

# The Scientific Ponzi Scheme

Kevin J.S. Zollman

July 26, 2019

## Abstract

Fraud and misleading research represent serious impediments to scientific progress. We must uncover the causes of fraud in order to understand how science functions and in order to develop strategies for combating epistemically detrimental behavior. This paper investigates how the incentive to commit fraud is enhanced by the structure of the scientific reward system. Science is an “accumulation process:” success begets resources which begets more success. Through a simplified mathematical model, I argue that this cyclic relationship enhances the appeal of fraud and makes combating it extremely difficult.

The replication crises in various sciences illustrate how many published studies are misleading. While often lumped together, there are two distinct categories of failed replications. First are genuine good faith errors. Scientists may do everything by the book, but their results fail to replicate because the standards for publication are too loose or simply because in any inductive field errors are inevitable (Heesen 2018). A second, more nefarious, category are the result of shoddy research practices, statistical manipulation, or outright fraud. Scholars who produce this kind of research know they are doing something wrong, but are doing it anyway.<sup>1</sup>

Both phenomena are interesting, but they are different in important ways. I will focus on the later: intentionally bad research. Fraud and other forms of scientific misconduct are a serious problem (Fanelli 2009; John, Loewenstein, and Prelec 2012; Fang, Steen, and Casadevall 2012). Before it is discovered, it is relied on by those doing honest science, thereby disrupting scientific progress. When the fraud is substantial, it makes headlines

---

1. In this paper “fraud” refers to both intentional data fabrication and engaging in sloppy research. For my purposes differences between these two practices can be ignored. There is an interesting middle category where scientists engage in statistical manipulation without realizing that this is bad research practice. Some of what I say is applicable to this middle category, but space prevents a complete discussion.

undermining the public’s trust in science. As a result, we must understand the causes of fraud and how it can be discouraged, prevented, and caught early.

There are four broad methods for addressing fraud. First, some suggest that scientists should value truth more and career advancement less.<sup>2</sup> Second, one can attempt to improve the filtering mechanisms present in peer review (van ’t Veer and Giner-Sorolla 2016). Third is to tolerate fraud, but improve the scientific community’s strategies for discovering and addressing it (Nosek, Spies, and Motyl 2012; Nissen et al. 2016).<sup>3</sup> The final category attempts to alter the incentives that cause scientists to commit fraud. By increasing the chance they will be caught, by increasing the punishment for committing fraud, or by reducing the potential benefits from fraudulent research, we might change the (explicit or implicit) mental calculus that leads to fraud.

Each of these solutions has their place; I will focus on the later. Ultimately, I offer a skeptical concern: many structural features of science make it difficult to effectively disincentivize fraud. The problem is shown through a simple mathematical model which is highly idealized. I don’t aim to provide a faithful representation of all the incentives scientists face. Instead, the model illustrates a problem which has gone unrecognized until now. I conclude from the model that some apparent solutions to fraud will fail because it takes time for fraud to be discovered and “punished” by the social norms of science.

## 1 Brian Wansink

Research conducted by Brian Wansink and collaborators made headlines by suggesting that subtle ways we interact with food have profound implications for how we eat. In a now retracted study, Wansink argued that the size of the bowl in which food is placed affects the amount of food consumed. His research had high scholarly impact. He has been cited over 28,000 times and has an h-index of 78.<sup>4</sup> His research regularly made it into popular outlets, he

---

2. One might imagine that making scientists care more about truth would combat fraud. It turns out, this is not as obvious as it sounds (Bright 2017) and might have other negative consequences (Zollman 2018).

3. There is overlap between this category and the next one. If we catch fraud faster, this also reduces the propensity to engage in fraud because, now, it has less benefit in terms of one’s career advancement (Bruner 2013; Romero 2017).

4. This means he has 78 articles that have each been cited at least 78 times. This is according to Google Scholar as of April 9, 2019.

wrote two bestselling popular books, and was appointed to a White House panel.

In 2016, Wansink posted an anecdote on his blog about the graduate student who “never said no” (Wansink 2016). Apparently unaware of what he was saying, Wansink praised a student for engaging in p-hacking and HARKing.<sup>5</sup> The graduate student took a data set that did not yield a statistically significant result, and went looking for what other things might be statistically significant within the dataset. For those unfamiliar with classical statistical methodology, this is bad. Searching for significance in large datasets is likely to yield false positives which are the result of statistical noise.

