Chapter 1: Key Methodological Issues in Researching Gifted Education and Advanced Academics

Matthew T. McBee

In The Development of the High Ability Child: Psychological Perspectives on Giftedness
 Edited By Erin M. Miller, Michael S. Matthews, Dante D. Dixon

doi: 10.4324/9781003025443-1

Abstract

This chapter addresses a number of crucial questions for people who wish to learn about giftedness and gifted education. How do you find reliable research on gifted children, and who does it? How can you evaluate the truth-value of claims made about giftedness and gifted education? What type of incentives confront researchers, and how do those affect the research they do? This chapter discusses the state of educational and psychological research (within which giftedness research is embedded) focusing on the human elements that can reduce the reliability and trustworthiness of research results. The causes and consequences of low-quality research are discussed, as well as methods for improving the quality and credibility of research. The chapter closes with advice on how to "read between the lines" to identify weak research findings that are unlikely to hold up.

Winston Churchill famously observed that "democracy is the worst form of government, except for all the others" (Langworth, 2009). This chapter has a similar thesis: research is the worst way of knowing things, except for every other way. I will not offer platitudes such as "believe science" or provide inviolable rules to differentiate trustworthy science from science fiction (McBee & Field, 2017). Indeed, my position may seem excessively skeptical, or perhaps excessively bearish on the reliability of research and its ability to differentiate truth from falsity. This is so. But my critiques of research can be made simply because research exposes itself, by definition, to the possibility of being wrong. Every other way of knowing -- induction from one's own experiences, adherence to unsupported theories that 'feel true', adoption of sociopolitical worldviews, or embracing the opinions of experts -- is just as (or more) capable of going astray. The problem with these methods is that they fail silently. Research, though flawed, is superior to these methods precisely because it is repeatedly made to crash against reality.

This chapter is not a substitute for rigorous study of research methods, the history and philosophy of social science, or statistics. Instead, it is a survey of issues to consider when delving into the diverse body of research and pseudo-research on gifted students; it is a meditation on applied epistemology. What constitutes knowledge pertaining to gifted students? Who produces it, and how do I know if it is reliable? What is evidence, and what should be my standard for accepting it? There are no consistent rules that definitively can distinguish reliable from unreliable work (though there are clues that will be discussed later in this chapter). Instead,

offered here is background knowledge and a set of heuristics to assist in judging claims about giftedness and gifted children. As statistician George Box observed; “Since all models are wrong the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad” (Box, 1976, p. 792). Let us seek to avoid the “tigers” -- ideas and concepts which are damagingly wrong, leading to the adoption of harmful practices or the discarding of helpful ones. I am less concerned about the “mice.”

The Big Picture: Who Does Gifted Education Research, and Where Do I Find It?

For a claim about giftedness or gifted children to qualify for the hallowed status of **supported by research**, there must at some point be data -- the records of careful and systematic observation -- at the end of the chain of reasoning. Observations derived from other sources -- for example intuition, informal observation, expert opinion, tradition, folk psychology, or even reasonable-sounding armchair theorizing -- simply do not achieve the same status as those derived from data. These are potent fodder for hypothesis formation, but cannot themselves qualify as evidence. This is because human cognition is notoriously unreliable, and it is exceptionally prone to noticing patterns that actually do not exist (see Friedman, 2017 for review). We can almost always find *ex post facto* explanations for the phenomena we observe, and even worse: the patterns we notice are biased by our preconceived notions and self-interest. Absent the structure imposed by the *careful study* requirement phrase in the definition of research, it would be all too easy for multiple observers to disagree about even the observation itself, to say nothing of its interpretation, implication, or meaning.

One must also bear in mind the distinction between *primary sources* and *secondary sources* when interacting with the gifted education literature. Primary sources are those that reference data directly -- in the form of statistical summaries of quantitative data or the quotes, field notes, or other primary artifacts of qualitative data. As the *open science* movement (McBee et al., 2018) continues to gather steam in the social sciences, we can increasingly expect research articles to link to the raw data itself. In the United States, the vast majority of gifted education primary-source research is published in the field’s specialty journals, which include *Gifted Child Quarterly*, *Journal of Advanced Academics* (formerly called the *Journal for Secondary Gifted Education*), *Journal for the Education of the Gifted*, and *Roeper Review*. There are also two international gifted education journals, *Gifted and Talented International* and *High Ability Studies* (formerly known as the *European Journal of High Ability*). Most original research in gifted education is published in one of these outlets.

Secondary sources are those that do not describe data per se, but instead describe, synthesize, or translate primary research documents into a more orderly, readable, or comprehensive form. Such is the case with textbooks as well as many other books, periodicals, and websites. For example, *Teaching for High Potential* and *Parenting for High Potential* seek to translate research findings into a form that can be consumed by teachers and parents, respectively. Most people who interact with the research in our field do it through the secondary source literature. One reason is that the primary source literature is more technical. But more importantly, nearly all of the primary research outlets mentioned in the previous paragraph are **paywalled**. This means that most non-university readers have no reasonable access to primary research on gifted students. And although extra-legal means of free access to paywalled articles have appeared, not everyone

knows about them or is comfortable using them. One hopeful development on this front is the recent appearance of preprint servers such as <http://psyarxiv.com> and <https://edarxiv.org/> which host free pre-publication versions of published articles. Unfortunately, only a small proportion of new work is being uploaded there, and there is no prospect for liberating access to the back catalog of already-published work via this mechanism.

Perils of Reliance on Secondary Sources

It is perilous to rely on secondary sources alone. The description of primary research can degrade like in a “game of telephone” in which multiple layers of summarizing and translation render the original author’s description nearly unrecognizable. Primary literature always describes the details of the design, sample size, and other features, which can help readers determine how much weight to give to the conclusion. Further, every research article must include a *limitations* section disclosing the known flaws of the research and how they might jeopardize the conclusion. These often are the first details to disappear when summaries of primary research are written. The summary of the research sometimes reads like a set of claims and counterclaims that all have the same implicit weight (for example, X and Y found that gifted students had higher satisfaction in cluster grouping, while Z did not) even when these studies might vary in rigor. Perhaps X’s claim was based on a sample of 4,000 students from 5 states, while Z’s was a convenience sample of 30 students in his own classroom.

Another issue is that sometimes secondary sources do not lead to a primary source bearing data or evidence. Secondary sources can become a recursive train to nowhere in which the original source is based on theory or expert opinion with no actual data or evidence at the basis. My colleague Matthew Makel has coined the phrase “handbooks all the way down” to describe this lack of grounding. I want to emphasize this issue because it is so common and, in my opinion, destructive to the intellectual foundation of the field. Once a person injects an evidence-free claim into the research literature, it can easily receive the credibility of actual evidence. Subsequent citations of that claim in secondary sources may not adequately indicate the questionable basis. Over time, the claims may become accepted as factual and may reach the same status as those findings supported by actual research. A perfect case study for this phenomenon is Silverman’s (1997) paper entitled “The Construct of Asynchronous Development”, which was published in the *Peabody Journal of Education*. This is an important and influential paper in our field which advances an alternative conceptualization of giftedness. But it is critical to realize that this paper is a theoretical piece which does not present any empirical data in support of its arguments. Though intriguing and stimulating to read, it is not empirical research. Simply appearing in an academic journal or a book does not make something *research*. Theorizing is just as important as verification in the scientific process, but making a claim is not the same as supporting it. I will pick on myself as well -- a book I co-authored (*Beyond Gifted Education*; Peters et al., 2013) similarly is not research. The evidence that both of these pieces offer in terms of supporting their core arguments is logical, rhetorical, and in some cases based on anecdote -- all of which is thin gruel in comparison to verified data derived from structured observation.

