FIVE MODELS OF SCIENCE, ILLUSTRATING HOW SELECTION SHAPES METHODS

PAUL E. SMALDINO*

To appear in *The Dynamics of Science: Computational Frontiers in History and Philosophy of Science*, edited by Grant Ramsey and Andreas De Block, University of Pittsburgh Press. Expected 2021.

1. INTRODUCTION

“The good thing about science is that it’s true whether or not you believe it.” This oft-repeated quote, attributed to the astrophysicist and TV presenter Neil DeGrasse Tyson, was seen everywhere at the March for Science, a set of gatherings held around the world on April 22, 2017. The quote has become a rallying cry for supporters of science—and of the application of scientific knowledge in daily life—against widespread science denialism. And of course, science should be defended. Carl Sagan, Tyson’s predecessor as host of *Cosmos*, noted that science not only increases our knowledge of the world but also serves as a bulwark against superstition and charlatanry (Sagan, 1996). However, there is a counterpoint to Tyson’s claim. Plenty of science, or at least scientific results, are *not* true.

During the first decade of the 21st century, the biotech company Amgen attempted to confirm the results of 53 oncology published papers deemed ‘landmark’ studies. Of these, they claim to have successfully replicated only six¹ (Begley & Ellis, 2012). In 2015, a team of 270 researchers calling themselves the Open Science Collaboration repeated 100 studies from published psychology papers. Of these, they successfully replicated only 39 results (Open Science Collaboration, 2015). In 2016, neuroscientists discovered design errors in the most popular statistical packages used to analyze fMRI data, indicating that as many as 70% of the results obtained using these packages may be false positives (Eklund et al., 2016). And in 2018, a team of social scientists targeted 20 high profile studies published in the prestigious journals *Science* and *Nature*, and successfully replicated only 12, and even among these, most of the effects turned out to be smaller than originally published (Camerer et al., 2018). Indeed, a survey conducted by *Nature* in 2016 revealed that a large proportion of empirical scientists, hailing from fields as diverse as chemistry, biology, physics, earth sciences, and medicine, had failed to replicate other researchers’ results (Baker, 2016).

*DEPARTMENT OF COGNITIVE AND INFORMATION SCIENCES, UNIVERSITY OF CALIFORNIA, MERCED
E-mail address: paul.smaldino@gmail.com.

¹Unfortunately, Amgen signed secrecy agreements with the scientists who produced the original work, so the specific studies that failed to replicate are unknown (Harris, 2017).

This is a problem. Our understanding of the world relies on facts. Charles Darwin understood the perniciousness of false facts, writing in *The Descent of Man*, “False facts are highly injurious to the progress of science, for they often endure long; but false views, if supported by some evidence, do little harm, for every one takes a salutary pleasure in proving their falseness; and when this is done, one path towards error is closed and the road to truth is often at the same time opened.” What he is saying in his overwrought Victorian prose is that we shouldn’t worry too much about false theories, because academics are competitive and love to take each other down a peg by demonstrating logical inconsistencies in one another’s theories. Since logic is a common language in science, the competition for theoretical explanations remains relatively healthy. However, any coherent explanation must rely on a firm foundation of facts. If our facts are false, we end up wasting our time arguing about how best to explain something that isn’t even true.

Science involves both theory building and fact finding. This chapter focuses on the fact-finding aspect, and as a shorthand the search for facts is what I will mean henceforth by the term “science.” In this sense, science can be viewed as a process of signal detection for facts. We wish to discover true associations between variables. However, our methods for measurement are imprecise. We sometimes mistake noise for signal, and vice versa.

How we conceptualize the scientific enterprise shapes how we go about the business of conducting research as well as how we strive to improve scientific practices. In this chapter, I’ll present several models of science. I’ll begin by showing ways in which the classic “hypothesis testing” model of science is misleading, and leads to flawed inferences. As a remedy, I’ll discuss models that treat science as a population process, with important dynamics at the group level that trickle down to the individual practitioners. Science that is robust and reproducible depends on understanding these dynamics so that institutional programs for improvement can specifically target them.

2. A FIRST MODEL OF SCIENCE: HYPOTHESIS TESTING

Early in our schooling, many of us are taught a simple and somewhat naïve model of science as “hypothesis testing” (Figure 1). The scientist comes up with a hypothesis about some natural system. She cannot directly infer the essential epistemic state of the hypothesis, whether it is true or false. Instead, she investigates the hypothesis by experimentation or other empirical means, which results either in a positive result in support of the hypothesis or a negative result indicating a lack of support. The alignment between her results and the epistemic state of the hypothesis is necessarily imprecise. There is some risk of a false positive, $\alpha = \Pr(+|F)$, as well as a false negative, $\beta = \Pr(-|T)$. These outcomes are sometimes called Type 1 and Type 2 errors, respectively². This uncertainty forces us to ask: how confident should our scientist be in her results?

²In reality, many scientific results carry explicit uncertainty, and are more probabilistic than implied by my simple dichotomous characterizations of hypothesis states and experimental outcomes. I use these dichotomies for ease of explanation, as an excessive appeal to nuance can hinder the development of theoretical explanations (Healy, 2017). Readers interested in nuance of this sort are directed to Gelman & Carlin (2014), as well as many other works in the philosophy of science.

	<div style="border: 1px solid black; border-radius: 50%; padding: 2px; display: inline-block;">T</div>	<div style="background-color: black; color: white; border-radius: 50%; padding: 2px; display: inline-block;">F</div>	
Result	$1 - \beta$	α	positive results
+			
-	β	$1 - \alpha$	negative results

FIGURE 1. A first model of science. Hypotheses are investigated and results, with characteristic error rates, are recorded. The real epistemic state of each hypothesis, true or false (T or F), is unknowable except through this sort of investigation.

Consider the following scenario. Dr. Pants investigates one of her many hypotheses. Using her well-tested method, the probability that the test will yield a false positive result is 5%. That is, $\Pr(+|F) = 0.05$. If the hypothesis is true, the probability that the test will correctly yield a positive result is 50%. That is, $\Pr(+|T) = 0.5$. The test is conducted, and the result is positive! Now, what is the probability that Dr. Pants' hypothesis is correct?

You may be tempted to answer 95%. After all, the probability of a false positive is 5%, and it's clear that $100 - 5 = 95$. If this is your answer, you are not alone. When a version of this question was posed to students with scientific training, 95% was indeed the most common answer, at least in years past (Gigerenzer & Hoffrage, 1995). Why is this wrong? Recall that we are looking for the probability that the hypothesis is true conditional on obtaining a positive result, $\Pr(T|+)$. Fortunately, we have a handy mathematical tool for computing exactly this sort of conditional probability. Using Bayes' Theorem, we can write out our conditional probability as follows:

$$(1) \quad \Pr(T|+) = \frac{\Pr(+|T) \Pr(T)}{\Pr(+|T) \Pr(T) + \Pr(+|F)(1 - \Pr(T))}$$

You'll notice right away that there's a term in this equation I haven't provided: $\Pr(T)$. This is the prior probability that *any* hypothesis being tested by Dr. Pants is true, often called the *base rate*. We ignore the base rate at our peril.