This blog post attracted critical attention and led to the identification of issues in a large number of Wansink’s publications. A series of published and informal exchanges took place while concerns over Wansink’s work grew. Cornell University, where Wansink was a professor, started an investigation. They concluded that Wansink “committed academic misconduct in his research and scholarship, including misreporting of research data, problematic statistical techniques, failure to properly document and preserve research results, and inappropriate authorship” (Kotlikoff 2018). Many of his papers were retracted, including six in one day (Bauchner 2018).<sup>6</sup> As a result, Wansink tendered his resignation.

The identified problems with Wansink’s are substantial. Numbers reported in published papers were inconsistent or straightforwardly impossible (Zee, Anaya, and Brown 2017). He engaged in self-plagiarism and data-duplication (Zee 2017), and there remain many questions without complete answers.

Much remains a mystery about the situation. How much of what he was doing was intentionally fraudulent? If he knew the level of misconduct he was engaging in, why did he draw attention to it with that blog post? Did he think he was invincible or did he truly not understand the gravity of what he was praising? About much of this, we can only speculate.

Whatever the cause, this case is worrisome. Wansink’s studies were

---

5. “p-hacking” is a name for a collection of statistically illegitimate practices where one attempts to generate statistically significant results by performing alternative tests, searching a large number of variables (without disclosing that one has done so), etc. HARKing stands for “hypothesizing after the results are known” which is a practice of looking for interesting results in a dataset and pretending that this was one’s hypothesis all along.

6. A search of the RetractionWatch database on April 9, 2019 revealed 40 articles had been retracted, corrected, or had an official expression of concern attached. Of course, not all studies on which he is an author are suspect. Some involved data collected and analyzed by others for which there is little cause for concern.

influential. He became well known through shoddy or fraudulent research. That fame created a feedback effect, it generated more resources that allowed him to do even more shoddy research which in turn generated even more resources. If he hadn't accidentally revealed his research practices in a blog post, he may have continued undetected for years.

Obviously, many scientists and lay people came to hold false (or at least unjustified) beliefs. Other scientists tried to build on Wansink's conclusions. People changed their behavior in response to his books, and policy was based on it. Several scientists spent innumerable hours painstakingly unraveling his fraud, time that could have been spent on more productive scientific effort. Scientific progress was undoubtedly harmed by him, not to mention the public's trust in research.

In the sections that follow, I will develop a formal model of fraud and sloppy science. I cannot say whether it is a model of Wansink. No one – perhaps not even Wansink himself – can tell us what motivated him. But his case is instructive: he was able to build a scientific empire in part through his malfeasance that he might not have built with more honest research. This possibility is concerning, and I will argue endemic to the way we reward science.

## 2 Modeling fraud

Fraud like Wansink's is hard to detect. Initially it appears important and reliable. It is likely to get published in a high profile journal. Some credit accrues for the researcher, but eventually the fraud gets discovered or the study fails to replicate. The researcher gets punished or loses standing. This evolving dynamic of credit is what I model.

To start with a stylized model, consider figure 1. A researcher is choosing between two different plans: fraudulent or honest research. For both options there is uncertainty about how it will turn out. When she factors in the uncertainty she comes up with the following expectations. Fraudulent research will be faster, she can complete it in one unit of time. Honest research will take longer (two units of time). This seems reasonable; at the extreme, one can simply make up data which should not take much time at all.

The fraudulent project will make a bigger splash. It will get into a better journal and will initially get cited more. At its peak, her fraudulent research will generate more credit-per-time than the non-fraudulent research. Eventually, however, she'll get caught. Perhaps people won't ever discover the