A distressing number of concepts in gifted education have been laundered into research through this, or similar, mechanisms. One example is the *Bright Versus Gifted* checklist that one often

encounters in the field, which I shall not reproduce here out of fear of further spreading a misguided idea. For an example of the influence this conception retains in our field, see Taibbi's (2012) *Psychology Today* piece. Scott Peters (2017) traced the origin of this idea to a 1989 article in *Challenge* magazine, which is not a research journal. This idea -- that there is a meaningful distinction between the truly *gifted* child, as opposed to the mere *high achiever* -- and that our field should be interested in the former but not the latter, is commonly described as *research based* when it is not. So one of the most important ideas of this chapter is:

There is no need to be a nihilist, but a little skepticism goes a long way.

I recognize that this task is made far more difficult by the paywalls restricting access to journal articles for non university-based individuals. Tools such as Google Scholar can help trace references back to their sources. Further, there are resources, such as Plucker and Callahan's (2020) *Critical Issues in Gifted Education: What The Research Says*, which does a nice job of identifying the empirically-supported ideas in our field. But unfortunately, this is not a job that can be completely outsourced. We must all be responsible for our own conceptions and our duty to align them ever more closely to reality.

Who Does Gifted Education Research?

Most gifted education research is performed by university faculty, graduate students, and K-12 personnel with university affiliation (perhaps during a degree program). These researchers tend to be found in education, educational psychology, or psychology departments, though recently some economists have become interested in the field. Unfortunately, most of the research output produced by these individuals -- which is the highest quality and most reliable of all the content produced in this area of inquiry -- disappears behind the paywall where its influence and reach is greatly limited. Most of the *writing* on giftedness, as found all over the internet in blog posts, websites, social media groups, and other such sources, is performed by laypersons with seemingly much lower evidentiary standards. The sad fact is that only a small percentage of the concepts, ideas, or perspectives one encounters in lay writing on the subject have any scientific or research-based validity whatsoever. There is an exceptional amount of nonsense flourishing in the wild under the rubric of giftedness.

Why do university types publish their work in paywalled journals? To understand that requires consideration of the incentives they confront. Writing accessible pieces for non-university audiences, engaging in public discussion and debate about giftedness and gifted education, and other types of what might be called "outreach" or "science communication" not only are not rewarded; they are actively punished. A researcher who spent their research time on those things rather than preparing and submitting work to academic journals would fall behind in the competition for publication productivity and grant funding - and thus run the risk of losing their position. Even tenured professors, whose jobs are not in jeopardy, would risk losing professional status ("retiring in place"), competitiveness for other positions, or future success gaining grants if they chose to allocate their effort away from status-generating activities. There is almost no incentive to engage in activities that do not promote career development.

For historical and now largely anachronistic reasons, private publishing companies provide the structure for typesetting, publishing, and distributing research work. When publishing transitioned from the intrinsically limited, bound printing model to online distribution during the first decade of the 2000s, publishers erected paywalls to protect their revenue streams. From the perspective of a university-based researcher, nothing had changed. After all, one would have needed library access to read printed and bound research journals; restricting online articles to the same audience of people was no different. But the transition to online dissemination presented an opportunity to radically democratize access to scholarly research; an opportunity that was largely not taken.

Setting Expectations About Research on Giftedness

There are four critically important issues to bear in mind when one evaluates the literature on giftedness. They initially may seem somewhat disconnected, but together they provide a critical background for reading and understanding the literature.

The Replication Crisis

In the early 2010s, a confluence of events in the social and biomedical sciences collectively demonstrated that the usual methods of ‘doing research’ in these fields were producing a much higher than expected rate of false positive claims. This was demonstrated most clearly by a paper on parapsychology published by social psychologist Daryl Bem (2011) in the *Journal of Personality and Social Psychology*. It was only Bem’s status as an eminent and legitimate scholar that made this publication possible; Bem secured an arrangement with the journal’s editors to subject this submission to the same peer review process as other, more topical submissions would receive. The editors agreed to publish the paper if the peer reviewers were satisfied (Judd & Grawonski, 2011).

Bem’s paper was published. It reported the results of nine experiments, involving over 1,000 subjects, testing for the effect of time-reversed responses to ‘classic’ social-cognitive psychological phenomena. The participants’ responses were observed *prior* to the presentation of the stimuli and found to be significantly different, at or beyond the customary $p = .05$ statistical threshold, in eight of the nine experiments. The mean effect size - a measure of the magnitude, and thus the real-world importance, of the difference between the “treatment” and “control” subjects -- across all the experiments was 0.22 standard deviations. This paper created a sensation in psychology because it demonstrated that the usual standards for scientific and statistical work in the field could utterly fail. And if these methods could be used to support a claim that ‘the future causes the past’, how many other false ideas might have received similar scientific support?

The answer, as it turned out, was quite a lot. One might initially think that the publication of a false claim would not be a huge problem because replications of that study would inevitably reveal it to be in error. This is precisely what happened when Pons and Fleishman claimed to have achieved cold fusion in the 1980s; within weeks multiple teams of physicists reported being unable to replicate cold fusion, and the idea was quickly abandoned (Browne, 1989). But the

social and biomedical sciences do not have a culture of replication, which has historically been viewed as a waste of time and resources, or as low-status make-work for insufficiently innovative researchers. Makel et al. (2012) analyzed the psychology literature and discovered that only 1-2% of published work was a replication of a previous study. The same analysis on education research revealed even worse findings. Only 0.13% of the research published in the top 100 education journals was identified as a replication of previous work (Makel and Plucker, 2014). Thus, erroneous research conclusions in these fields have a good chance of remaining unchallenged and uncorrected.

When psychologists began to invest in replication studies in the mid 2010s, they made a series of alarming discoveries. Many social psychology findings, which had previously been understood as well-established, began to fall apart under replication. These included such phenomena as ego depletion (Hagger et al., 2016), the facial feedback hypothesis (Wagenmakers et al., 2016), and social priming (Doyen et al., 2012). Table 1.1 summarizes the results of some recent large-scale replication efforts in the social and biomedical sciences. Overall, only 35% of the replicated studies have successfully produced evidence in line with the original claims. An excellent timeline of the replication crisis through 2016 was compiled by Gelman (2016b). A discussion of the causes of the replication crisis, and some proposed reforms, are given by McBee et al. (2018), van der Zee and Reich (2019), and Makel et al. (2019).

Study	Field	Replicability
Prinz et al. (2011)	Biomedical	13/52 = 25%
Begley & Ellis (2012)	Biomedical	6/59 = 11%
RP: Psychology (2015)	Psychology	35/97 = 36%
Camerer et al. (2016)	Economics	11/18 = 61%
Camerer et al. (2018)	Social science	13/21 = 62%
Many Labs 1 (Klein et al., 2014)	Psychology	10/13 = 77%
Many Labs 2 (Klein et al., 2018)	Psychology	14/28 = 54%
Many Labs 3 (Ebersole et al., 2016)	Psychology	3/10 = 30%
Registered Replication Reports	Psychology	1/9 = 11%
RP: Cancer Biology (preliminary)	Biomedical	7/16 = 44%

Table 1.1. *Summary of Large-Scale Social and Biomedical Science Replication Efforts*

Note: **Overall replication rate: 35%**. RP stands for “Reproducibility Project.” PR: Psychology refers to Open Science Collaboration (2015). RP: Cancer Biology refers to Errington et al., (2014). Registered Replication Reports refers to the series of articles published in *Perspectives on Psychological Science* and includes Alogna et al. (2014), Hagger et al. (2016), Eerland et al. (2016), Cheung et al. (2016), Wagenmakers et al. (2016), Bouwmeester et al. (2017), O’Donnell et al. (2018), McCarthy et al. (2018), and Verschuere et al. (2018).