3. A SECOND MODEL OF SCIENCE: HYPOTHESIS SELECTION AND INVESTIGATION

Imagine now that Dr. Pants tests not one but one hundred hypotheses. Of these, ten are true, and 90 are false. If you want a more concrete example, imagine Dr. Pants runs a behavioral genetics lab. She is looking for single nucleotide polymorphisms (SNPs) that correlate with a heritable behavioral disorder. She tests 100 SNPs, of which ten are actually associated with the disorder. Thus, the base rate is $b = 0.1$. If this seems low, consider that for many disciplines, the base rate may actually be much lower. Every association tested, every statistical test run is a hypothesis that may be supported. Dr. Pants tests her hypotheses using the method described in the previous paragraph, with $\alpha = .05$ and $\beta = .5$. So, what is the probability that a hypothesis with a positive result actually reflects

FIGURE 2. The importance of base rate. Left: 100 hypotheses are tested, of which 10 are true (the base rate is $b = 0.1$). Right: 50% of the true hypotheses and 5% of the false hypotheses yield positive results, producing a posterior probability that a positive result is actually true of approximately $\Pr(T|+) = 0.5$.

a true hypothesis? In this case, it's 50%, not 95% (Figure 2). And the lower the base rate, the lower this posterior probability gets. Worse yet, in reality we can never know for certain the epistemic states of our hypotheses, nor can we easily estimate the base rate. Our results are all we have.

So now we have a second model of science that includes the process of hypothesis selection as well as the experimental investigation of that hypothesis³ (Figure 3). We can capture this model in terms of the posterior probability that a positive result indicates a true hypothesis using the notation introduced so far:

$$(2) \quad \Pr(T|+) = \frac{(1 - \beta)b}{(1 - \beta)b + \alpha(1 - b)}$$

This Bayesian model of science was introduced by Ioannidis (2005) in his now classic paper, "Why most published research findings are false." The analysis is straightforward. If the base rate, b is low, then even a moderate false positive rate (such as 5%) will lead to a low posterior probability, and large number of false positives.

One concern about this model is that it treats each hypothesis in isolation. It ignores the social and public aspect of science. Scientists don't just produce results, they also try to publish them, and some results are easier to publish than others. Once published, results can then be replicated, and with new information comes the opportunity for new estimates of the epistemic states of the underlying hypothesis.

³The model portrays hypothesis selection as preceding investigation. In reality, new hypotheses do sometimes arise at various stages in the scientific process, but for simplicity we organize the model of science in this stylized way.

FIGURE 3. A second model of science. Investigation is preceded by hypothesis selection. The inner circles indicate the real epistemic value of each hypothesis. Black indicates false, white indicates true. The gray outer circle represents the fact that these epistemic values are unknown before investigation.

4. A THIRD MODEL OF SCIENCE: THE POPULATION DYNAMICS OF HYPOTHESES

The first two models of science both portray a science in which each hypothesis is investigated in isolation. But consider what happens to a result once the hypothesis has been investigated. The researcher will sometimes decide to publish the result. I say “sometimes,” because some results are never published, especially when they don’t support the hypotheses being tested. These results end up in the “file drawer” (Rosenthal, 1979). Once published, the studies supporting a given hypothesis can be replicated, whether by other labs or by the one that generated the original result.

Our third model conceptualizes hypothesis testing as a dynamical system involving a large number of hypotheses being tested by a large number of scientists (Figure 4). A scientist first selects a hypothesis to test. A novel hypothesis is true with probability b , the base rate. The hypothesis is investigated, producing results. These results can then be disseminated to the scientific community via publication. This stage is important, because not all results are published with equal probability. Novel, positive results are usually the easiest to publish. Negative results are published at much lower rates (Fanelli, 2012), possibly due to being rejected by journal editors, but also because they are viewed

FIGURE 4. A third model of science. After hypothesis selection and investigation, results are communicated. Some results end up published, and become part of the literature, which can accrue through replication. This is indicated by the rectangle containing concentric circles. Each layer represents a positive (white) or negative (black) results. Results that are not published end up in file drawers, unknown to the scientific community.

as carrying low prestige for researchers and are therefore never submitted (Franco et al., 2014). Once findings are published, they can be replicated. The results can then be added to the literature, but only if they are published. As results accrue, each hypothesis is associated with a record of positive and/or negative results in the published literature. Because some types of results are more likely than others to be published, the published literature likely reflects a biased record of investigation.

This dynamical model was introduced and analyzed in an earlier paper (McElreath & Smaldino, 2015). Our analysis focused on the probability that a hypothesis was true, conditional on its publication record⁴. For simplicity, we operationalized the publication record as a tally of the net positive findings—that is, the number of positive results minus the number of negative results in the published literature. Although this conditional probability was influenced to some degree by all of the models parameters, we found that the two parameters exerting the largest influence—by far—were the base rate, b , and the false positive rate, α . If the base rate is high (so that most tested hypotheses are true) and

⁴Two related studies are worth mentioning here. First, Romero (2016) examined how meta-analysis of a publication record could fail to correctly identify real effects under constraints of limited budgets, investigator biases, and publication bias. Second, Nissen et al. (2016) examined how the accumulation of biased results can lead to the “canonization of false facts.”

the false positive rate is low (so that most positive results reflect true hypotheses), then a single positive result likely reflects a true hypothesis. However, as base rate decreases and false positive rate increases—to values that, I must add, I view as quite realistic for many disciplines—then more and more successful replications are necessary to instill the same amount of confidence in the truth of a hypothesis.

Above all, this indicates that replication is important for a healthy science (Smaldino, 2015). Indeed, our analysis showed that replication studies are valuable even when they use designs with different methodological power than the original investigations. More than that, we shouldn't be surprised that some results fail to replicate. Some erroneous results are inevitable. When methods are imperfect, both false positives and false negatives may be common. That said, the model also illustrates that improvements to the practices and culture of science should focus on factors that increase the base rate of true hypotheses and lower the rate of false positives results, so as to decrease the number of false facts in the published literature.

A number of factors lead to false discovery. False facts are more common when:

- Studies are underpowered, because small sample sizes lead to false positives and ambiguous results.
- Negative results aren't published, distorting the publication record by eliminating disconfirmatory evidence.
- Statistical techniques are misunderstood, leading to false positives and ambiguous results.
- Surprising results are the easiest to publish, because such results have a low base rate of being true, given priors to the contrary.

Although the factors in this list may be new to some readers, scientists have, in general, been aware of these issues for decades. Why, then, isn't science better? Understanding how scientific practice—and not just scientific knowledge—changes over time requires a new model that includes the scientists themselves in the model dynamics. Before introducing such a model, I'll need to say a few words about some of the incentives that structure human social behavior.