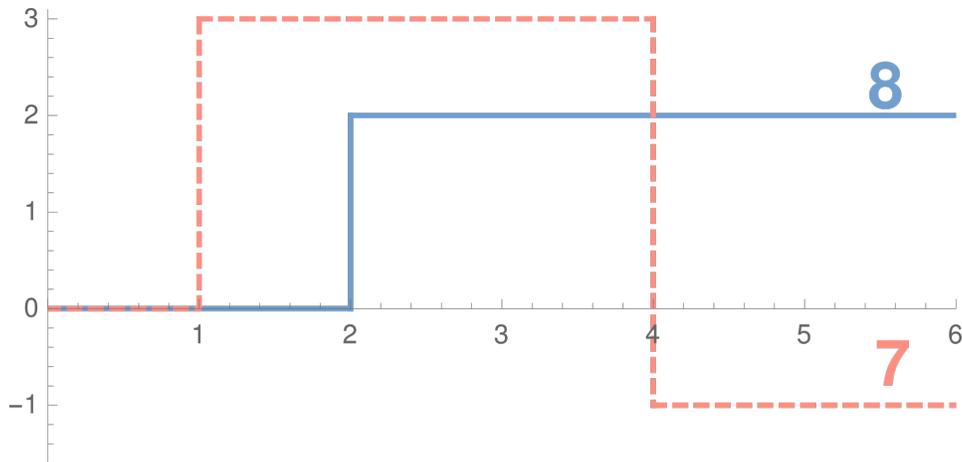


Figure 1: A simple illustration of the choice between fraudulent and non-fraudulent science. Time is on the  $x$ -axis and credit at a time is on the  $y$ -axis. The blue, solid line represents honest science and the red, dashed line represents fraudulent science. The total credit is written next to the relevant line.

outright fraud, but the result won't replicate. After some time, she will cease getting credit. We will even suppose that she gets "punished," the accumulated credit she receives is less than if she had done nothing at all.<sup>7</sup> That won't happen for the honest science. It will accumulate less credit at its peak, but it will continue to be cited for long into the future.

We will suppose a few things about how our scientist decides. First, she cares only for credit and nothing else. This is not true of actual scientists, of course. Where possible, we would prefer that scientists not be forced to choose between their career and their values, so we would like our scientific institutions to incentivize good science. The extreme case of a scientist motivated solely by the credit focuses our attention on the incentive system. Second, our scientist has a finite time horizon over which she evaluates the outcomes. Our scientist doesn't care about her legacy after she is dead. This may not be true of some scientists but might be true of others. (We will revisit this assumption.) Third, our scientist maximizes her expected credit. There are reasons to be skeptical that this accurately models real decision making, but it turns out that more realistic decision procedures will make

<sup>7</sup> There is evidence that scientists are "punished" for retractions (Lu et al. 2013; Stern et al. 2014), but for the argument that follows, this is not critical.

this paper’s central problem worse rather than better. This assumption will also be reevaluated once the conclusions are on the table.

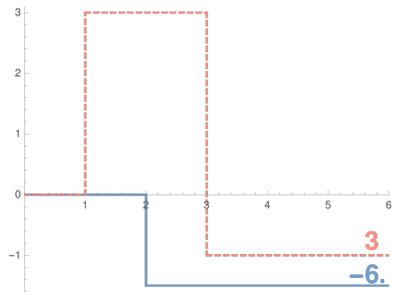
All of this is meant to be (a) incredibly stylized to demonstrate a point and (b) represent the expectation of what is a chancy process. Honest science often fails to replicate, and fraud is sometimes never discovered. This model is an illustration of a type of choice scientists face. In the example in figure 1, if the scientist cares about her accumulated credit for the first 8 units of time, she will choose honest science. The net credit is higher for honest science than for dishonest science.

Already in this simple model we have a few parameters: how much credit does one get, how long does one get it, how likely are we to discover fraudsters, and how bad is it to be discovered. By modifying these parameters, we can represent some of the ways that people are incentivized to engage in fraudulent science.

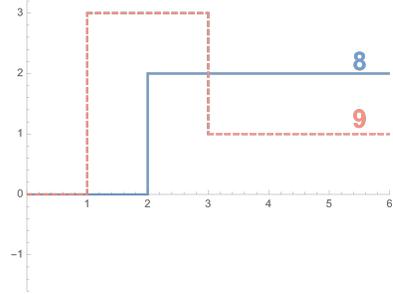
Figure 2 shows how we might represent several commonly discussed motivations for fraud. Briefly they are:

- **Desperation.** Scientists facing an imminent up-or-out decision like hiring or tenure might regard fraud as their only option. If the consequences are sufficiently dire, committing fraud and then being caught later might be better than honest research (DuBois et al. 2013).
- **Insufficient punishment.** The negative consequences are insufficient to make fraud unprofitable.
- **Unlikely to be caught.** Some types of fraud are essentially impossible to detect. In some fields, attempts to replicate are rare, and so a fraudulent study may go undetected (Romero 2017).
- **Noble lies.** The scientist believes that their understanding of the phenomena is correct and that their fraudulent results will replicate in honest studies (Bright 2017).
- **Rush to publish** Scientists engage in sloppy research to publish faster than a competitor (Heesen 2018).