The observant reader will note the absence of any education replication projects in the table. We do not know what the replication rate of education research in general (or giftedness research specifically) might be, but since it relies on the same methods, norms, and incentives as general social science work, we might expect a similar degree of replicability. Educational studies are much harder to replicate than simple lab studies involving undergraduate volunteers because it is difficult to collect data in K-12 settings. Further, educational researchers have not yet faced a Bem-like situation, and are largely aware of the replication crisis as something that has occurred elsewhere rather than something truly threatening to the status quo. However, there are signs that education research will face the same issues. A replication of Steele and Aronson's (1995) stereotype threat effect, which purports to explain race and gender differences in academic performance, found no evidence that it accounted for disparities in math achievement by gender (Flore et al., 2019). A recent large-scale, high-profile replication of Dweck's *growth mindset* theory (Yeager et al, 2019) found that a mindset intervention had no effect for average or high-achieving students, but did increase the average achievement of low-achieving students by 0.1 GPA units. By comparison, websites such as MindsetWorks.com claim that growth mindset research started a revolution that "changed education forever." It would not be unreasonable to expect that only a third to two thirds of educational research is trustworthy.

Most Educational Interventions Have Modest Effects

The previous section attempted to bracket expectations about the trustworthiness of published research. Even if only a third of published educational research holds up, that is (in my estimation) at least ten to one hundred times better reliability than what one might find in a random parenting blog or gifted education Facebook group. And at least it exposes the details of the reasoning chain to scrutiny.

Another type of adjustment of expectations needs to occur in order for readers to establish appropriate expectations about the effect of educational interventions. That is - most interventions do not work. And the ones that do work tend to exhibit weak effect sizes. I highly recommend spending some time browsing the studies compiled by the Institute for Education Sciences *What Works Clearinghouse* (WWC, <https://ies.ed.gov/ncee/wwc/>) in the United States. The WWC reports results from a carefully curated set of education studies which meet certain standards of rigor (e.g., randomized experiment or quasi-experiment). The studies are organized by topic and presented in a standard format. The effect sizes for the studies are expressed on a common metric for ease of comparability; for many studies, this is an *improvement index* value, computed as the expected percentile improvement for treated versus control students. So, for example, an intervention with an improvement index of 10 would be expected to increase the average achievement of treated students by 10 percentile units relative to untreated students.

There are a few striking aspects about the WWC. First is that so few studies qualify. The WWC has reviewed over 10,741 studies, but only 204 (about 1.9%) of these were deemed rigorous enough for inclusion in the Clearinghouse. Many of the topic areas summarized do not include any studies at all. Second is that the effect sizes tend to be so small. Most positive effects reported by the WWC are small; the largest ones tend to be in the 8-12 percentile range. A little time with the WWC can help to establish reasonable expectations for the impact of educational interventions in general. Educational outcomes are resistant to improvement.

Weaker Studies Tend to Show Stronger Effects

Another pattern that is crucial to understand is that, when examining studies published in peer-reviewed academic journals, weaker studies tend to show larger and stronger effects than stronger studies do. One easy-to-measure aspect of a study's strength is its sample size (n). Figure 1.1 plots the sample size against the meta-analytic effect size, in this example for studies in three areas of biomedical science. The relationship is striking; small- n studies generally reported effect sizes in the 0.6 - 1.0 range, whereas large- n studies reported effect sizes in the 0.15 - 0.25 range.

Figure 1.1 The relationship between sample size and the reported effect size in biomedical science, as presented by Dumas-Mallet et al. (2017)

Note: This figure was presented as Figure 4a in the original publication. Published by the Royal Society under the terms of the Creative Commons Attribution License

<http://creativecommons.org/licenses/by/4.0/>

Perhaps you are thinking, “Wait a minute. Maybe scientists reserve large sample sizes to study phenomena that they expect to be weak or subtle.” This is absolutely a reasonable hypothesis, and there may be some truth to it -- but it cannot be the whole story of what is happening here. This pattern is so consistent as to almost be a physical law. It is observed even within the same phenomenon. For example, the growth mindset work discussed earlier showed strong effects on achievement in small studies (see Yeager & Dweck, 2012, for review). Because initial studies tend to be small, and follow-up studies larger, this pattern makes it appear that effects are weakening over time. In a piece for *New Yorker* magazine, Jonah Lehrer (2010) used the phrase “the truth wears off” to describe this pattern. But the truth isn’t wearing off. It is rather that a mirage disappears when you look at it more closely.

More on Incentives

Absent any filtering, one would expect there to be no relationship between estimated effect size and study rigor; at least not within a particular topic area. Small studies would produce more variable results, to be sure, but when averaging over these effect sizes (as in meta-analysis), size / rigor should be independent of effect size. In reality, there are multiple aspects of filtering that are linked to incentives operating at different levels of the system.

- There are at least two ways for a study to be interesting: rigor and findings. A large, rigorous study that finds nothing can be interesting, as can a small study that produces an exciting finding. But a small, null study is not very exciting.
- ‘Interesting’ articles get read and cited more than uninteresting articles.
- Academic journals print only a small proportion of the submitted articles, ranging from a ~50% acceptance rate for low-prestige journals to <5% for the most prestigious journals.
- A journal’s prestige is measured by the *impact factor*, which is calculated on the basis of citations. High-prestige journals will attract more interesting studies, leading to higher prestige. There is a positive feedback loop with selection pressure.
- Researchers need to amass publications in high-prestige journals.
- Researchers are rewarded much more thoroughly for making new discoveries than for running rigorous studies that find nothing. The former is a path to scientific fame and celebrity, with all its attendant rewards (plum jobs, grants, editorships, speaking and book contracts, the best graduate students, etc).

So journals need a steady supply of interesting (as defined above) research to publish. This is why it is, even now, quite difficult to publish replication studies in many fields (and why replications have been so historically underutilized). Null findings are largely uninteresting unless they result from extremely rigorous investigations with high prior probability. So null findings are rarely published (Ferguson & Heene, 2012). Large effect sizes are interesting. And they are much easier to generate in small studies rather than large ones.

Most quantitative research in gifted education is analyzed using frequentist statistical methods. These methods have built-in compensation for variability as a function of sample size. The standard for judging an estimate (e.g., a mean, a treatment effect, a correlation) as ‘statistically significant’ adjusts with the sample size such that a larger value is required for significance in small samples than in large ones. This standard is set such that only 5% of observed results would cross this threshold by chance when there is no effect. These *false positives* are known as Type-I errors in statistics. Figure 1.2 displays the results of a large set of simulated studies when there is no underlying effect. Each dot represents the absolute value of the effect size estimated from the study. The curved line shows the threshold for statistical significance. Only 5% of the dots lie above this threshold, even though the dots have much more vertical dispersion in small samples than in large ones.

Thus far, there is no reason to expect any relationship between sample size and effect size. Sure, smaller samples are more variable, therefore having a higher chance of yielding a large estimated

effect, but this is cancelled out perfectly by the higher threshold they must reach to trigger statistical significance. The problem occurs when a selection mechanism is applied to the results.