5. A BRIEF INTERLUDE ON INCENTIVES

Science is the search for truth about the natural world, for a better understanding of our universe. Scientists, however, are also human beings who need steady employment and the resources to conduct their research. Obtaining those jobs and securing that funding is far from trivial these days. There are currently far more PhDs looking for employment in academia than there are permanent positions for them to fill. In several disciplines, including biomedicine and anthropology, the creation of new PhDs outpaces the creation of new faculty positions by a factor of five (Ghaffarzadegan et al., 2015; Speakman et al., 2018). More generally, the number of open faculty positions in scientific disciplines is only a small fraction of the number of total PhDs awarded each year (Cyranoski et al., 2011; Schillebeeckx et al., 2013). This creates a bottleneck at which selection is non-random.

In academic science, this selection pressure is often linked to an individual’s publication history, as evinced by the clichéd admonition to “publish or perish.”

Successful scientists are certainly publishing more. Since just the early 2000s, the number of publications at the time of hiring for new faculty has more than doubled in fields such as evolutionary biology (Brischoux & Angelier, 2015) and cognitive psychology (Pennycook & Thompson, 2018). A large study of over 25,000 biomedical scientists showed that scientists who ended up become PIs were consistent publishing more papers and placing them in higher-impact journals than those researchers who ended up leaving academia (van Dijk et al., 2014).

It may not be immediately obvious that preferential reward for productivity and impact factor are bad things. Indeed, it seems that we should *want* scientists to be productive and we should *want* their work to have a wide impact. Don’t we want our scientists to be awesome? The difficulty is that awesomeness is in reality quite complicated and multidimensional. The importance of research may not be manifest for quite some time, and a lack of productivity can just as easily reflect careful study of a difficult problem as it can a lack of drive. This difficulty becomes a serious problem when awesomeness is assessed with crude, quantitative metrics like paper count, journal impact factor, and *h*-indices. It has been widely noted by savvy social scientists that, as Campbell (1976, p. 49) noted, “The more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor.” When incentives to publish drive scientists, science itself may become distorted.

There is evidence that scientists do, in fact, respond to incentives. In China, as in several other countries, PIs are often give cash rewards for publishing in top English-language journals. This system began in the early 1990s with small rewards, but the size of the rewards has grown tremendously. As of 2016, Chinese researchers were paid, on average, \$984 for a paper in *PLOS ONE*, \$3,513 for a paper in the *Proceedings of the National Academy of Sciences*, and a whopping \$43,783 for a first-author paper in *Science* or *Nature* (Quan et al., 2017). Correspondingly, in the years between 2000 and 2009, Chinese submissions to the journal *Science* nearly quadrupled (Franzoni et al., 2011). China was recently declared the world’s largest producer of scientific papers (Tollefson, 2018). Such cash-for-papers incentives can be found in several other countries, including India, Korea, Malaysia, Turkey, Venezuela, and Chile (Quan et al., 2017). The West is not immune either. For example, I recently had dinner with some American psychologists, who told me with pride about how much their graduate students published. Their program provided a cash prize of several hundred dollars for the best student paper each year. When I asked how they assessed the best paper, they told me that a first-author publication in a top journal was the best indicator. “Do you read all the papers?” I asked. The answer was no; the journal’s reputation was deemed a sufficient mark of quality. It is not hard to see how students in this program are incentivized not only to produce papers, but to produce a particular type of paper.

Evidence that scientists respond to incentives can be more subtle. Vinkers et al. (2015) looked at relative word frequencies in PubMed abstracts between 1974 and 2014. They

found dramatic increases in the frequencies of positive, congratulatory words. Frequencies of the words “innovative” and “groundbreaking” had each increased 2500%. Frequency of “novel” had increased 4000%. And frequency of “unprecedented” had increased 5000%. There are, of course, two possible explanations for this shift in word frequencies. The first is that contemporary scientific research is actually 25 times more innovative than it was 40 years ago. The other, a smidge more likely, is that scientists are responding to incentives to distinguish their work as important and pathbreaking.

A system that rewards novel, innovative results can—and does—incentivize cheating. Recent examples include Jan Schön in physics, Diederik Stapel in psychology, and Brian Wansink in nutrition science. A personal favorite is a case of fraud uncovered by the editors of the *British Journal of Clinical Pharmacology*. The authors of a paper claiming impressive results suggested as reviewers several prominent scholars in their field. These scholars were contacted as reviewers, and all returned glowing reviews within just a few days. One of the editors grew suspicious at the quick responses from the busy big-shot scientists, and contacted them at the email addresses listed on their university web pages. They were all surprised by the emails, because none of them had heard of the paper in question. The explanation: when the authors submitted their manuscript, they had provided fake email addresses for their suggested reviewers, and submitted forged reviews of their own paper⁵ (Cohen et al., 2016).

Fraud surely happens, but it’s also probably the exception rather than the rule. Most scientists are well-meaning people who want to learn about the world. The problem is that incentives for maximizing simple quantitative metrics, which act as proxies for more meaningful but multifaceted concepts like productivity and influence, can be detrimental even if all actors are well-intentioned. To help explain why, we’ll turn to a new model of science that includes the scientists as well as the hypotheses.

6. A FOURTH MODEL OF SCIENCE: VARIATION, HERITABILITY, AND SELECTION

Science is a cultural process that, like many cultural processes, evolves through a Darwinian process (Richerson & Boyd, 2005; Mesoudi, 2011; Smaldino, 2014; Smaldino & McElreath, 2016). Philosophers of science including Campbell (1965), Popper (1979), and Hull (1988) have discussed how scientific theories evolve by variation and selective retention. But scientific methods can also evolve. Darwinian evolution requires three conditions to occur:

- (1) There must be variation.
- (2) That variation must have consequences for survival or reproduction.
- (3) Variation must be heritable.

Research practices and methods certainly vary. That variation leads to differences in the sorts of results that are produced and, consequently, the publications that arise from those results. These publications have consequences in determining who is successful in terms of getting hired and promoted, securing grants, attracting graduate students and postdocs, and placing those trainees in positions heading their own research groups. And variation

⁵It should go without saying that you should never do this.

FIGURE 5. A fourth model of science. Dynamics occur in two phases: Science and Evolution. In the evolution stage, labs compete for research positions. A lab’s methods are represented by its shading, and its prestige is represented by its size. When a new position opens, it is more likely to be filled by someone using the methods of more prestigious labs.

in practice is partly heritable, in the sense that trainees acquire research habits and statistical procedures from mentors and peers. Researchers also acquire research practices from successful role models in their fields, even if they do not personally know them. Therefore, when researchers are rewarded primarily for publishing, habits that promote publication are likely to be passed on.

