradition encompassing both parametric and nonparametric techniques as well as Bayesian and classical approaches to probability.

Teaching Estimation Thinking

Introducing estimation into the undergraduate curriculum is overdue. Since the 5th edition of the APA Publication Manual, the APA (2001) has recommended estimation as “the best reporting strategy” (p. 22; see Fidler, 2010). Specifically, the APA (2020) enjoins authors to “wherever possible, base discussion and interpretation of results on point and interval estimates” (p. 88; see Applebaum et al., 2018). Following on these recommendations, many journals require the use of the estimation approach (see Giofrè, Cumming, Fresc, Boedker, & Tressoldi, 2017),

and textbooks are increasingly teaching estimation as the default choice for reporting and interpreting results (e.g., Harrington, 2020; Morling, 2020).

Introductory coursework is the ideal time to foster estimation thinking. Teachers can use the prompt, “How much?” to help students consider the magnitudes of effects and to seek context. Using the prompt, “How wrong?” can encourage students to embrace uncertainty and to introduce the key idea of sampling variation. Finally, prompting students with, “What else is known?” helps them see science as cumulative and integrative process rather than a series of “one-and-done” demonstrations. These three questions instill a nuanced view of science, where any one study is tenuous, and yet the cumulative evidence from a body of research can be compelling. This is a sophisticated epistemic viewpoint that avoids both excessive confidence and undue cynicism.

Certain psychology courses afford little time to cover statistical inference. Will students miss out if teachers emphasize estimation over decision-making with p values? No. In fact, the limited coverage of statistics provided in many textbooks is often incorrect, so it probably instills misconceptions rather than accurate knowledge (Cassidy et al., 2019). By emphasizing estimation, students gain practice in nuanced quantitative reasoning and also gain a foundation for understanding decision-making. Specifically, students will come to learn that a “significance test” is equivalent to checking for the null hypothesis in the confidence interval. Estimation gives students a foundation for understanding statistical significance that can help them avoid common misconceptions (e.g., thinking that statistical significance means a certain result is replicable).

Start With Opinion Polls

Teachers can introduce their students to the basics of estimation through political or attitudes polling. Students are used to seeing single polling results accompanied by a margin of

error. The polling estimate (52% support a new referendum) is “how much,” and the margin of error (plus or minus 4%) is “how wrong.” Online poll aggregators such as fivethirtyeight.com synthesize polls conducted on the same topic, which address “what else is known.”

Use Visualizations

Data visualizations from polling aggregators can help make the reality of sampling variation salient and illustrate how consensus can emerge across multiple polls. Teachers can go further by introducing the forest plot, which summarizes results across a meta-analysis. Forest plots allow students to visualize all three estimation-thinking points. The effect size of each study shows “how much.” Each study includes a line indicating the confidence interval, or “how wrong.” Finally, showing every study and plotting a meta-analytic effect size indicates “what else is known.” In sum, a forest plot illustrates variation across studies as well as the emerging consensus across results.

As an example, Figure 3 shows a forest plot from a meta-analysis of the relation between sleep quality and children’s access to media devices at bedtime (Carter et al., 2016). Despite considerable variability across studies, all of them show that access to a media device at bedtime is associated with inadequate sleep quality. Overall, children with access to a media device were 2.17 times more likely to have poor sleep. Even combining across studies, there is still uncertainty, with the 95% CI ranging from 1.41 to 3.32 times the risk. Values outside this CI are also possible. Students could debate the meaning of this relation and whether parents should ban devices from the bedroom.

Overall, working with students to understand a variety of data visualizations is time well spent. Students can develop the ability to interpret, use, and think critically about quantitative

information in an increasingly data-visualized world. A second advantage is that students can practice navigating through a figure using the caption and axis labels.

Explore Simulations

Interactive simulations can help students conceptualize sampling variation. “The Dance of the Means” lets students explore how samples drawn from the same population “dance” around the population mean (Cumming, 2012; <https://tinyurl.com/danceofthameans>; see also <https://rpsychologist.com/d3/CI/>). Such simulations illustrate that even when scientists replicate a study, results vary due to sampling error, with the degree of variation determined in part by sample size. Moreover, students witness how scientists can take this variation into account, reporting an estimate of uncertainty (a confidence interval) that allows most studies to capture the truth despite sampling variation.