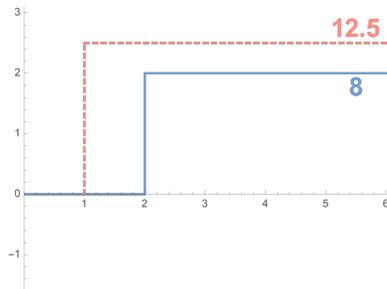
All of these examples compared two options (a) publish one fraudulent study or (b) publish one honest study. We will now turn our attention to a more complex decision that will occupy the remainder of our discussion. Imagine a scientist who is considering a career choice: should she habitually engage in fraud or honest science? Shifting our focus identifies another cause of fraud, it’s more efficient (Heesen 2018).



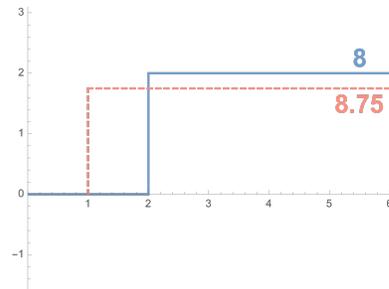
(a) Desperation



(b) Insufficient punishment for fraud



(c) Unlikely to be caught and noble lies



(d) Rush to publish

Figure 2: Illustrations of various motivations for fraud.

Suppose, for example, that a scientist could produce three sloppy studies in the time that it takes to produce two well designed ones. While it might not seem reasonable to choose a single fraudulent study over a single honest one (in terms of credit received), the added value of conducting three studies may render the choice of a fraudulent career superior.

### 3 The Scientific Ponzi Scheme

So far we’ve treated credit as a good that scientists consume. We’ve talked about it as like food: a scientist receives it, they enjoy it, and then it disappears. In the model thus far, it has no further downstream effects on the scientist’s career. But this is not how real credit works. Scientific esteem changes a scientist’s career. It determines what job she’s offered, when she gets promotion, and whether she receives research grants. This creates an important cyclical relationship between esteem and productivity. Higher productivity produces more esteem (*ceteris paribus*) which produces more resources which increases productivity which increases esteem further (Lator and Woolgar 1979; Grimes, Bauch, and Ioannidis 2018). Feedback mechanisms like this are often called cumulative advantage processes (Price 1976; Zuckerman 1998).

More formally, we can imagine the process in the following way.<sup>8</sup> First, let’s begin with an honest scientist. An honest scientist begins her career with a certain amount of resources  $r_0$ . Those resources determine a level of productivity  $p_0 = p(r_0)$ . The productivity translated into some amount of credit  $c_0 = c(p_0)$ . This credit is then turned into resources for the next period of scientific research  $r_1 = r(c_0)$ .

These three functions allow us to describe a dynamical system that determines a scientist’s lifetime productivity by specifying their accumulated productivity at a time:

$$p_{t+1} = p(r(c(p_t))) + p_t \tag{1}$$

To make life easy we will simply represent the  $p(r(c(\cdot)))$  function as  $e(\cdot)$ .  $e$  incorporates the compound affect of credit, resources, and productivity. That allows us to rewrite this difference equation as  $p_{t+1} = e(p_t) + p_t$

---

8. While developed independently, this model shares many of the same characteristics of the model presented in (Grimes, Bauch, and Ioannidis 2018). The focus on the analysis is different, and they do not explore how the cumulative advantage process can turn unprofitable fraud into profitable fraud – the central focus of this paper. Interested readers will find consonant results in their paper.