Figure 1.2 Relationship between sample size and the effect size when the true effect is zero

Note: Based on 5,000 simulated studies generated under a condition of no true effect. The results vary due to statistical noise. Each dot represents one study's result. The line displays the threshold absolute effect size for statistical significance; points above the line would be deemed significant at or beyond the .05 level. The figure illustrates how (a) results from small studies have a much larger 'spread' of effect size estimates, and (b) the small-sample statistically significant studies would have much larger effect sizes than the large-sample significant studies.

Journals are incentivized to publish interesting results. Results that show no effect are generally not considered very interesting (unless they are surprising). Therefore, journals tend to be filled with statistically significant results (van Zwet & Cator, 2020). Look again at Figure 1.2, but imagine only seeing the dots above the curved significance threshold line. The significant results from small samples have much larger effect sizes than those from large samples. When you examine only significant results, you will induce a negative relationship between sample size and effect size.

Gelman and Carlin (2014) coined the term *Type-M error* to refer to this inflation of effect sizes in studies that have been selected for significance (the "M" standard for "magnitude"). The smaller the study, the larger the Type-M error is likely to be. This implies that the published literature overestimates effect sizes, sometimes severely. Unfortunately, this overestimation also contaminates summaries of the literature. For example, meta-analysis is a technique for

computing the average weighted effect size over a set of studies. A meta-analysis of a set of inflated effect sizes yields an inflated meta-analytic estimate. There is what is known as a *garbage-in, garbage-out* problem. When large-scale multi-lab replication results have been compared to meta-analytic estimates, the replication effect size is typically one third of its meta-analytic counterpart (Kvarven et al., 2020).

Selection Effects Bias the Literature

The selection of statistically significant articles for publication is often called the *file drawer* problem (Rosenthal, 1979). It is important to note that the selection process happens at multiple levels, not just at the level of journal publication. In fact, selection bias operates at every level of the system. For an individual researcher operating within a competitive academic culture, a non-publishable study represents an enormous waste of effort. Studies with null results may not be completed if an interim analysis shows little hope of positive findings. If data collection is complete, the research paper may never be written. (After all, writing the paper takes time and effort that would better be spent on a new study). In both cases, the record of the study ever having been attempted disappears. This pressure also affects the way researchers allocate their time and resources. If a researcher has the resources to collect data from 1,000 students in a year, they could choose to do one study of $n = 1,000$ or twenty studies of $n = 50$. Running a large set of small studies maximizes the probability of producing a few publishable studies. If the one large study fails, the researcher (and his or her students) will have no research output for the year, imperiling their ability to continue doing research. The incentives strongly contradict the notion of putting ‘all the eggs in one basket’, even though scientific progress would be greatly improved by running fewer studies with higher quality.

Worse yet, researchers may apply selection methods *within* the study itself to attempt to achieve significance. Such efforts are known as *questionable research practices* (QRPs; Simmons et al., 2011) or *p-hacking*, and include practices such as collecting additional data to try to push a marginal result over the line or trying out different statistical analysis methods until a significant result is obtained. The influence of these methods is much stronger on small studies, because they have a high degree of variability in their results due to their inherent instability. Because statistical significance and effect size are linked, any p-hacking motivated to achieve statistical significance also inflates effect sizes. This is a serious problem that is worthy of deep consideration for anyone who wishes to read or (especially) do research. Beyond these within-study selection methods there is even yet another type of post-publication selection bias -- studies making positive claims are cited at much higher rates (Murray et al., 2019; de Vries, 2019).

The end result of all of these complementary biases is a poor replicability rate for the published literature, such as the 35% overall value reported in Table 1. When interested readers turn to the primary literature to assess, for example, the evidence base for growth mindset interventions, they will find study after study reporting positive effects. What they do not see are the studies that might have been done that did not find benefits for these interventions. If I show you five studies that all support mindset interventions, you are likely to form the impression that research supports their effectiveness. If I show you the same five positive studies, along with twenty-five studies with null results, you would likely reach a different conclusion. Your overall estimate of

both the probability of existence as well as the magnitude of the effect should hinge on how many null or near-zero results are in the set. But this is precisely the information that is hidden when each stakeholder in the academic research system follows its own private incentives.

What Should We Do?

In summary, there are multiple types of selection effects that make it so research results are not readily interpretable at face value. These effects not only bias the primary literature toward false positives, but also inflate the reported effect sizes for true positive effects.

1. Researchers who allocate effort to many small studies instead of fewer large studies.
2. Researchers who select for positive results within a study (e.g., via QRPs or p-hacking).
3. Researchers who fail to write up or submit null results for publication.
4. Journals that reject papers reporting null results.

It might appear that researchers bear nearly all the blame for this tragic state of affairs, but it is crucial to understand that the underlying driving factor actually is market demand. Consumers of research (the market) conflate entertainment value with scientific worth. The artificial scarcity of publishing capacity, which is necessary to create a status hierarchy among researchers, means that only the “best” articles are selected for publication. And since readers generally find “something” to be more interesting than “nothing”, journals that publish affirmative content are rewarded with increased subscriptions, downloads, and citations. This pressure is passed down to researchers, who create the interesting content; if they don’t create it, they are replaced by other researchers who will. The missing null literature, which renders those interesting, positive findings interpretable, is a tragedy of the commons. Until recently, its harm has been too diffuse and abstract to elicit a strong reform response from any part of the system. Even now, reformers appear to be only a small (but vocal) minority of scholars.

Taken together, the issues discussed in this section paint a somewhat dismal picture of the status quo of education research. Obviously, findings from primary research filter down into secondary sources, such as textbooks, so these latter sources cannot be any more reliable than their foundational material. Even though academic journal articles have gone through a gatekeeping process involving editors and peer reviewers, outsourcing one’s skeptical discernment to any third party is still a bad idea.

I believe that one can derive insight, knowledge, and wisdom from published academic research, by treating it carefully. The first and most important consideration is simply understanding that the literature is strongly biased toward affirmative claims; that some uncomfortably large proportion of what is published likely would not replicate, and that effect size estimates are inflated. With this knowledge, one can apply some basic heuristics to try to ‘de-bias’ the research base. The easiest one is to mentally divide the claimed effect size by three. Would this effect still be interesting if it were one-third as large? As discussed, many true effect sizes are zero, and most of the others are modest. Educational outcomes are strongly resistant to change. So when reading about some type of quick, minimal intervention that claims to have a strong effect on student achievement, weigh the relative likelihood of two possible explanations in your mind:

1. The actual effect is either zero or quite modest, but the results reported in this study are inflated.
2. This is actually one of the few, the proud, the strong educational effects.

Bear in mind that strong effects usually require a large, meaningful intervention that differs substantially from the status quo and that is delivered over a sustained period of time with high dosage. For example, the Knowledge is Power Program (KIPP, Tuttle et al., 2015) dramatically increases the time students spend in school over multiple academic years. According to the WWC, its effect on language arts, math, science, and social studies are 8, 12, 11, and 5 improvement index points, respectively.