Integrate Estimation Thinking Throughout the Semester

How might teachers reinforce estimation thinking throughout the semester? First, seek opportunities to discuss effect sizes and confidence intervals. For example, while discussing the relation between testosterone and aggression, rather than adopting qualitative language (testosterone *is* associated with aggression in humans), teachers can model estimation thinking by focusing on effect sizes (testosterone predicts about 2% of variation in human aggression; Book et al., 2001).

Second, highlight uncertainty by providing margins of error or confidence intervals for classic studies. In fact, it can be useful to have students rate their confidence in a study from a qualitative textbook presentation and again after being presented with the effect size and uncertainty. For example, if teachers discuss the facial feedback study (Strack, et al., 1998), they might share that the original effect was quite uncertain: On a 10-point scale, ratings increased by

0.82 with a margin of error of 0.87 points, giving a 95% confidence interval of $[-0.05, 1.69]$. In this light, it is actually not surprising that a large-scale replication found an increase of only .03 points; Wagenmakers et al., 2016). Reflecting on uncertainty can help students understand the essential role of replication in science and also develop intuitions for which initial results are most likely to hold up to replication.

A third way to reinforce estimation thinking is to use forest plots and meta-analyses throughout the semester (e.g., Wagenmakers et al. [2016] included a forest plot of replication studies similar to Figure 3). Most textbooks mention only one or two key studies as evidence for a theory. A forest plot can meaningfully extend that coverage and show what happened next. In fact, if teachers help students understand uncertainty in the initial studies, students might even demand the meta-analysis!

One difficulty with these recommendations is that, for now, textbooks tend to omit effect sizes and uncertainty, presenting key studies in purely qualitative terms. Fortunately, they often provide key graphs or summary data from which teachers can extract effect sizes and uncertainty (where they do not, Google Scholar can help surface this information). There are also tools for obtaining effect sizes, confidence intervals, and figures from summary data (see, e.g., <http://thenewstatistics.com>; Cumming & Calin-Jageman, 2017).

Critical thinking involves looking for context, and effect sizes are difficult to discuss in isolation. For example, a recent meta-analysis showed that electroconvulsive therapy produces more improvement than no treatment by an average of 9.7 points on the Hamilton Rating Scale for Depression (UK ECT Review Group, 2017). To contextualize these results, students need to learn about the scale: that scores range from 0 to 52 and that a score of at least 21 is required to indicate depression (Hamilton, 1960; Sharp, 2015). It also helps to provide context from other

effect sizes (Figure 4). Providing context can help students think deeply about the scientific issues involved: Are these comparisons meaningful across different types of patients? Is treatment with the biggest effect size recommended, or should costs and side effects also be considered?

Closing

In this piece, we have provided teachers with content updates about methodological changes of the past several years. We suggested that teachers should emphasize two fundamental themes: research transparency and estimation thinking. By introducing the value of transparency and the practice of estimation thinking early on, teachers can better prepare students to reason scientifically and quantitatively, preparing them for a variety of future paths.