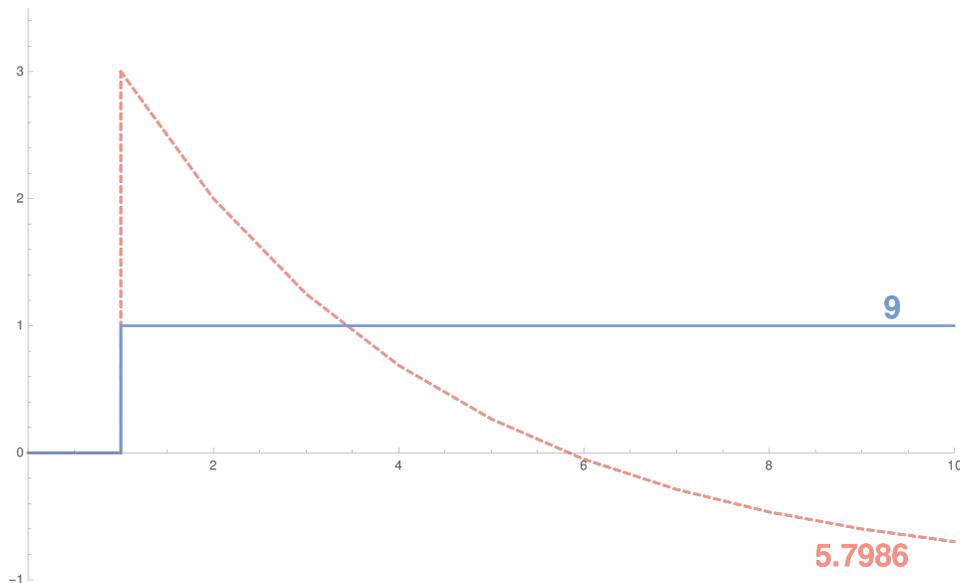


Figure 3: An illustration of a single shot decision where the scientists would prefer to honestly report their results.

We can include fraud in this model by breaking productivity into two parts:  $p^+$  is the positive aspects of a scientist's career achievements (those results which continue to influence a scientist's status positively) and the negative parts of a scientist's past productivity  $p^-$  (those results which have been discovered to be in error or fraudulent).

We can assume that there is a function  $d$  which determines the rate by which positive results are overturned. So, the change in credit is determined both by what new results are produced and how many old results are overturned. Credit is now a function of two arguments  $c(p^+, p^-)$ . Credit increases in the first argument but is reduced by the second.

The dynamical system is now defined by a slightly more complex equation:<sup>9</sup>

$$p_{t+1}^+ = e(p_t^+ - d(p_t^+), p_t^- + d(p_t^+)) + p_t^+ - d(p_t^+) \quad (2)$$

and

$$p_{t+1}^- = p_t^- + d(p_t^+) \quad (3)$$

---

9. As a technical matter, this later system can incorporate the former by setting  $d(\cdot) = 0$ .

To illustrate the central concern of this paper, consider the decision in figure 3. A scientist has conducted a study and the results are underwhelming. She can publish the results, but the paper will receive little attention. Instead of publishing her boring paper, she can alter the data to make the result more exciting. For example, she might p-hack the data, exclude certain data points, or add in fraudulent ones. Doing so will produce a big splash – generating three times the credit – but will also eventually be discovered. As time progresses, there is an increasing probability that her fraud will be discovered.<sup>10</sup> When it is eventually discovered, this will result in her receiving “negative” credit of -1.

In this decision, the credit-maximizing scientist would prefer to be honest. While she would receive an initial big splash from engaging in fraud, the probability of being discovered combined with the cost far outweighs the benefit. Considered as a one-shot decision, our scientist is properly incentivized to be honest.

Now, what happens if we embed this decision the accumulation process? Before her fraud is discovered, a fraudster can secure more resources which will allow her to engage in more fraud. This changes the calculation and results in the expected credit pictured in figure 4. Once we include the possibility of repeated fraud, underwritten by the larger resources that fraud allows one to accumulate, fraud becomes profitable.<sup>11</sup>

Our researcher is engaging in what I’m calling *the scientific Ponzi scheme*. Like more traditional Ponzi schemes, she is using ill-gotten gains to secure further “investment.” This will eventually come crashing down, but it may be rational nonetheless to engage in the process if one does not care sufficiently about the future.

Consider first the expectation lines (the red and blue lines) in figure 4(a). This shows how, when compounded over time, an individually unprofitable instance of fraud can become profitable. When considered as a single shot decision (as pictured in figure 3) we would conclude that fraud doesn’t pay. If we considered only that case, we might conclude that the social mechanisms in place would deter fraud. However, when considered from the perspective of an accumulation process, they are insufficient. A lifetime of fraud, in this

---

10. In this model, we assume that in each time period there is a 25% chance that the fraud is discovered. This probability is smoothly distributed over the unit of time.

11. We assume that the probability of a project being overturned is independent across time and across different projects. These assumptions might be false. If a scientist habitually commits fraud in the same way, the rate of fraud detection might increase. (Once I’ve figured out your trick, it’s easier to find in other publications.) Since this paper is a proof-of-possibility, I do not think this assumption is critical.