Identifying Credible Research

Simine Vazire, a personality psychologist and leader in the open science movement, has asked, “do we want to be incredible or credible?” (Vazire, 2020). This framing cleverly captures both the appeal and the harm caused by the many biasing factors discussed in the previous section. The literature is *incredible* in the sense that it presents surprising finding after surprising finding that all appear to have the potential to revolutionize education. For example, growth mindset researchers suggested that just a subtle change from praising effort instead of ability can cause students to develop a growth mindset and thereby increase their achievement (Yeager and Dweck, 2012). But it is also *incredible* in the sense that it is *not credible*. When investigated rigorously, growth mindset interventions have no benefits for most students (Yeager et al. 2019). And if *credibility* is the opposite of untrustworthiness, it is also the opposite of the whiz-bang, one-weird-trick, low-effort/high-reward promises that have seduced so many educators and educational researchers. Credible research usually isn’t very surprising, but it can be counted on. One can implement credible interventions and expect them to work. And one can build on credible research to identify new principles and applications. Addressing the critique that credible science is boring, Andrew Gelman (2016a) wrote:

Daryl Bem claimed that Cornell students had ESP abilities. *If true*, this would indeed be interesting, given that it would cause us to overturn so much of what we thought we understood about the world. On the other hand, **if false, it’s pretty damn boring**, just one more case of a foolish person believing something he wants to believe. (no page, emphasis original)

Research isn’t supposed to be exciting. It’s just supposed to be true. And Gelman is right; it is hard to see much value in counterintuitive, surprising, clever, but ultimately incorrect research findings. Science fiction is an enjoyable genre, but it works better with characters and a plot.

Where Does Credibility Come From?

I think of research credibility as arising from three fundamental sources: technical rigor, procedural rigor, and transparency. Technical rigor comprises those elements that are typically taught in research methods courses, for example the design of the study, the constraining or

elimination of threats to validity, blinding procedures, ethics, measurement, sampling process and sample size, and data analysis. Technical rigor is often the only area in which people receive any training. Textbooks often treat it as if it is the only type of rigor needed. In doing so, they ignore the human element of the research process.

Procedural rigor involves creating and closely following a workable plan for the research. It involves data management, version control of data and analysis files, training of research staff, and coding of data. It also involves clearly thinking through and preemptively solving problems that may arise during the study -- or designing the study such that those problems do not occur. Procedural rigor is about constraining the researcher's "degrees of freedom" - the ability to make mid-study alterations in the research plan. These must be limited because once the data start being observed, the researchers' decision-making process becomes suspect. There is a definite chance that the researcher will either consciously or unconsciously make decisions to enhance the probability that the study produces a desired result. The highest level of procedural rigor occurs when researchers produce a detailed preregistration of the study and are able to follow it without deviation (McBee & Field, 2017).

The final aspect of credibility is transparency. Transparency means *showing your work*, allowing for every aspect of the inferential chain to be exposed for verification and criticism. This has been a core scientific value for hundreds of years. The motto of the Royal Society, the world's oldest scientific society, is *Nullius in Verba*, which translates to "take no one's word." Lupia and Elman (2014) wrote:

The credibility of scientific claims comes, in part, from the fact that their meaning is, at a minimum, available for other scholars to rigorously evaluate [...] Such open access to the origins of others' claims is the hallmark of scientific ways of knowing. Accordingly, when social scientists fail to document their assumptions, decisions, and actions [...] it limits others' ability to understand the meaning of scientists' claims. (p. 20)

Transparency of data, study materials, and data analysis code are three hallmarks of the open science movement (McBee et al, 2018). But the transparency also encompasses disclosing the true process used in the research, warts and all. For example, imagine how differently you might interpret the conclusions from these two processes:

1. We planned a sample size of 150 students. When data collection was complete, we performed a blinded, pre-planned analysis of the data. Our complete results are presented in the paper. We did not examine any other analysis other than that which was reported.
2. We initially planned a sample size of 70 students. When this was finished, the result was close to significant but not quite there. We began collecting more data and reanalyzing the data in increments of 20 students. Once we hit a sample of 150, the result for our main hypothesis was no longer near significance. However, we noticed another significant effect in the data analysis that we found exciting. We revised the framing of the paper around this finding.

The first situation is fully confirmatory, while the second is completely exploratory. Further, the main finding in the second paper resulted from HARKing (“hypothesizing after results are known”; Kerr, 1998). In the first situation, the confirmatory nature of the data analysis allows the false positive rate of the statistical analysis to be constrained to some tolerably low value. In the second situation, conversely, the false positive risk cannot be controlled (De Groot, 1969). In the long run, it is 100%. That is, given enough unconstrained freedom, one can *always* achieve statistical significance.

Credibility depends on these three aspects, all of which are necessary but not sufficient conditions. The problem is, typically, only technical rigor can be ascertained from reading a study. This is also the most difficult aspect to judge, as it can require a great deal of expertise to determine whether a particular statistical analysis was properly done or whether the design adequately protects against certain types of confounding. The procedural aspect of rigor is almost impossible to ascertain from a direct reading of a journal article, unless it is a preregistered study -- in which case, the publication can be compared to the preregistration document and reviewed for consistency. The best one can do is to amass circumstantial evidence by carefully ‘reading between the lines’ while looking for certain signs that might betray the use of questionable research practices. Transparency can be greatly improved by open data, open code, and open materials. But procedural transparency is essentially totally absent in standard, non-preregistered research articles. Readers have no choice but to take the authors at their word. So much for *nullius in verba*!

There is even a fourth aspect of credibility, outside of the internal validity structure of the research itself, that concerns the inference from the study’s results to the status of the theoretical propositions being evaluated. Studies provide the strongest evidence when they evaluate *severe tests* of theoretical statements (Mayo, 2018a). Mayo (2018b) wrote that “A claim is severely tested to the extent it has been subjected to and passes a test that probably would have found flaws, were they present.” (no page). A severe test is one in which an observation is evaluated that would only be expected to occur under the theory being evaluated, but not other theories. Affirmative study results provide corroborating evidence of theory only to the extent that they are severe tests. This idea connects directly to Karl Popper’s notion that a scientific theory is distinguishable from pseudo-science to the extent that it makes *risky predictions* -- predictions that are capable of being contradicted by evidence.

Strategies for Learning to ‘Read Between the Lines’ to Evaluate Credibility.

Learning to evaluate the technical rigor of a study requires a substantial understanding of research methods, measurement, and statistics. These skills are well-covered in many other publications and so I will not discuss them here. Instead, I will try to give some insight into some things I look for when evaluating research that has not been preregistered. Altmejd et al. (2019) showed that machine learning algorithms can be trained to identify the replicability of research articles with 70% accuracy. Dreber et al. (2015) found a similar 71% rate using a betting market. Clearly, the signs are there if we know what to look for. These methods are definitely not foolproof. They may fail to detect bogus articles (false negatives) and erroneously flag other trustworthy articles as likely untrue (false positives). But I believe that they likely result in much

more highly-calibrated judgments than simply accepting every published study's claims at face value.

1. *Does the research question or hypothesis seem like something that was formulated before or after seeing the data?* In his seminal article on HARKing, Kerr (1998) wrote about the “too convenient qualifier” that often distinguishes hypotheses that were formulated after the data were examined. This can often manifest as a hypothesized effect for one subgroup but not another, or an interaction between variables but not an overall effect (e.g., a main effect). Is it truly reasonable to expect that the effect would manifest for only a specific subset of the participants? Is it believable that the researchers would have been able to identify this specific subgroup in advance of seeing the data? One way to tell is by the sample size -- detecting subgroup effects or interactions takes much larger samples than detecting main effects. A possible example of this type of HARKing is from Bem's (2011) precognition study. In the first experiment, there were three types of stimuli presented to subjects. The ‘psi effect’ was observed only for one of the three classes of stimuli.