References

- American Psychological Association. (2001). *Publication Manual of the American Psychological Association* (5th ed.). Washington, DC: Author.
- American Psychological Association. (2020). *Publication Manual of the American Psychological Association* (7th ed.). Washington, DC: Author.
- Appelbaum, M., Cooper, H., Kline, R. B., Mayo-Wilson, E., Nezu, A. M., & Rao, S. M. (2018). Journal article reporting standards for quantitative research in psychology: The APA Publications and Communications Board task force report. *American Psychologist*, 73(1), 3–25. doi:[10.1037/amp0000191](https://doi.org/10.1037/amp0000191)
- Bacon, F. (1620/1889). *New organon, or true directions concerning the interpretation of nature*. Oxford: Clarendon Press.
- Baumeister, R. (2019). *Self-Control, Ego Depletion, and Social Psychology's Replication Crisis* [Preprint]. doi: [10.31234/osf.io/uf3cn](https://doi.org/10.31234/osf.io/uf3cn)
- Baumeister, R. F. (2014). Self-regulation, ego depletion, and inhibition. *Neuropsychologia*, 65, 313–319. Doi: [10.1016/j.neuropsychologia.2014.08.012](https://doi.org/10.1016/j.neuropsychologia.2014.08.012)
- Bech, P., Cialdella, P., Haugh, M. C., Birkett, M. A., Hours, A., Boissel, J. P., & Tollefson, G. D. (2000). Meta-analysis of randomised controlled trials of fluoxetine v. placebo and tricyclic antidepressants in the short-term treatment of major depression. *British Journal of Psychiatry*, 176(5), 421–428. Doi: 10.1192/bjp.176.5.421
- Bem, D. J. (2004). Writing the empirical journal article In *The compleat academic: A career guide* (2nd ed.), by Darley, John M (ed.) ; Zanna, Mark P (ed.) ; Roediger, Henry L., III (ed.) . 185-219, Chapter xii, 422 Pages. 2nd ed.. American Psychological Association, 2004.
- Bem, D. J. (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, 100(3), 407–425. Doi: [10.1037/a0021524](https://doi.org/10.1037/a0021524)
- Benjamin, A. S. (2019). Editorial. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(2), 193–195. Doi: [10.1037/xlm0000688](https://doi.org/10.1037/xlm0000688)
- Bergelson, E., Bergmann, C., Byers-Heinlein, K., Cristia, A., Cusack, R., Dyck, K., ... Yurovsky, D. (2017). *Quantifying sources of variability in infancy research using the infant-directed speech preference* [Preprint]. doi: [10.31234/osf.io/s98ab](https://doi.org/10.31234/osf.io/s98ab)
- Boebel, W., Wagenmakers, E. J., Belay, L., Verhagen, J., Brown, S., & Forstmann, B. U. (2015). A purely confirmatory replication study of structural brain-behavior correlations. *Cortex*, 66, 115–133. doi: [10.1016/j.cortex.2014.11.019](https://doi.org/10.1016/j.cortex.2014.11.019)
- Book, A. S., Starzyk, K. B., & Quinsey, V. L. (2001). The relationship between testosterone and aggression: A meta-analysis. *Aggression and Violent Behavior*, 6(6), 579–599. doi: [10.1016/S1359-1789\(00\)00032-X](https://doi.org/10.1016/S1359-1789(00)00032-X)
- Border, R., Johnson, E. C., Evans, L. M., Smolen, A., Berley, N., Sullivan, P. F., & Keller, M. C. (2019). No support for historical candidate gene or candidate gene-by-interaction hypotheses for major depression across multiple large samples. *American Journal of Psychiatry*, 176(5), 376–387. doi: [10.1176/appi.ajp.2018.18070881](https://doi.org/10.1176/appi.ajp.2018.18070881)
- Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2010). Power Posing: Brief Nonverbal Displays Affect Neuroendocrine Levels and Risk Tolerance. *Psychological Science*, 21(10), 1363–1368. doi: [10.1177/0956797610383437](https://doi.org/10.1177/0956797610383437)

- Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2010). Power posing: Brief nonverbal displays affect neuroendocrine levels and risk tolerance. *Psychological Science*, 21(10), 1363–1368. <http://dx.doi.org/10.1177/0956797610383437>
- Carter, B., Rees, P., Hale, L., Bhattacharjee, D., & Paradkar, M. S. (2016). Association Between Portable Screen-Based Media Device Access or Use and Sleep Outcomes: A Systematic Review and Meta-analysis. *JAMA Pediatrics*, 170(12), 1202. [doi: 10.1001/jamapediatrics.2016.2341](https://doi.org/10.1001/jamapediatrics.2016.2341)
- Caruso, E. M., Vohs, K. D., Baxter, B., & Waytz, A. (2013). Mere exposure to money increases endorsement of free-market systems and social inequality. *Journal of Experimental Psychology: General*, 142(2), 301–306. [doi: 10.1037/a0029288](https://doi.org/10.1037/a0029288)
- Cassidy, S. A., Dimova, R., Giguère, B., Spence, J. R., & Stanley, D. J. (2019). Failing Grade: 89% of Introduction-to-Psychology Textbooks That Define or Explain Statistical Significance Do So Incorrectly. *Advances in Methods and Practices in Psychological Science*, 251524591985807. [doi: 10.1177/2515245919858072](https://doi.org/10.1177/2515245919858072)
- Chambers, C. (2017). *The Seven Deadly Sins of Psychology: A Manifesto for Reforming the Culture of Scientific Practice*. [doi: 10.1515/9781400884940](https://doi.org/10.1515/9781400884940)
- Cuddy, A. J. C., Schultz, S. J., & Fosse, N. E. (2018). P-curving a more comprehensive body of research on postural feedback reveals clear evidential value for power-posing effects: Reply to Simmons and Simonsohn (2017). *Psychological Science*, 29(4), 656–666. <http://dx.doi.org/10.1177/0956797617746749>
- Cumming, G. (2012). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. New York: Routledge.
- Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, 25(1), 7–29. <http://dx.doi.org/10.1177/0956797613504966>
- Cumming, G., & Calin-Jageman, R. (2017). *Introduction to the new statistics*. Routledge.
- DeSteno, D., Duong, F., Lim, D., & Kates, S. (2019). The Grateful Don't Cheat: Gratitude as a Fount of Virtue. *Psychological Science*, 30(7), 979–988. [doi: 10.1177/0956797619848351](https://doi.org/10.1177/0956797619848351)
- Dominus, S. (2017, October 18). When the Revolution Came for Amy Cuddy. *The New York Times*. Retrieved from <https://www.nytimes.com/2017/10/18/magazine/when-the-revolution-came-for-amy-cuddy.html>
- Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., ... Nosek, B. A. (2016). Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67, 68–82. [doi: 10.1016/j.jesp.2015.10.012](https://doi.org/10.1016/j.jesp.2015.10.012)
- Engber, D. (2017). Daryl Bem proved ESP is real: Which means science is broken. *Slate*. <https://slate.com/health-and-science/2017/06/daryl-bem-proved-esp-is-real-showed-science-is-broken.html>
- Feynman, R. (1974). *Cargo cult science*. Commencement address delivered at California Institute of Technology.
- Fidler, F. (2010). The American Psychological Association Publication Manual Sixth Edition: Implications for Statistics Education. *Proceedings of ICOTS-8, Eighth International Conference on Teaching Statistics*, 8.
- Fisher, R.A. (1926). The arrangement of field experiments. *Journal of the Ministry of Agriculture of Great Britain*, 33, 503–513.

- Fiske, A. P. P. S. T. (2016). A Call to Change Science's Culture of Shaming. *APS Observer*, 29(9). Retrieved from <https://www.psychologicalscience.org/observer/a-call-to-change-sciences-culture-of-shaming>
- Flake, J. K., Pek, J., & Hehman, E. (2017). Construct Validation in Social and Personality Research: Current Practice and Recommendations. *Social Psychological and Personality Science*, 8(4), 370–378. doi: [10.1177/1948550617693063](https://doi.org/10.1177/1948550617693063)
- Fraley, R. C., & 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. doi: [10.1371/journal.pone.0109019](https://doi.org/10.1371/journal.pone.0109019)
- Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on “Estimating the reproducibility of psychological science.” *Science*, 351(6277), No Pagination Specified.
- Giofrè, D., Cumming, G., Fresc, L., Boedker, I., & Tressoldi, P. (2017). The influence of journal submission guidelines on authors' reporting of statistics and use of open research practices. *PLOS ONE*, 12(4), e0175583. doi: [10.1371/journal.pone.0175583](https://doi.org/10.1371/journal.pone.0175583)
- Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., ... Zwienenberg, M. (2016). A Multilab Preregistered Replication of the Ego-Depletion Effect. *Perspectives on Psychological Science*, 11(4), 546–573. doi: [10.1177/1745691616652873](https://doi.org/10.1177/1745691616652873)
- Horder, J., Matthews, P., & Waldmann, R. (2011). Placebo, Prozac and PLoS: significant lessons for psychopharmacology. *Journal of Psychopharmacology*, 25(10), 1277–1288. doi: [10.1177/0269881110372544](https://doi.org/10.1177/0269881110372544)
- Hamilton, M. (1960). A Rating Scale for Depression. *Journal of Neurology, Neurosurgery & Psychiatry*, 23(1), 56–62. doi: [10.1136/jnnp.23.1.56](https://doi.org/10.1136/jnnp.23.1.56)
- Harrington, M. (2020). *The design of experiments in neuroscience* (3rd Ed.) Sage.
- Johnsen, T. J., & Friberg, O. (n.d.). The Effects of Cognitive Behavioral Therapy as an Anti-Depressive Treatment is Falling: A Meta-Analysis. doi: [10.1037/bul0000015](https://doi.org/10.1037/bul0000015)
- Kerr, N. L. (1998). HARKing: Hypothesizing After the Results are Known. *Personality and Social Psychology Review*, 2(3), 196–217. doi: [10.1207/s15327957pspr0203_4](https://doi.org/10.1207/s15327957pspr0203_4)
- Kharabian Masouleh, S., Eickhoff, S. B., Hoffstaedter, F., Genon, S., & Alzheimer's Disease Neuroimaging Initiative. (2019). Empirical examination of the replicability of associations between brain structure and psychological variables. *ELife*, 8, e43464. doi: [10.7554/eLife.43464](https://doi.org/10.7554/eLife.43464)
- King, J. B., Prigge, M. B. D., King, C. K., Morgan, J., Weathersby, F., Fox, J. C., ... Anderson, J. S. (2019). Generalizability and reproducibility of functional connectivity in autism. *Molecular Autism*, 10(1), 1–23. doi: [10.1186/s13229-019-0273-5](https://doi.org/10.1186/s13229-019-0273-5)
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Bahník, Š., Bernstein, M. J., ... Nosek, B. A. (2014). Investigating Variation in Replicability: A “Many Labs” Replication Project. *Social Psychology*, 45(3), 142–152. doi: [10.1027/1864-9335/a000178](https://doi.org/10.1027/1864-9335/a000178)
- Kornell, N., Hays, M. J., & Bjork, R. A. (2009). Unsuccessful retrieval attempts enhance subsequent learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(4), 989–998. doi: [10.1037/a0015729](https://doi.org/10.1037/a0015729)
- Kosıte, D., König, L. M., De-loyde, K., Lee, I., Pechey, E., Clarke, N., ... Hollands, G. J. (2019). Plate size and food consumption: A pre-registered experimental study in a general population sample. *International Journal of Behavioral Nutrition and Physical Activity*, 16(1), 75. doi: [10.1186/s12966-019-0826-1](https://doi.org/10.1186/s12966-019-0826-1)