If we want to understand how we might minimize false discoveries, we need a model of science that includes variation among scientists. This model has two phases: Science and Evolution (Figure 5). In the Science phase, each research lab chooses and investigates hypotheses and tries to publish their results, just as in our third model of science. However, the methods used by each lab can differ, which affects the rate at which they conduct research and the probability of certain results. More specifically, consider a population of labs, all conducting research. We make the following assumptions:

- Each lab has characteristic methodological power, $\Pr(+|T)$.
- Increasing power also increases false positives, unless effort is exerted. This is because it is easy to have perfect power if every result is positive, but correctly eliminating the false hypotheses requires additional work⁶.
- Additional effort also increases the time between results, because each study requires more work.
- Negative results are harder to publish than positive results.
- Labs that publish more are more likely to have their methods “reproduced” in new labs.

⁶This is essentially the logic of signal detection theory.

This model was first presented and analyzed in another paper with Richard McElreath (Smaldino & McElreath, 2016). First, we found that if effort is held constant and power is allowed to evolve, power evolves to its maximum value and the false discovery rate (the proportion of published results that are incorrect) skyrockets. Everything is deemed “true,” and we have no information about anything. This scenario is pretty unrealistic. We have fairly good ways of assessing the power of research methods, and no one would ever allow this to happen. However, *effort* is notoriously difficult to assess. If we hold power constant and allow effort to evolve, we find that effort reliably evolves to its minimum value, and once again the false discovery rate balloons. To reiterate, this dynamic requires no cheating or strategizing on the part of our agents, only that publication is a determinant of job placement. We have referred to this dynamical process as “the natural selection of bad science” (Smaldino & McElreath, 2016).

What does this mean? It means that if our model of science is at least moderately realistic, and incentives for publishing do drive selection on research methods, then we should see evidence for impediments to the improvement of scientific methods on the time scale of generations. If, on the other hand, incentives are rewarding methodological rigor, we should see a steady increase in the quality of methods for scientific inquiry.

In 1967, Paul Meehl cautioned about the misuse of p -values, warning that scientists were wrongly interpreting their meaning and consequently generating lots of false positives (Meehl, 1967). In 2016, the American Statistical Association published their “Statement on p -values,” cautioning about their misuse, and warning that scientists were wrongly interpreting their meaning and consequently generating lots of false positives. They bemoaned, “Let us be clear. Nothing in the ASA statement is new. Statisticians and others have been sounding the alarm about these matters for decades, to little avail.” (Wasserstein & Lazar, 2016, p. 130).

In 1962, Jacob Cohen published a meta-analysis of abnormal and social psychology experiments, noting the frustratingly low statistical power of most published research (Cohen, 1962). He cautioned that many studies were not sufficiently powered to adequately provide confirming or disconfirming evidence, leading to an excess of spurious results. In the late 1980s, two studies provided new meta-analyses investigating whether there had been any improvement to the average statistical power of psychological research (Sedlmeier & Gigerenzer, 1989; Rossi, 1990). They found no improvement. Recently, Richard McElreath and I updated those studies and confirmed that, on average, there was no improvement to the average statistical power in the social and behavioral sciences through 2011, with an average power to detect small effects of 0.24⁷ (Smaldino & McElreath, 2016). Szucs & Ioannidis (2017) provided a focused study of 10,000 papers published in psychology, medicine, and cognitive neuroscience journals between 2011 and 2014, and similarly found very low power in all three fields.

The natural selection of bad science appears to be pernicious. I previously noted the importance of replication for assessing the true epistemic value of hypotheses. Could replication similarly help to curb the degradation of methods? One particularly interesting,

⁷Of course, many individual studies had quite high power.

if extreme, suggestion came from Rosenblatt (2016), who proposed that the authors of each published paper, or their host institutions, sign a contract committing them to pay a fine if their studies fail to replicate. Let me be clear: this is a terrible idea. As stated earlier, occasional failure to replicate is to some extent the price of doing business in scientific research. However, it is one of the more concrete suggestions for using replication to improve science. So, we put it—or something like it—into the model.

Under our replication extension, all labs committed a proportion r of their investigations to replicating previously published results. We assumed that all replications were publishable regardless of the result, and carried half of the prestige carried by a novel positive finding⁸. If another lab replicated a finding successfully, the lab that published it originally got a small boost in prestige. If another lab *failed* to replicate a finding successfully, the original authors suffered a tremendous loss of prestige. To be honest, we thought this extreme intervention would curb the decline in effort and the runaway false discovery rate. In hindsight, it is clear why it didn't. Although some labs did suffer a huge loss of prestige, the most successful labs were still those who cut corners and avoided being caught.

Incentive structures that push scientists to boost quantitative metrics like publication counts and impact factors can lead to the degradation of methods. This dynamic requires no fraud or ill intent on the part of individual actors, only that successful individuals transmit their methods⁹. From this, we might conclude that changing individual behavior—each of us improving our methods—is not sufficient to improve scientific methods; this requires institutional change. Specifically, it requires that the selection bottlenecks of hiring and promotion are not overly focused on those metrics, but can instead provide a more nuanced assessment of researcher quality that maintains high methodological integrity.

Unfortunately, institutional change is far from easy. For the most part, institutions are not *meant* to change easily. They provide a stable framework that structures social interactions and exchanges, and ensure some consistency to the operation of a society in the absence of enforcement by specific individuals (North, 1990). This means that we run into trouble when our institutions are unhealthy. If we are to change the institutional incentives for publishing in academic science, we should be aware that such change will likely be slow. Is there anything else that can be done in the short run?

There are many efforts currently underway to improve the norms and institutions of academic science regarding rigor and reproducibility, often under the banner of the “Open Science” movement (Nosek et al., 2012; Munafò et al., 2017). Some of these new norms include preregistration and registered reports (Nosek & Lakens, 2014; Chambers, 2017),

⁸The results are qualitatively the same as long as replications are worth more than nothing and less than the prestige of a novel finding.

⁹Selection for publications is of course not the only force shaping the cultural evolution of science, and methodological rigor is not the only behavior that cultural selection acts upon. For example, Holman & Bruner (2017) showed how industry funding can shape results to align with corporate interests in a manner reminiscent of Noam Chomsky's famous quip to a journalist: “I'm sure you believe everything you're saying. But if you believed something different, you wouldn't be sitting here.” Likewise, O'Connor (2018) used a model similar to ours to study how selective forces shape risk taking and conservatism in hypothesis selection.

preprints (Bourne et al., 2017; Smaldino, 2017b), double blind and open peer review (Mulligan et al., 2013; Okike et al., 2016; Tomkins et al., 2017), and better training in methods, statistics, and philosophy of science. At the same time, funding agencies are increasingly paying attention to what gets funded, and some have been shifting how they fund new research projects. How do these developments influence the conclusions from our fourth model of science?