2. *Do the researchers seem like they would have taken “no” for an answer?* Given that researcher degrees of freedom are ubiquitous in research that is not preregistered, readers need to understand how much they were used. One ‘tell’ I look for is *spin*. This is most readily observed in the discussion section of the article. If the authors engage in a discussion of how their theory is probably true (or of how the intervention they evaluated is probably effective) in the face of underwhelming findings, I discount my appraisal of the likely truth-value of the findings they reported. I am much more impressed by a straightforward appraisal of evidence (e.g., “we thought this would work, but it didn't”) rather than “it didn't work so well this time, but here's why we still believe in it.” In other words, if the authors don't believe their own negative results, why should I believe their positive results? This is also why I can find qualitative research using grounded theory to be persuasive, but I am not very moved by qualitative work performed under theoretical perspectives that seem founded on strong worldview commitments.

3. *Is there evidence of outcome switching?* Where the presence of ‘spin’ gives you information about the researchers' potential preference for a certain set of results, assessing researcher degrees of freedom gives you information about their opportunity to bias the conclusion (either consciously or unconsciously) toward some preferred result. One manifestation of these degrees of freedom occurs when researchers have a large set of outcome variables available for reporting, and selectively report the ones that “worked.” For example, if the researchers are assessing academic achievement, how many different measures were available to them, and how many of these did they report? Do they report results only for the state mathematics achievement test, but not other subjects or other tests? Other times, the researchers may report a primary outcome on something other than the logical variable one would expect to assess given the research question. For example, the paper may report a significant increase in student or teacher satisfaction, but not report any achievement outcomes. You might ask yourself, “if I was studying this intervention, would I really care more about satisfaction than achievement?” If the answer is not ‘yes’, you might correctly surmise that the achievement outcome is not being reported because it was a null effect. And further, the satisfaction outcome being reported may well be a secondary

outcome that would not even have been examined if the achievement outcome had been successfully impacted by the intervention.

4. *Does the finding have high prior probability?* Ioannidis (2005) showed that not all statistically significant p -values have the same interpretation - even those with exactly the same value. The most important determinant of whether a significant p -value is likely to correspond with a true effect is the prior probability of that effect. Is the effect plausible? Do you find it surprising? Does the effect have a clear mechanism of action? Is the magnitude of its benefit commensurate with the intensity and dosage of its delivery? I am much more inclined to believe that grade-skipping (a full-time intense academic intervention whose dosage amounts to over a thousand hours per year) produces large effects on achievement than I am to believe the same for a one-hour per month pull-out enrichment program. This is also why I think so-called “nudge” interventions are completely bogus. I am not saying that we should discount any scientific finding that clashes with our preexisting beliefs. I am merely stating that extraordinary claims require extraordinary evidence.

5. *Are the numbers internally consistent?* The numbers reported in quantitative research papers follow certain regularities. For example, the sample sizes for subgroups should sum to the total sample size. The weighted mean of subgroup means should equal the overall mean. The test statistic equals the point estimate divided by its degrees of freedom. The p -value is directly calculable from the test statistic and degrees of freedom. Nuijten et al. (2016) analyzed 250,000 psychology papers and found that 50% contain at least one reported p -value that is inconsistent with its test statistic and degrees of freedom, with 12% containing a gross error. Nick Brown and colleagues (2017, 2019, 2020) have uncovered numerous examples of scientific misconduct (or fraud) by examining reported numbers for internal consistency.

6. *Do the authors have a conflict of interest?* One might expect the originator of a theory or model under study to have a higher stake in how research results turn out than a third party. After all, their reputation and professional status is at stake. And indeed, ‘first party replications’ (those that involve the original progenitors of an idea) succeed at higher rates (Makel et al., 2012). I am much more apt to believe claims after they are replicated by a third party.

7. *Are the key p -values barely significant?* One telltale sign of p -hacking is a pattern of barely significant p -values. When effects are real, p -values should follow a right-skewed distribution. There should be many more small values than large ones (Simonsohn et al, 2014). This even applies to the ‘significant’ range of values between 0 and .05. For example, there should be more p -values in the [0, .025] range than in the [.025, .05] range. There should be more p -values in [0, .01] than in [.04, .05]. A natural pattern of p -values might look like: .001, .005, .01, .0004, .034. An unnatural pattern of p -values might look like: .043, .021, .035, .047, .039. Once you understand this, you might be amazed how often you see this pattern in classic social psychology papers - most of which have not replicated.

8. *Are the results too good to be true?* If the results seem too good to be true, they probably aren’t true. Social science research is inherently noisy. The constructs are often vaguely defined. Manipulations or interventions may not be delivered with perfect fidelity. Measurement is

imperfect. And humans are incredibly variable. We should expect this messiness to manifest itself in research results. Schimmack (2012) described how excess success diminishes the credibility of research papers. For example, if a study reports six experiments, where each one is powered at 60%, the probability that all six would succeed *even if their underlying hypothesis is true* is less than 5%. It is far more likely that some of these experiments would fail by chance. So, ironically we should find 4 out of 6 or 5 out of 6 successful experiments more convincing than 6 out of 6. The latter is so unlikely that it should make you suspect that it is the result of some chicanery. Having reason to believe that researchers would not hide their failures should bolster your confidence when they show you their successes. This also applies to effect sizes that are too large to be believable. Often the blame does not lie with researchers themselves. Peer reviewers and editors are notorious for demanding that non-successful studies or hypothesis tests be removed. This is scientific malpractice that is typically justified for reasons of “space” (even though publication is online and electrons are free) or of “telling a clear story.” But as we have discussed, you cannot interpret successes without seeing the failures.

9. *Do we see the deadly “small n , big effect size” combination?* Large effect sizes may be inherently suspicious given everything that has been discussed thus far. I am far more willing to entertain the possibility when they come from large, rigorous studies than from small ones. For example, if the Yeager et al. (2019) pre-registered national experiment on mindset had found a moderate or large effect size for the intervention, I would have been willing to update my beliefs. When an $n = 60$ study finds a large effect, it does not move my belief in the slightest.

Conclusion and Coda

As I write this chapter, a movement toward abandoning testing on the grounds of equity is ascendent in the United States. Graduate programs are abandoning the GRE; colleges are making the SAT optional. Racial and socioeconomic inequities are easily observable in test scores. The rationale for eliminating them is to increase the fairness of admissions. But I fear that this effort will backfire. In the absence of tests, universities will become more reliant on letters of recommendation, personal statements, research experience, and prior academic performance. It seems self-evident to me that disparities also exist in all of these things, and that they will probably be much more severe than those seen in test scores. After all, which students will be able to get professional help crafting personal statements, intern in well-known laboratories, secure recommendation letters from famous academics, or attend elite universities? Tests make disparate outcomes visible, and for that they are being shunned in favor of admissions criteria whose biases are much harder to evaluate.

In the same manner, research is the worst way of knowing (except for all the others). Research is far from perfect, but at least we can see and evaluate its flaws. Some research is misleading, but some of it is trustworthy. And some flawed research is right in spite of itself. While I am bearish on the present state of social research and education research, I am optimistic about its future. Science is self-correcting. It may not happen quickly enough for our taste, but it is happening. The replication crisis has blossomed into a credibility revolution (Vizire, 2020). Someday soon, the revolution will begin in earnest in educational research; the new generation of researchers demands it (see van der Zee and Reich, 2018; McBee et al, 2018; Makel et al, 2019). I am certain that the quality of research in 2030 will be better than it is now (2020), and in 2040 better still. I

cannot say the same about the opinions of self-appointed gifted consultants, memes in Facebook groups, or opinions espoused in parenting blogs. What will never, ever change is the responsibility of each consumer of research to critically evaluate claims made about reality, and the necessity of adopting a high standard of evidence before accepting those claims.