- Lilienfeld, S. O. [Ed, & Waldman, I. D. [Ed (Eds.). (2017). Maximizing the reproducibility of your research. In *Psychological science under scrutiny: Recent challenges and proposed solutions* (pp. 3–21, Chapter xxv, 371 Pages). Retrieved from <https://search.proquest.com/psycinfo/docview/1915377042/E28C1674F6FF4E8DPQ/1>
- Lindsay, D. S. (2015). Replication in Psychological Science. *Psychological Science*, 26(12), 1827–1832. doi: [10.1177/0956797615616374](https://doi.org/10.1177/0956797615616374)
- Lindsay, D. S. (2017). Preregistered Direct Replications in *Psychological Science*. *Psychological Science*, 28(9), 1191–1192. doi: [10.1177/0956797617718802](https://doi.org/10.1177/0956797617718802)
- Makel, M. C., Plucker, J. A., & Hegarty, B. (2012). Replications in Psychology Research: How Often Do They Really Occur? *Perspectives on Psychological Science*, 7(6), 537–542. doi: [10.1177/1745691612460688](https://doi.org/10.1177/1745691612460688)
- Morling, B. (2020). *Research methods in psychology: Evaluating a world of information* (4th Ed.) W.W. Norton.
- Naro, M. (2016) Repeat after me. *The Nib.com*. <https://thenib.com/repeat-after-me>
- Nelson, L. D., Simmons, J., & Simonsohn, U. (2018). Psychology's Renaissance. *Annual Review of Psychology*, 69(1), 511–534. doi: [10.1146/annurev-psych-122216-011836](https://doi.org/10.1146/annurev-psych-122216-011836)
- Normile, C. J., Bloesch, E. K., Davoli, C. C., & Scherr, K. C. (2019). Introducing the new statistics in the classroom. *Scholarship of Teaching and Learning in Psychology*, 5(2), 162–168. doi: [10.1037/stl0000141](https://doi.org/10.1037/stl0000141)
- Nosek, B. A., Beck, E. D., Campbell, L., Flake, J. K., Hardwicke, T. E., Mellor, D. T., ... Vazire, S. (2019). Preregistration Is Hard, And Worthwhile. *Trends in Cognitive Sciences*, 23(10), 815–818. doi: [10.1016/j.tics.2019.07.009](https://doi.org/10.1016/j.tics.2019.07.009)
- Nuzzo, R. (2015). How scientists fool themselves – and how they can stop. *Nature News*, 526(7572), 182. doi: [10.1038/526182a](https://doi.org/10.1038/526182a)
- Ogan, A., Aleven, V., & Jones, C. (n.d.). *Advancing Development of Intercultural Competence through Supporting Predictions in Narrative Video*. 22.
- Open Science Collaboration. (2012). An Open, Large-Scale, Collaborative Effort to Estimate the Reproducibility of Psychological Science. *Perspectives on Psychological Science*, 7(6), 657–660. doi: [10.1177/1745691612462588](https://doi.org/10.1177/1745691612462588)
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716–aac4716. doi: [10.1126/science.aac4716](https://doi.org/10.1126/science.aac4716)
- Ranehill, E., Dreber, A., Johannesson, M., Leiberg, S., Sul, S., & Weber, R. A. (2015). Assessing the Robustness of Power Posing: No Effect on Hormones and Risk Tolerance in a Large Sample of Men and Women. *Psychological Science*, 26(5), 653–656. doi: [10.1177/0956797614553946](https://doi.org/10.1177/0956797614553946)
- Rohrer, D., Pashler, H., & Harris, C. R. (2015). Do subtle reminders of money change people's political views? *Journal of Experimental Psychology: General*, 144(4), e73–e85. doi: [10.1037/xge0000058](https://doi.org/10.1037/xge0000058)
- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86(3), 638–641. doi: [10.1037/0033-2909.86.3.638](https://doi.org/10.1037/0033-2909.86.3.638)
- Schuch, F. B., Vancampfort, D., Richards, J., Rosenbaum, S., Ward, P. B., & Stubbs, B. (2016). Exercise as a treatment for depression: A meta-analysis adjusting for publication bias. *Journal of Psychiatric Research*, 77, 42–51. doi: [10.1016/j.jpsychires.2016.02.023](https://doi.org/10.1016/j.jpsychires.2016.02.023)
- Sharp, R. (2015). The Hamilton Rating Scale for Depression. *Occupational Medicine*, 65(4), 340–340. doi: [10.1093/occmed/kqv043](https://doi.org/10.1093/occmed/kqv043)