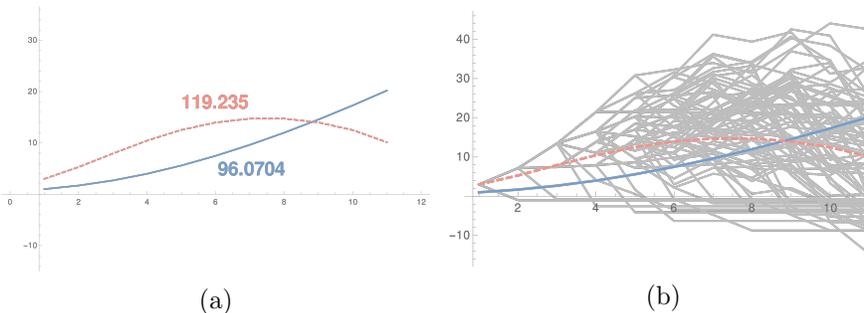


Figure 4: The scientific Ponzi scheme.  $d(p^+) = 0.25p^+$  and  $e(p^+, p^-) = \ln(3p^+ - p^- + 1)$  for fraudulent science. For honest science,  $d(p^+) = 0$  (results are never overturned) and  $e(p^+, p^-) = \ln(p^+ + 1)$ . The red dotted line is the expected payoff for fraud, the blue is the expected payoff for honesty. (a) contains only the expectation lines along with the total value of each strategy. The grey lines in (b) represent individual simulated trajectories through the space.

case, does pay.

Not all cases of unprofitable fraud can be so transformed. If the punishment for a single instance of discovered fraud is sufficiently high, then the Ponzi scheme won't pay either. Why don't we just ramp up the punishment for fraud? We could make the punishment so severe that even the Ponzi scheme doesn't make fraud viable. From a modeling perspective, this is certainly possible. However, there are concerns about how this might be implemented.

First, it is often difficult to distinguish unintentional errors from outright fraud. Even in a case so "cut-and-dry" as Wansink, we still can't eliminate the possibility that many of his errors came about through ignorance. Other cases are even more fraught. Second because scientists often have personal relationships with one another, harsh punishments might disincentivize publication of failed replication. If I knew that my study would destroy the career of a friend, I might opt to leave it in my file drawer. This would have a different negative epistemic consequence that is not being modeled here.

As a final point, consider the individual simulated trajectories pictured in figure 4(b). Looking at individual cases, some fraudulent scientists will appear to be superstars, even after a substantial amount of time has passed. Even in cases where fraud doesn't pay on average, there will be some instances where fraudulent scientists do very well because they've gotten

lucky.<sup>12</sup>

In some cases, this high variance feature of fraud might be attractive. If one’s expectation from engaging in honest research is low, one might opt for the high-variance but lower-expectation strategy in order to maximize one’s chances to clear a hurdle like tenure (Tsetlin, Gaba, and Winkler 2004).<sup>13</sup> This possibility is a central focus of (Grimes, Bauch, and Ioannidis 2018).

## 4 Revisiting the assumptions

In illustrating the scientific Ponzi scheme, we made several assumptions which might not be true of real scientists.

We assumed that scientists care about all time periods equally. Both traditional economics models and behavioral economics suggest that people value current consumption over future consumption (Frederick, Loewenstein, and O’donoghue 2002). Because fraud pays more at the outset, adding in discounting or so-called “present bias” would make fraud *more* profitable.

There is a second behavioral phenomenon that points in the other direction. People also appear to have a preference for “increasing” sequences of outcomes.<sup>14</sup> This phenomena pushes in the other direction since honest science is always increasing in credit and fraudulent science is not. However, preference for increasing sequences is not overriding, and so it seems likely that we could still construct examples if there was some preference for increasing sequences of credit.

A second assumption is that scientists stop caring about credit after a certain period. If scientists continued to accumulate credit long after they die, and scientists don’t discount the future very much at all, the desire to engage in the Ponzi scheme might vanish. It is unclear how much actual scientists care about their legacy relative to current reputation. I am not aware of any significant research into this question. In some fields, the speed of progress is so fast that very few scientist have a prospect for much of a legacy – only those who rise to the level of major awards like the Nobel prize can expect to be remembered after their death. In other fields, legacy might

---

12. Although produced through a completely different mechanism, this possibility is consonant with a process whereby scientists can become famous principally through luck discussed in (Heesen 2016).