References

- Alogna, V. K., Attaya, M. K., Aucoin, P., Bahník, Š., Birch, S., Birt, A. R., ... Brown, C. (2014). Registered Replication Report: Schooler & Engstler-Schooler (1990). *Perspectives on Psychological Science*, 9(5), 556–578. <http://dx.doi.org/10.1177/1745691614545653>
- Altmejd, A., Dreber, A., Forsell, E., Huber, J., Imai, T., Johannesson, M., et al. (2019) Predicting the replicability of social science lab experiments. *PLoS ONE* 14(12): e0225826. <http://dx.doi.org/10.1371/journal.pone.0225826>
- Begley, C. G., & Ellis, L. M. (2012). Raise standards for preclinical cancer research. *Nature*, 483(7391), 531–533. <http://dx.doi.org/10.1038/483531a>
- 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. <http://dx.doi.org/10.1037/a0021524>
- Bouwmeester, S., Verkoeijen, P. P. J. L., Aczel, B., Barbosa, F., Bègue, L., Brañas-Garza, P., ... Espín, A. M. (2017). Registered Replication Report: Rand, Greene, and Nowak (2012). *Perspectives on Psychological Science*, 12(3), 527–542. <http://dx.doi.org/10.1177/1745691617693624>
- Box, G. E. P. (1976). Science and statistics., *Journal of the American Statistical Association*, 71(356), 791–799, <http://dx.doi.org/10.1080/01621459.1976.10480949>.
- Brown, N. (2017, March 22). Strange patterns in some results from the Food and Brand Lab. [blog post]. <https://steamtraen.blogspot.com/2017/03/strange-patterns-in-some-results-from.html>
- Brown, N. (2019, February 19). Just another week in real-world science [blog post]. <https://steamtraen.blogspot.com/2019/02/just-another-week-in-real-world-science.html>
- Brown, N. (2020 April 21). Some issues in recent gaming research [blog post]. <https://steamtraen.blogspot.com/2020/04/some-issues-in-recent-gaming-research.html>
- Browne, M. (1989, May 3). Physicists Debunk Claim Of a New Kind of Fusion. *The New York Times*. Retrieved December 12, 2020
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T.-H., Huber, J., Johannesson, M., ... Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280), 1433–1436. <http://dx.doi.org/10.1126/science.aaf0918>
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., ... Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*. <http://dx.doi.org/10.1038/s41562-018-0399-z>

- Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1), 67–90. [http://dx.doi.org/10.1016/0149-7189\(79\)90048-X](http://dx.doi.org/10.1016/0149-7189(79)90048-X)
- Cheung, I., Campbell, L., LeBel, E. P., Ackerman, R. A., Aykutoğlu, B., Bahník, Š., ... Yong, J. C. (2016). Registered Replication Report: Study 1 From Finkel, Rusbult, Kumashiro, & Hannon (2002). *Perspectives on Psychological Science*, 11(5), 750–764. <http://dx.doi.org/10.1177/1745691616664694>
- De Groot, A. D. (1969). *Methodology: Foundations of inference and research in the behavioral sciences*. Mouton.
- Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., Nosek, B. A., & Johannesson, M. (2015). *Proceedings of the National Academy of Sciences*, 112(50) 15343-15347; <http://dx.doi.org/10.1073/pnas.1516179112>
- Dumas-Mallet E., Button K. S., Boraud T., Gonon F., Munafò M. R. (2017). Low statistical power in biomedical science: a review of three human research domains. *Royal Society Open Science*, 4, 160254. <http://dx.doi.org/10.1098/rsos.160254>
- Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., . . . Brown, E. R. (2016). Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67, 68–82. <http://dx.doi.org/10.1016/j.jesp.2015.10.012>
- Eerland, A., Sherrill, A. M., Magliano, J. P., Zwaan, R. A., Arnal, J. D., Aucoin, P., ... Prenoveau, J. M. (2016). Registered Replication Report: Hart & Albarracín (2011). *Perspectives on Psychological Science*, 11(1), 158–171. <http://dx.doi.org/10.1177/1745691615605826>
- Errington, T. M., Iorns, E., Gunn, W., Tan, F. E., Lomax, J., & Nosek, B. A. (2014). Science forum: An open investigation of the reproducibility of cancer biology research. *E-life*, 3, e04333. <http://dx.doi.org/10.7554/eLife.04333>
- Federal Aviation Administration. (2019). *Pilot's handbook of aeronautical knowledge*. [Washington, D.C.]: U.S. Dept. of Transportation, Federal Aviation Administration.
- Ferguson, C. J., & Heene, M. (2012). A Vast Graveyard of Undead Theories: Publication Bias and Psychological Science's Aversion to the Null. *Perspectives on Psychological Science*, 7(6), 555–561. <http://dx.doi.org/10.1177/1745691612459059>
- Flore, P. C., Mulder, J., & Wicherts, J. M. (2019). The influence of gender stereotype threat on mathematics test scores of Dutch high school students: a registered report. *Comprehensive Results in Social Psychology*, 1–35. <http://dx.doi.org/10.1080/23743603.2018.1559647>
- Friedman, H. H. (2017), Cognitive Biases that Interfere with Critical Thinking and Scientific Reasoning: A Course Module. Available at SSRN: <https://ssrn.com/abstract=2958800> or <http://dx.doi.org/10.2139/ssrn.2958800>
- Gelman, A. (2016a, 23 June). It comes down to reality and it's fine with me cause I've let it slide. [blog post]. Retrieved from <https://statmodeling.stat.columbia.edu/2016/06/23/it-comes-down-to-reality-and-its-fine-with-me-cause-ive-let-it-slide/>

- Gelman, A. (2016b, 21 September). What has happened down here is that the winds have changed. [blog post]. Retrieved from <https://statmodeling.stat.columbia.edu/2016/09/21/what-has-happened-down-here-is-the-winds-have-changed/>
- Gelman, A., Carlin, J. (2014). Beyond power calculations: Assessing type-S (sign) and type-M (magnitude) errors. *Perspectives on Psychological Science*, 9, 641–651. <http://dx.doi.org/10.1177/1745691614551642>
- Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., ... Bruyneel, S. (2016). A Multilab Preregistered Replication of the Ego-Depletion Effect. *Perspectives on Psychological Science*, 11(4), 546–573. <http://dx.doi.org/10.1177/1745691616652873>
- Hitchens, C. (2007). *God is not great: How religion poisons everything*. McClelland & Stewart.
- Ioannidis, J. (2005). Why Most Published Research Findings Are False. *PLoS Med*, 2(8): e124. <http://dx.doi.org/10.1371/journal.pmed.0020124>
- Judd, C. M. & Gawronski, B. (2011). Editorial comment. *Journal of Personality and Social Psychology*, 100(3), 406. <http://dx.doi.org/10.1037/0022789>
- Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*. 2(3): 196–217. http://dx.doi.org/10.1207/s15327957pspr0203_4
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Bahník, Š., Bernstein, M. J., . . . Nosek, B. A. (2014a). Investigating variation in replicability: A “Many Labs” replication project. *Social Psychology*, 45, 142–152. <http://dx.doi.org/10.1027/1864-9335/a000178>
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., ... Bahník, Š. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443–490. <http://dx.doi.org/10.1177/2515245918810225>
- Kvarven, A., Strømland, E. & Johannesson, M. (2020). Comparing meta-analyses and preregistered multiple-laboratory replication projects. *Nature Human Behavior*, 4, 423–434. <http://dx.doi.org/10.1038/s41562-019-0787-z>
- Langworth, R. (2009, June 26). Democracy is the worst form of government... [blog post]. Retrieved from <https://richardlangworth.com/worst-form-of-government>
- Lehrer, J. (2010, December 6). The truth wears off. *New Yorker*. Retrieved from <https://www.newyorker.com/magazine/2010/12/13/the-truth-wears-off>
- Lupia, A., & Elman, C. (2014). Openness in political science: Data access and research transparency. *Political Science & Politics*, 47, 19–42. <http://dx.doi.org/10.1017/s1049096513001716>
- McBee, M., Makel, M. C., Peters, S. J., & Matthews, M. S. (2018). A Call for Open Science in Giftedness Research. *Gifted Child Quarterly*, 62(4), 374–388. <http://dx.doi.org/10.1177/0016986218784178>. Preprint: <https://psyarxiv.com/nhuv3/>
- McBee, M., Field, S. (2017). Confirmatory study design, data analysis, and results that matter. In Makel, M. C., Plucker, J. A. (Eds.), *Toward a more perfect psychology: Improving trust,*