7. A FIFTH MODEL OF SCIENCE: FOLLOW THE MONEY

Our fourth model of science makes several pessimistic—if realistic—assumptions about the way academic science works in our era. However, changes in just the last few years prompt us to challenge some of these. I want to focus on three specific assumptions, and discuss what happens when we relax or alter them.

Assumption 1: Publishing negative results is difficult or confers little prestige. This assumption is realistic, because negative results are rarely published (Fanelli, 2012) or even submitted (Franco et al., 2014). However, there is an increasingly large push to publish negative results. Many journals now accept registered reports, in which the research plan is peer reviewed before a study is conducted. Once approved, the paper’s acceptance is contingent only on adherence to the submitted plan, and not on the character of the results (Nosek & Lakens, 2014; Chambers, 2017). A recent study by Allen & Mehler (2019) found that among studies using registered reports, 61% of results did not support the authors’ original hypotheses, compared to estimates of 5-20% of null findings in the wider literature¹⁰. What if publication bias against negative results were eliminated?

Assumption 2: Publishing positive (confirmatory) results is always possible. This assumption ignores the corrective role of peer review in maintaining high quality research. The assumption is realistic, because there is little evidence that peer reviewers can act as effective gatekeepers against false discovery. The many failed replications discussed in this chapter’s Introduction testify to that. Peer review may in many cases be more about maintaining group norms than about weeding out error. There is widespread evidence that peer reviewers can be biased toward prestigious individuals and institutions, and against authors who are women and underrepresented minorities (Budden et al., 2008; Tomkins et al., 2017). If peer review was reliable, we should expect consistency between reviewer recommendations. Instead, a number of studies have found low correlation between reviewer decisions on grant panels (Cole et al., 1981; Marsh et al., 2008; Mutz et al., 2012), conference proceedings (Langford & Guzdial, 2015; Deveugele & Silverman, 2017), and journals articles (Peters & Ceci, 1982; Cicchetti, 1991; Nicolai et al., 2015).

Nevertheless, we increasingly see efforts to improve the conditions that facilitate effective peer review. Registered reports remove biases based on the novelty or expectedness of a study’s results (Nosek & Lakens, 2014; Chambers, 2017). Double-blind peer review aims

¹⁰Of course, this result might be at least partly driven by researchers being more likely to use registered reports when they are less confident in their hypotheses.

to reduce biases (Mulligan et al., 2013; Okike et al., 2016; Tomkins et al., 2017). Journals increasingly require or incentivize open data and methods, which improves the ability of peer reviewers to assess results, and the increased use of repositories such as OSF and GitHub has helped to facilitate this behavior. Open peer review and the increased use of preprint servers also allow for a greater number of critical eyes to read and comment on a manuscript before it is published (Bourne et al., 2017; Smaldino, 2017b). And better training in statistics, logic, and best research practices—as evidenced by the popularity of books, MOOCs, podcasts, symposia and conferences on open science—may promote more informed reviews. What if peer review was effective at filtering out false discovery?

Assumption 3: Research productivity is constrained only by the ability to complete projects. This assumption ignores the role of funding, which is required for much scientific research. This assumption was justified by the desire to ignore differences in access to funding and focus on the bottlenecks at hiring and promotion. Moreover, if one assumes that success in securing grant funding results from success in the quantity and prestige of one’s publications, then including explicit funders in the model is unnecessary. Instead, what if funders ignored publication records, or even focused on funding projects with the most rigorous methods?

The norms of hiring and promoting researchers based on simple metrics are entrenched in deeply rooted tradition and diffuse across many academic institutions; they will not be changed quickly or easily. In contrast, the recent changes highlighted above are occurring rapidly, due to greater top-down control from journals and funders. To investigate the consequences of these changes, we will once again revise our model of science.

We again consider a finite population of labs. Each lab has a characteristic methodological rigor (or lack thereof), which is linked to the false positive rate of the results they obtain. In our fourth model, a lab’s productivity was limited only by its rigor. This time, investigating hypotheses requires funding. Each lab is initialized with some startup funds it can use to conduct research. Once these funds are exhausted, additional funds must be acquired from grant agencies.

To our two phases of Science and Evolution, we add a third: Grant Seeking (Figure 6). In the Grant Seeking phase, a subset of labs apply for funding, and the one that best matches the funding agency’s allocation criteria is awarded a grant. We might consider any number of strategies. My colleagues and I have considered those based on publication quantity, funding labs at random, and targeting those labs with the most rigorous methods. The Science phase looks quite similar to our previous models, having three phases of hypothesis selection, investigation, and communication. Here we may also take the opportunity to study changes to peer review and publication bias as discussed. In the communication phase, positive results are always published, and negative results are published with probability p . Erroneous results (in which the result does not reflect the real epistemic state of the hypothesis) are successfully blocked during peer review with probability r . The Evolution phase works exactly as it did in the previous model, such that labs with more publications are most likely to transmit their methods to the next generation. This

FIGURE 6. A fifth model of science. In addition to the Science and Evolution phases, labs also compete for grant funding, which enables them to conduct more research.

is worth repeating: the selection pressure for publication quantity is still present. For a detailed analysis of this model, see Smaldino et al. (2018). Here, I summarize our main results.

First, we can ask whether, in the absence of any contributions from funding agencies, curbing publication bias and improving peer review can promote substantial improvements to reproducible science. There is bad news, then good news, and then bad news again. The bad news is that, taken one at a time, each of these improvements must be operating at nearly maximum levels for any improvements to occur. That is, negative results must be published at equal rates as positive results, and peer reviewers must be nearly perfect in detecting false discoveries. The good news is that the effects of these two interventions are additive, so that moderate improvement to both publication bias and peer review *can* decrease the rates of false discovery to some extent. The bad news (again) is that this effect operates on the published literature, so that more published results are true, but does little to improve the quality of the scientists who produce that published research, at least in terms of methodological rigor. We get bad scientists, but institutions that don't allow them to publish their worst work. This is doubly troubling if we then expect those same corner-cutting researchers to perform exemplary peer review.

We next turned to an exploration of funding strategies. We first studied very simple strategies and found that a strategy of purely random funding allocation is little better than directly funding labs based on publication history. We did find that if funding agencies could effectively target those research groups using the most rigorous methods, the degradation of research quality can be completely mitigated. This is, however, a big “if.” Rigor is notoriously difficult to assess, and it is probably quite unrealistic to assume that funders could consistently and accurately infer the quality of a lab’s methods. So it appears at first glance that random allocation is unhelpful, and that funding focused on rigor works but is probably a pipe dream.