- Simmons, J. P., & Simonsohn, U. (2017). Power Posing: *P*-Curving the Evidence. *Psychological Science*, 28(5), 687–693. [doi: 10.1177/0956797616658563](https://doi.org/10.1177/0956797616658563)
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, 22(11), 1359–1366. [doi: 10.1177/0956797611417632](https://doi.org/10.1177/0956797611417632)
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2012). A 21 Word Solution. *SSRN Electronic Journal*. [doi: 10.2139/ssrn.2160588](https://doi.org/10.2139/ssrn.2160588)
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2018). False-Positive Citations. *Perspectives on Psychological Science*, 13(2), 255–259. [doi: 10.1177/1745691617698146](https://doi.org/10.1177/1745691617698146)
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). *p*-Curve and Effect Size: Correcting for Publication Bias Using Only Significant Results. *Perspectives on Psychological Science*, 9(6), 666–681. [doi: 10.1177/1745691614553988](https://doi.org/10.1177/1745691614553988)
- Strack, F. (2016). Reflection on the Smiling Registered Replication Report. *Perspectives on Psychological Science*, 11(6), 929–930. [doi: 10.1177/1745691616674460](https://doi.org/10.1177/1745691616674460)
- Strack, F., Martin, L. L., & Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology*, 54(5), 768–777. [doi: 10.1037/0022-3514.54.5.768](https://doi.org/10.1037/0022-3514.54.5.768)
- Świątkowski, W., & Dompnier, B. (2017). Replicability Crisis in Social Psychology: Looking at the Past to Find New Pathways for the Future. *International Review of Social Psychology*, 30(1), 111–124. [doi: 10.5334/irsp.66](https://doi.org/10.5334/irsp.66)
- 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(3), 252–260. [doi: 10.1056/NEJMsa065779](https://doi.org/10.1056/NEJMsa065779)
- UK ECT Review Group. Efficacy and safety of electroconvulsive therapy in depressive disorders: a systematic review and meta- analysis. *The Lancet*, 361(9360), 799–808. [doi: 10.1016/S0140-6736\(03\)12705-5](https://doi.org/10.1016/S0140-6736(03)12705-5)
- Vazire, S. (2018, January). *The Credibility Revolution in Psychological Science*. Keynote presentation at National Institute for the Teaching of Psychology, St. Pete Beach, Florida, USA.
- Vohs, K. D. (2015). Money priming can change people's thoughts, feelings, motivations, and behaviors: An update on 10 years of experiments. *Journal of Experimental Psychology: General*, 144(4), e86–e93. [doi: 10.1037/xge0000091](https://doi.org/10.1037/xge0000091)
- Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R. B., ... Zwaan, R. A. (2016). Registered Replication Report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, 11(6), 917–928. [doi: 10.1177/1745691616674458](https://doi.org/10.1177/1745691616674458)
- Wallace, B. C., Dahabreh, I. J., Trikalinos, T. A, Lau, J., Trow, P., & Schmid, C. H. (2012). Closing the gap between methodologists and end-Users: R as a computational back-end. *Wiley Interdisciplinary Reviews Computational*, 49(5), 1–15.
- Walum, H., Waldman, I. D., & Young, L. J. (2016). Statistical and Methodological Considerations for the Interpretation of Intranasal Oxytocin Studies. *Biological Psychiatry*, 79(3), 251–257. [doi: 10.1016/j.biopsych.2015.06.016](https://doi.org/10.1016/j.biopsych.2015.06.016)