- accuracy, and transparency in research* (pp. 59-78). Washington, DC: American Psychological Association.
- Makel, M. C., Smith, K. N., McBee, M., Peters, S. J., & Miller, E. M. (2019). A Path to Greater Credibility: Large-Scale Collaborative Education Research. *AERA Open*.
<http://dx.doi.org/10.1177/2332858419891963>
- Mayo, D. (2018a). *Statistical Inference as Severe Testing: How to Get Beyond the Statistics Wars*. Cambridge University Press
- Mayo, D. (2018b). The meaning of my title ‘Statistical Inference as Severe Testing: How to Get Beyond the Statistics Wars’ [blog post]. Retrieved from
<https://errorstatistics.com/2018/05/19/the-meaning-of-my-title-statistical-inference-as-severe-testing-how-to-get-beyond-the-statistics-wars/#:~:text=That's%20what%20it%20means%20to,foundwe%20flaws%2C%20were%20they%20present.&text=The%20probability%20that%20a%20method,data%20is%20an%20error%20probability>.
- Nuijten, M. B., Hartgerink, C. H. J., van Assen, M. A. L. M., Epskamp, S., & Wicherts, J. M. (2016). The prevalence of statistical reporting errors in psychology (1985-2013). *Behavior Research Methods*, 48(4), 1205-1226. <http://dx.doi.org/10.3758/s13428-015-0664-2>
- McCarthy, R. J., Skowronski, J. J., Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., ... Bakos, B. E. (2018). Registered Replication Report on Srull and Wyer (1979). *Advances in Methods and Practices in Psychological Science*, 251524591877748.
<http://dx.doi.org/10.1177/2515245918777487>
- O'Donnell, M., Nelson, L. D., Ackermann, E., Aczel, B., Akhtar, A., Aldrovandi, S., ... Babincak, P. (2018). Registered Replication Report: Dijksterhuis and van Knippenberg (1998). *Perspectives on Psychological Science*, 13(2), 268–294.
<http://dx.doi.org/10.1177/1745691618755704>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716–aac4716. <http://dx.doi.org/10.1126/science.aac4716>
- Peters, S. J., Matthews, M. S., McBee, M. T., & McCoach, D. B. (2014). *Beyond gifted education: Designing and implementing advanced academic programs*. Prufrock.
- Peters, S. J. (2017, September 5). Bright versus gifted: An unnecessary distinction. [blog post]. Retrieved from <https://www.nagc.org/blog/bright-vs-gifted-unnecessary-distinction>
- Plucker, J. A., & Callahan, C. M.-. (2020). *Critical issues and practices in gifted education: a survey of current research on giftedness and talent development*. Prufrock Press Inc.
- Popper, K. (1959). *Logic of Scientific Discovery*. Hutchison.
- Prinz, F., Schlange, T., & Asadullah, K. (2011). *Believe it or not: how much can we rely on published data on potential drug targets?* *Nature Reviews Drug Discovery*, 10(9), 712.
<http://dx.doi.org/10.1038/nrd3439-c1>
- Rohrer, J. M. (2018). Thinking Clearly About Correlations and Causation: Graphical Causal Models for Observational Data. *Advances in Methods and Practices in Psychological Science*, 27–42. <http://dx.doi.org/10.1177/2515245917745629>

- Rosenthal, R. (1979). File drawer problem and tolerance for null results. *Psychological Bulletin*, 86(3), 638-641. <http://dx.doi.org/10.1037/0033-2909.86.3.638>
- Schimmack, U. (2012). The ironic effect of significant results on the credibility of multiple-study articles. *Psychological Methods*, 17(4), 551–566. <https://dx.doi.org/10.1037/a0029487>
- Silverman, L. (1997). The construct of asynchronous development. *Peabody Journal of Education*, 72(3-4), 36-58.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, 22(11), 1359–1366. <https://dx.doi.org/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(2), 534–547. <https://dx.doi.org/10.1037/a0033242>
- Szabos, J. (1989). Bright child, gifted learner. *Challenge*, 34. Good Apple.
- Taibbi, C. (2012, January 29). "Bright Child" vs. "Gifted Learner": What's the Difference? *Psychology Today*. Retrieved from <https://www.psychologytoday.com/us/blog/gifted-ed-guru/201201/bright-child-vs-gifted-learner-whats-the-difference>
- Tuttle, C. C., Gleason, P., Knechtel, V., Nichols-Barrer, I., Booker, K., Chojnacki, G., ... & Goble, L. (2015). Understanding the Effect of KIPP as It Scales: Volume I, Impacts on Achievement and Other Outcomes. *Final Report of KIPP's "Investing in Innovation Grant Evaluation."* Mathematica Policy Research, Inc.
- van der Zee, T., & Reich, J. (2018). Open Education Science. *AERA Open*. <https://dx.doi.org/10.1177/2332858418787466>
- van Zwet, E. & Cator, E. (2020). The significance filter, the winner's curse and the need to shrink. *ArXiv*. Retrieved from <https://arxiv.org/abs/2009.09440>
- Vazire, S. (2020, January). Do we want to be credible or incredible? *Association for Psychological Science Observer*. Retrieved from <https://www.psychologicalscience.org/observer/do-we-want-to-be-credible-or-incredible>
- Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R., McCarthy, R. J., ... Bakos, B. E. (2018). Registered Replication Report on Mazar, Amir, and Ariely (2008). *Advances in Methods and Practices in Psychological Science*, 1(3), 299-317. <http://dx.doi.org/10.1177/2515245918781032>
- Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R. B., ... Blouin-Hudon, E.-M. (2016). Registered Replication Report: Strack, Martin, and Stepper (1988). *Perspectives on Psychological Science*, 11(6), 917–928. <http://dx.doi.org/10.1177/1745691616674458>
- Yeager, D. S., & Dweck, C. S. (2012). Mindsets that promote resilience: When students believe that personal characteristics can be developed. *Educational Psychologist*, 47(4), 302-314. <http://dx.doi.org/10.1080/00461520.2012.722805>

Yeager, D. S., Hanselman, P., Walton, G. M., Murray, J. S., Crosnoe, R., Muller, C., ... Dweck, C. S. (2019). A national experiment reveals where a growth mindset improves achievement. *Nature*. 573. <http://dx.doi.org/10.1038/s41586-019-1466-y>