These results were discouraging, to say the least. However, we then started paying more attention to the emerging literature on *modified funding lotteries*, which incorporate aspects of funding strategies focused on both randomness and rigor. Recently, a number of scholars and organizations have supported a type of lottery system for allocating research funds (Barnett, 2016; Fang & Casadevall, 2016; Bishop, 2018; Avin, 2018; Gross & Bergstrom, 2019), usually proposing that a baseline threshold for quality must first be met in order to qualify projects for consideration in the lottery. Although rigor may be difficult to assess precisely, at least *some* information about the integrity of a research lab is often available. Such lotteries may confer advantages not directly related to reproducibility, including (i) promoting a more efficient allocation of researchers’ time (Gross & Bergstrom, 2019); (ii) increasing the funding of innovative, high-risk-high-reward research (Fang & Casadevall, 2016; Avin, 2018); and (iii) reducing gender and racial bias in funding, as well as systematic biases arising from repeat reviewers or proposers coming from elite institutions (Fang & Casadevall, 2016). Such biases can lead to cascading successes that increase the funding disparity between those who, through luck, have early successes, and those that don’t (Bol et al., 2018). However, the potential influence of modified lotteries on reproducibility had not previously been studied.

We investigated a funding strategy in which funds were awarded randomly to the pool of *qualified* applicants. Applicants were qualified if their methodological rigor (equivalent to the inverse of their characteristic false positive rate) did not fall below a threshold. We found that this strategy could be extremely effective at reducing false discoveries, even when using fairly modest thresholds (such as restricting funding to labs with false positive rates below 30%). Even better, when modified lotteries were paired with improvements to peer review and publication bias, the model produced dramatic improvements to both the scientific literature *and* the scientists producing that literature. This indicates that funders who prioritize research integrity over innovation or productivity may be able to exert a positive influence over the landscape of scientific research above and beyond the individual labs they fund.

Many of the interventions heralded by the Open Science movement—including registered reports, preprints, open data, and the like—have undeniable value. This model indicates that these interventions are likely to be insufficient to sustain the persistence of high quality research methods as long as there are strong incentives for maximizing simple quantitative metrics like publication quantity and impact factor, which act as proxies for desirable but complex and multi-faceted traits. On the other hand, the model also provides room

for cautious optimism. Even in the face of strong selective pressures for publication at the key bottlenecks of hiring and promotion, science may nevertheless be improved by countervailing pressures at other bottlenecks, such as the competition for funding, if they promote rigor at the cost of productivity.

8. DISCUSSION

This is a chapter about how institutional incentives shape behavior in academic science. Methods are shaped by cultural selection for practices that help researchers optimize the criteria on which they are judged, hired, and promoted. Selection can shape practices even in the absence of strategic behavior to change those practices. If methods are heritable, selection is sufficient to be damaging. The improvements promoted by the Open Science movement, as well as by well-intentioned funding agencies, are important. The models indicate that they can do some good. Beyond what is captured by the models, these practices may produce normative shifts by becoming associated with prestige, and by promoting the informal punishment of transgressors. However, the models also indicate that Open Science practices are not sufficient if selection continues to favor easily measured evaluation metrics over more holistic, multi-dimensional assessments of quality. This conclusion forces us to consider exactly what properties we want in our academic scientists.

This is also a chapter about cultural evolution. In the last few decades, a new interdisciplinary field has emerged. It has provided formal models, increasingly backed by empirical research, of how individuals maintain cooperative participation (e.g. Boyd & Richerson (1992); Hooper et al. (2010)), how they acquire and transmit cultural information (e.g. Henrich & Gil-White (2001); Kendal et al. (2018)), and how the population dynamics of cultural traits unfold as a result (e.g. Boyd & Richerson (1985, 2002); Mesoudi (2011); Turchin et al. (2013); Waring et al. (2017)). In October 2018, the Cultural Evolution Society had its second meeting in Tempe, Arizona, with over 200 participants representing psychology, anthropology, archaeology, behavioral ecology, genetics, linguistics, economics, sociology, engineering, and mathematics. It behooves those who are interested in the science and sociology of science to pay attention to this field, for its primary focus is cultural stability and the dynamics of cultural change. As a participant, it also appeared to me that much of the science presented was of unusually high quality. It is possible that, when one has to present work to those unfamiliar with the methodological norms of a small subfield, there is a strong incentive to be extraordinarily thorough and transparent. Although field-specific expertise is invaluable in assessing research, it may also be that cross-disciplinary communication has an important role to play in maintaining methodologically rigorous research.

This is also a chapter about models. I have presented a series of five models, each of increasing complexity, to help us understand and explain the process and cultural phenomenon of scientific research. How we model science shapes our ability to identify both problems and solutions. Even at their most complex, models involve drastic oversimplification. The models I have presented focus on hypothesis testing, the fact-finding portion of science, and ignore the critical role of theory building. In these models, hypotheses are

independent of another, rather than interconnected. Hypotheses are formulated as clearly true or false, and results are formulated as unambiguously positive or negative. The later models characterize competition as being solely about publication, whereas network effects and research topics also drive success. Perhaps most importantly, the models ignore innovation and the social significance of results. Taken in isolation, these models represent a fairly crude way of thinking about science. However, the point of a model is not to capture all the nuances of a system. The point of a model is to be stupid (Smaldino, 2017a). By being stupid, a model clarifies the aspects of the system we should be paying attention to, and makes clear the aspects we do not include, forcing us to consider their influence on a system we now at least partially understand. Models are not the sum total of our understanding, but they can scaffold our imaginations toward richer and deeper understanding of complex systems (Haldane, 1964; Schank et al., 2014). The models I have presented have focused on the factors that make positive results more or less likely to represent true facts. That is an important question about how science works, but it is far from the only question. A more complete understanding of the system requires many models with many perspectives and many different stupid oversimplifications. With them we can consider, for example, how false facts are canonized through publication bias (Nissen et al., 2016; Romero, 2016), how funding allocation affects the efficiency of research effort (Avin, 2018; Gross & Bergstrom, 2019), how group loyalties and gatekeeping institutions can stifle innovative paradigms (Akerlof & Michailat, 2018), how scientists select important research questions (Strevens, 2003; Weisberg & Muldoon, 2009; Thoma, 2015; Alexander et al., 2015; Bergstrom et al., 2016; O’Connor, 2018; Zollman, 2018) and how we might develop better theories (Stewart & Plotkin, 2020).

To some extent, this is a chapter about how incentives for publication ruin everything, and how those incentives have to change. However, it should not be taken as a story about how we academics are powerless in the face of the mighty incentives. It’s true that we inherit the culture into which we are born and develop, but it’s also true that we collectively create the culture in which we participate. Collectively, we have the power to change that culture.

REFERENCES

- Akerlof, G. A. & Michailat, P. (2018). Persistence of false paradigms in low-power sciences. *Proceedings of the National Academy of Sciences*, 115(52), 13228–13233.
- Alexander, J. M., Himmelreich, J., & Thompson, C. (2015). Epistemic landscapes, optimal search, and the division of cognitive labor. *Philosophy of Science*, 82(3), 424–453.
- Allen, C. & Mehler, D. M. A. (2019). Open science challenges, benefits and tips in early career and beyond. *PLOS ONE*, 17(5), e3000246.
- Avin, S. (2018). Policy considerations for random allocation of research funds. *RT. A Journal on Research Policy and Evaluation*, 6, 1.