Table 1

Summary of Practices Related to Research Transparency

Term or practice	Description	Context, Significance, or Impact	Teaching notes for General Psychology
Questionable Practices			
Underreporting of null findings	When a researcher conducts multiple studies or includes multiple dependent measures in a single study, they sometimes report only the studies or dependent measures that showed the predicted effect. This means that null findings are never reported or are underreported.	This practice misleads the reader about the strength and consistency of the evidence. It also misleads by presenting exploratory research as confirmatory.	When teaching about the theory-data cycle, ask students, What should a scientist do if their study does not support the hypothesis? Should they hide that result or make it publicly available?
p-hacking	During data analysis a researcher may unintentionally exploit flexibility in data decisions, such as which outliers or covariates to include or which statistical tests to run. Decisions that led to a significant p value remain in the research report, while less successful analyses go unreported.	P-hacking inflates the Type I error rate, so p-hacked results are more likely to be flukes that cannot be replicated.	Have your students try the p-hacking demonstration available at fivethirtyeight.com : “hack your way to scientific glory.” Teach estimation thinking, rather than significance testing (see Table 2).
HARKing Hypothesizing after the results are known	After finding an unexpected result, researchers go back and revise the hypothesis, pretending they had predicted the result all along.	Theoretically unexpected results are potentially interesting, but they can mislead about the importance or magnitude of an effect. Therefore, unexpected results should be treated with caution until they are replicated. HARKing misleads readers about	You can teach HARKing as the opposite of the theory-data cycle. A hypothesis should be like placing an advance bet: you bet that the study will come out a particular way and scientists accept the

		the strength of the evidence.	results no matter how they come out.
New Practices			
Full transparency via open materials and open data	<p>Full transparency includes sharing every dependent variable and every statistical choice explored in the study.</p> <p>Open materials refers to the practice of sharing full experimental protocols and stimuli.</p> <p>Open data refers to the practice of sharing one's data files and codebooks so that other researchers can independently analyze the results.</p>	<p>Being transparent about all the steps taken in a particular study provides a more honest account of the strength of the evidence.</p> <p>Open materials allow other scientists to conduct direct replication studies.</p> <p>Open data allows other scientists to replicate a study's analyses independently. Open data also makes p-hacking more apparent. Open data allows people to explore other hypotheses in the data, ensuring that our participants' time and data has been fully used.</p>	<p>In an introductory class, discuss the benefits of open data for the general public, especially given that research is often supported by tax dollars.</p> <p>In upper-level statistics and methods courses, students can attempt to replicate published studies with open materials, or replicate published analyses with an open data set.</p>
Preregistration and registered reports	<p>Preregistration is when researchers publicly submit their study's procedures, specific hypotheses, and planned data analyses in advance. The plan is posted online with a time- and date-stamp before data collection begins.</p> <p>In a registered report, the preregistered plan is peer reviewed and the journal offers conditional acceptance of the resulting</p>	<p>Preregistration makes the research team's hypotheses and data analytic decisions public in advance. It helps counter HARKing and p-hacking.</p> <p>Registered reports explicitly shift publication criteria away from the strength of the result and toward the importance of the question and the quality of the methods.</p>	<p>Introduce preregistration during theory-data cycle. Students can practice a form of preregistration by making their own predictions during class. After hearing about a study's method, students can submit a prediction (perhaps using a personal response system) before the instructor displays</p>