- Baker, M. (2016). Is there a reproducibility crisis? a nature survey lifts the lid on how researchers view the 'crisis' rocking science and what they think will help. *Nature*, 533(7604), 452–455.
- Barnett, A. G. (2016). Funding by lottery: Political problems and research opportunities. *mBio*, 7(4), e01369–16.
- Begley, C. G. & Ellis, L. M. (2012). Drug development: Raise standards for preclinical cancer research. *Nature*, 483, 531–533.
- Bergstrom, C. T., Foster, J. G., & Song, Y. (2016). Why scientists chase big problems: individual strategy and social optimality. *arXiv preprint arXiv:1605.05822*.
- Bishop, D. (2018). Luck of the draw. <https://www.natureindex.com/news-blog/luck-of-the-draw>.
- Bol, T., de Vaan, M., & van de Rijdt, A. (2018). The Matthew effect in science funding. *Proceedings of the National Academy of Sciences*, 115(19), 4887–4890.
- Bourne, P. E., Polka, J. K., Vale, R. D., & Kiley, R. (2017). Ten simple rules to consider regarding preprint submission. *PLoS Computational Biology*, 13(5), e1005473.
- Boyd, R. & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Boyd, R. & Richerson, P. J. (1992). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology*, 13(3), 171–195.
- Boyd, R. & Richerson, P. J. (2002). Group beneficial norms can spread rapidly in a structured population. *Journal of Theoretical Biology*, 215(3), 287–296.
- Brischoux, F. & Angelier, F. (2015). Academia's never-ending selection for productivity. *Scientometrics*, 103(1), 333–336.
- Budden, A. E., Tregenza, T., Aarssen, L. W., Koricheva, J., Leimu, R., & Lortie, C. J. (2008). Double-blind review favours increased representation of female authors. *Trends in Ecology & Evolution*, 23(1), 4–6.
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B. A., Pfeiffer, T., et al. (2018). Evaluating the replicability of social science experiments in nature and science between 2010 and 2015. *Nature Human Behaviour*, 2, 637–644.
- Campbell, D. T. (1965). Variation and selective retention in socio-cultural evolution. In H. Barringer, G. Blanksten, & R. Mack (Eds.), *Social change in developing areas: A reinterpretation of evolutionary theory* (pp. 19–49). Cambridge, MA: Schenkman.
- Campbell, D. T. (1976). Assessing the impact of planned social change. Technical report, The Public Affairs Center, Dartmouth College, Hanover, New Hampshire, USA.
- Chambers, C. (2017). *The seven deadly sins of psychology: A manifesto for reforming the culture of scientific practice*. Princeton University Press.
- Cicchetti, D. V. (1991). The reliability of peer review for manuscript and grant submissions: A cross-disciplinary investigation. *Behavior and Brain Sciences*, 14, 119–186.
- Cohen, A., Pattanaik, S., Kumar, P., Bies, R. R., De Boer, A., Ferro, A., Gilchrist, A., Isbister, G. K., Ross, S., & Webb, A. J. (2016). Organised crime against the academic peer review system. *British Journal of Clinical Pharmacology*, 81(6), 1012–1017.

- Cohen, J. (1962). The statistical power of abnormal-social psychological research: A review. *Journal of Abnormal and Social Psychology*, 65(3), 145–153.
- Cole, S., Simon, G. A., et al. (1981). Chance and consensus in peer review. *Science*, 214(4523), 881–886.
- Cyranoski, D., Gilbert, N., Ledford, H., Nayar, A., & Yahia, M. (2011). Education: the phd factory. *Nature*, 472(7343), 276–279.
- Darwin, C. (1871). *The descent of man and selection in relation to sex*. Murray.
- Deveugele, M. & Silverman, J. (2017). Peer-review for selection of oral presentations for conferences: Are we reliable? *Patient Education and Counseling*, 100(11), 2147–2150.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 201602413.
- Fanelli, D. (2012). Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3), 891–904.
- Fang, F. C. & Casadevall, A. (2016). Research funding: The case for a modified lottery. *mBio*, 7(2), e00422–16.
- Franco, A., Malhotra, N., & Simonovits, G. (2014). Publication bias in the social sciences: Unlocking the file drawer. *Science*, 345(6203), 1502–1505.
- Franzoni, C., Scellato, G., & Stephan, P. (2011). Changing incentives to publish. *Science*, 333(6043), 702–703.
- Gelman, A. & Carlin, J. (2014). Beyond power calculations: Assessing type s (sign) and type m (magnitude) errors. *Perspectives on Psychological Science*, 9(6), 641–651.
- Ghaffarzadegan, N., Hawley, J., Larson, R., & Xue, Y. (2015). A note on phd population growth in biomedical sciences. *Systems Research and Behavioral Science*, 32, 402–405.
- Gigerenzer, G. & Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: frequency formats. *Psychological Review*, 102(4), 684.
- Gross, K. & Bergstrom, C. T. (2019). Contest models highlight inherent inefficiencies of scientific funding competitions. *PLoS Biology*, 17(1), e3000065.
- Haldane, J. B. S. (1964). A defense of beanbag genetics. *Perspectives in Biology and Medicine*, 7(3), 343–360.
- Harris, R. (2017). *Rigor mortis: How sloppy science creates worthless cures, crushes hope, and wastes billions*. Basic Books.
- Healy, K. (2017). Fuck nuance. *Sociological Theory*, 35(2), 118–127.
- Henrich, J. & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22(3), 165–196.
- Holman, B. & Bruner, J. (2017). Experimentation by industrial selection. *Philosophy of Science*, 84(5), 1008–1019.
- Hooper, P. L., Kaplan, H. S., & Boone, J. L. (2010). A theory of leadership in human cooperative groups. *Journal of Theoretical Biology*, 265(4), 633–646.
- Hull, D. L. (1988). *Science as a process: An evolutionary account of the social and conceptual development of science*. Chicago, IL: University of Chicago Press.

- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124.
- Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social learning strategies: Bridge-building between fields. *Trends in Cognitive Sciences*, 22(7), 651–665.
- Langford, J. & Guzdial, M. (2015). The arbitrariness of reviews, and advice for school administrators. *Communications of the ACM*, 58(4), 12–13.
- Marsh, H. W., Jayasinghe, U. W., & Bond, N. W. (2008). Improving the peer-review process for grant applications: reliability, validity, bias, and generalizability. *American Psychologist*, 63(3), 160.
- McElreath, R. & Smaldino, P. E. (2015). Replication, communication, and the population dynamics of scientific discovery. *PLoS One*, 10(8), e0136088.
- Meehl, P. E. (1967). Theory-testing in psychology and physics: A methodological paradox. *Philosophy of Science*, 34(2), 103–115.
- Mesoudi, A. (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences*. University of Chicago Press.
- Mulligan, A., Hall, L., & Raphael, E. (2013). Peer review in a changing world: An international study measuring the attitudes of researchers. *Journal of the American Society for Information Science and Technology*, 64(1), 132–161.
- Munafò, M. R., Nosek, B. A., Bishop, D. V., Button, K. S., Chambers, C. D., du Sert, N. P., Simonsohn, U., Wagenmakers, E.-J., Ware, J. J., & Ioannidis, J. P. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1(1), 0021.
- Mutz, R., Bornmann, L., & Daniel, H.-D. (2012). Heterogeneity of inter-rater reliabilities of grant peer reviews and its determinants: A general estimating equations approach. *PLoS One*, 7(10), e48509.
- Nicolai, A. T., Schmal, S., & Schuster, C. L. (2015). Interrater reliability of the peer review process in management journals. In *Incentives and Performance* (pp. 107–119). Springer.
- Nissen, S. B., Magidson, T., Gross, K., & Bergstrom, C. T. (2016). Publication bias and the canonization of false facts. *eLife*, 5, e21451.
- North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press.
- Nosek, B. A. & Lakens, D. (2014). Registered reports. *Social Psychology*, 45(3), 137–141.
- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific Utopia: II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7(6), 615–631.
- O'Connor, C. (2018). The natural selection of conservative science. *Studies in History and Philosophy of Science Part A*.
- Okike, K., Hug, K. T., Kocher, M. S., & Leopold, S. S. (2016). Single-blind vs double-blind peer review in the setting of author prestige. *JAMA*, 316(12), 1315–1316.
- Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- Pennycook, G. & Thompson, V. A. (2018). An analysis of the Canadian cognitive psychology job market (2006–2016). *Canadian Journal of Experimental Psychology*, 72(2),

71–80.

- Peters, D. P. & Ceci, S. J. (1982). Peer-review practices of psychological journals: The fate of published articles, submitted again. *Behavior and Brain Sciences*, 5, 187–255.
- Popper, K. (1979). *Objective knowledge: An evolutionary approach*. Oxford: Oxford University Press.
- Quan, W., Chen, B., & Shu, F. (2017). Publish or impoverish: An investigation of the monetary reward system of science in china (1999-2016). *Aslib Journal of Information Management*, 69(5), 486–502.
- Richerson, P. J. & Boyd, R. (2005). *Not By Genes Alone: How Culture Transformed Human Evolution*. University of Chicago Press.
- Romero, F. (2016). Can the behavioral sciences self-correct? A social epistemic study. *Studies in History and Philosophy of Science Part A*, 60, 55–69.
- Rosenblatt, M. (2016). An incentive-based approach for improving data reproducibility. *Science Translational Medicine*, 8, 336ed5.
- Rosenthal, R. (1979). The “file drawer problem” and tolerance for null results. *Psychological Bulletin*, 86(3), 638–641.
- Rossi, J. S. (1990). Statistical power of psychological research: What have we gained in 20 years? *Journal of Consulting and Clinical Psychology*, 58(5), 646.
- Sagan, C. (1996). *The demon-haunted world: Science as a candle in the dark*. Ballantine Books.
- Schank, J. C., May, C. J., & Joshi, S. S. (2014). Models as scaffolds for understanding. In L. R. Caporael, J. R. Griesemer, & W. C. Wimsatt (Eds.), *Developing scaffolds in evolution, culture, and cognition* (pp. 147–167). Cambridge, MA: MIT Press.
- Schillebeeckx, M., Maricque, B., & Lewis, C. (2013). The missing piece to changing the university culture. *Nature Biotechnology*, 31(10), 938–941.
- Sedlmeier, P. & Gigerenzer, G. (1989). Do studies of statistical power have an effect on the power of studies? *Psychological Bulletin*, 105(2), 309–316.
- Smaldino, P. E. (2014). The cultural evolution of emergent group-level traits. *Behavioral and Brain Sciences*, 37(3), 243–95.
- Smaldino, P. E. (2015). A theoretical lens for the reproducibility project. <http://smaldino.com/wp/a-theoretical-lens-for-the-reproducibility-project/>.
- Smaldino, P. E. (2017a). Models are stupid, and we need more of them. In R. R. Vallacher, A. Nowak, & S. J. Read (Eds.), *Computational social psychology* (pp. 311–331). Routledge.
- Smaldino, P. E. (2017b). On preprints. <http://academiclifehistories.weebly.com/blog/on-preprints>.
- Smaldino, P. E. & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, 3, 160384.
- Smaldino, P. E., Turner, M. A., & Contreras Kallens, P. A. (2018). Open science and modified funding lotteries can impede the natural selection of bad science. *Royal Society Open Science*, 6, 190194.
- Speakman, R. J., Hadden, C. S., Colvin, M. H., Cramb, J., Jones, K., Jones, T. W., Lulewicz, I., Napora, K. G., Reinberger, K. L., Ritchison, B. T., et al. (2018). Market

- share and recent hiring trends in anthropology faculty positions. *PLoS ONE*, 13(9), e0202528.
- Stewart, A. J. & Plotkin, J. B. (2020). The natural selection of good science. *arXiv preprint arXiv:2003.00928*.
- Strevens, M. (2003). The role of the priority rule in science. *The Journal of Philosophy*, 100(2), 55–79.
- Szucs, D. & Ioannidis, J. P. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biology*, 15(3), e2000797.
- Thoma, J. (2015). The epistemic division of labor revisited. *Philosophy of Science*, 82(3), 454–472.
- Tollefson, J. (2018). China declared world’s largest producer of scientific articles. *Nature*, 553(7689), 390.
- Tomkins, A., Zhang, M., & Heavlin, W. D. (2017). Reviewer bias in single- versus double-blind peer review. *Proceedings of the National Academy of Sciences*, 114(48), 12708–12713.
- Turchin, P., Currie, T. E., Turner, E. A., & Gavrillets, S. (2013). War, space, and the evolution of old world complex societies. *Proceedings of the National Academy of Sciences*, 110(41), 16384–16389.
- van Dijk, D., Manor, O., & Carey, L. B. (2014). Publication metrics and success on the academic job market. *Current Biology*, 24(11), R516–R517.
- Vinkers, V., Tjink, J., & Otte, W. (2015). Use of positive and negative words in scientific PubMed abstracts between 1974 and 2014: Retrospective analysis. *British Medical Journal*, 351, h6467.
- Waring, T. M., Goff, S. H., & Smaldino, P. E. (2017). The coevolution of economic institutions and sustainable consumption via cultural group selection. *Ecological Economics*, 131, 524–532.
- Wasserstein, R. L. & Lazar, N. A. (2016). The ASA’s statement on p-values: Context, process, and purpose. *The American Statistician*, 70(2), 129–133.
- Weisberg, M. & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2), 225–252.
- Zollman, K. J. (2018). The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1), 5–33.