	study, no matter how it turns out		the results.
--	--------------------------------------	--	--------------

Table 2*A brief comparison of the new statistics to null-hypothesis significance testing (NHST)*

Concept	Null Hypothesis Significance Testing (NHST) approach	Estimation (New Statistics) approach	Teaching notes for the estimation approach:
Research question	Binary and qualitative: Is there a difference?	Quantitative: How big is the difference?	Teach students to ask, “How much?” Supplement textbook coverage with effect size information.
Statistical reporting	Test statistics and p -value	Effect size with confidence interval	Teach students to ask, “How wrong?” Use simulations to show how point estimates and confidence intervals vary from sample to sample.
Conclusions in context	Reject the null or fail to reject. A single result seems definitive, need for replication not clear.	The confidence interval indicates the range of values still compatible with the data. Highlights uncertainty, fostering replication and meta-analysis.	Teach students to ask “What else is known?” Use forest plots to teach both quantitative reasoning and meta-analytic thinking.
Error Framework	False positives (Type I) and False negatives (Type II)	Sign error (Type S: Is the difference in the wrong direction?) and Magnitude Error (Type M: Were we wrong about the strength of the effect?).	

Figure 1*Concepts discussed in this article*

Note. Badge images in the center column were developed by the Association for Psychological Science (<https://osf.io/tvyxz/wiki/home/>).

Figure 2

Preregistration can be added as a step in the theory-data cycle

Note. Here, preregistration is added to an existing figure from an open introduction to psychology text.

Source of original image: <https://courses.lumenlearning.com/wmopen-psychology/chapter/outcome-the-scientific-method/>

Figure 3

Forest plot of a meta-analysis of the relationship between poor sleep quality and access to a nighttime media device

Studies

Arora et al., 2013
 Arora et al., 2014
 Chahal et al., 2013
 Gamble et al., 2014
 Gradisar et al., 2013
 Kubiszewski et al., 2013
 Lemola et al., 2015

Note. The effect size here is an odds ratio; it is the chance of having poor sleep *with* access to a device compared to those *without* access. Values bigger than 1 indicate devices are associated with an increased risk of poor sleep, numbers less than 1 indicate devices are associated with a decreased risk; a 1 represents no association. Each square represents the effect size from a single study (with larger squares indicating larger samples); each line represents uncertainty from that study (95% confidence interval). The studies with bigger sample sizes have less uncertainty and therefore shorter lines. The diamond represents the meta-analytic effect size and uncertainty, with the middle of the diamond indicating the observed change in risk and the width of the diamond indicates the remaining uncertainty (95% confidence interval). Overall, those with access to a device at bedtime were 2.17 times as likely to have poor sleep, but there is still considerable uncertainty about this risk (95% CI[1.42, 3.32]). This figure is replotted from a meta-analysis by Carter et al. (2016; all studies above fully cited there), using OpenMeta[Analyst] (Wallace et al., 2012)

Figure 4*Effect size “zoo” for depression treatments*

Note. This figure shows meta-analytic effects of different interventions for depression: Electroconvulsive therapy (ECT), exercise, cognitive-behavioral therapy (CBT), and the antidepressant Prozac (Fluoxetine). Each square shows the meta-analytic effect size relative to control/placebo treatment; the lines show 95% confidence intervals. The meta-analyses reported effect sizes in standardized units (Cohen's d). In addition, the HAM-D axis converts these values into expected improvements in the Hamilton Depression Rating Scale (conversion based on the meta-analysis by UK ECT Group, 2003 which translated from d to HAM-D scores using a control-group standard deviation of 10.68).