13. This can hold even if we assume that agents are trying to maximize expected utility, because thresholds can create apparent (but not actual) non-expected utility maximizing behavior.

14. See for example (Loewenstein and Sicherman 2002), but there are some complications (Frederick and Loewenstein 2008).

be more of a serious possibility.

Finally, there is substantial literature that questions the degree to which individuals evaluate uncertain outcomes according to their mathematical expectation. It is not at all clear how this modification would affect these models. Both fraud and honest research is fraught with uncertainty, and it is not clear whether modifying our scientist to be more realistic would make one more appealing than the other. Nor is it at all clear how we would do so, since the uncertainty in science spans the different categories studied in behavioral economics. As a result, I see no reason that this would introduce reasons to reject the conclusion here.

## 5 Conclusion

One broad strategy for combating fraud and sloppiness in science would be to alter the social incentives to discourage such behavior. Doing so, one might hope, would alter a scientist's decision calculus in a way that reduces fraud over all. Through a proof-of-possibility model, this paper argues that this strategy is more complicated than it appears. The scientific Ponzi scheme, would allow for scientists to combine discentivized fraudulent behavior in a way that would be, all-things-considered, profitable for the scientist in terms of career advancement. Because science is an accumulation process, where early success results in more resources for later science, early fraud might show larger returns than would be suggested by looking at this as a "single-shot" decision.

Some have suggested that science should move from a proposal and grant system to a prize-for-past-work system (Charlton and Andras 2008). While not conclusive, I think this paper suggests a problem with that approach. If the prizes are too proximate to recent work, it may encourage the Ponzi scheme rather than combat it. Beyond prizes, (Grimes, Bauch, and Ioannidis 2018) show, in a very similar model, that more competitive grant races attenuate this problem by rewarding fewer researchers.

Since this is merely a proof of possibility, it does not eliminate the incentive-based strategy for combating fraud. It does suggest that this strategy is more difficult than one might hope. Any approach we take is unlikely to be completely effective: some scientists will be incentivized to commit fraud and some will do so. It will always be critical that the institutions of science be robust, and that we understand that fraud will occur even in "healthy" scientific communities.

## References

- Bauchner, Howard. 2018. “Notice of Retraction: Wansink B, Cheney MM. Super Bowls: Serving Bowl Size and Food Consumption.” *JAMA* 320 (16): 1648. ISSN: 0098-7484. doi:10.1001/jama.2018.14249.
- Bright, Liam Kofi. 2017. “On Fraud.” *Philosophical Studies* 174 (2): 291–310.
- Bruner, Justin P. 2013. “Policing Epistemic Communities.” *Episteme* 10 (04): 403–416. ISSN: 1742-3600. doi:10.1017/epi.2013.34.
- Charlton, Bruce G., and Peter Andras. 2008. “Stimulating revolutionary science with mega-cash prizes.” *Medical Hypotheses* 70 (4): 709–713. ISSN: 03069877. doi:10.1016/j.mehy.2008.01.001.
- DuBois, James M., Emily E. Anderson, John Chibnall, Kelly Carroll, Tyler Gibb, Chiji Ogbuka, and Timothy Rubbelke. 2013. “Understanding Research Misconduct: A Comparative Analysis of 120 Cases of Professional Wrongdoing.” *Accountability in Research* 20 (5-6): 320–338. ISSN: 08989621. doi:10.1080/08989621.2013.822248.
- Fanelli, Daniele. 2009. “How many scientists fabricate and falsify research? A systematic review and meta-analysis of survey data.” *PLoS ONE* 4 (5). ISSN: 19326203. doi:10.1371/journal.pone.0005738.
- Fang, Ferric C., R. Grant Steen, and Arturo Casadevall. 2012. “Misconduct accounts for the majority of retracted scientific publications.” *Proceedings of the National Academy of Sciences* 109 (42): 17028–17033. ISSN: 0027-8424. doi:10.1073/pnas.1220833110.
- Frederick, Shane, and George Loewenstein. 2008. “Conflicting motives in evaluations of sequences.” *Journal of Risk and Uncertainty* 37 (2-3): 221–235. ISSN: 08955646. doi:10.1007/s11166-008-9051-z.
- Frederick, Shane, George F Loewenstein, and Ted O’donoghue. 2002. “Time Discounting and Time Preference: A Critical Review.” *Journal of Economic Literature* 40 (2): 351–401. ISSN: 0022-0515. doi:10.1257/002205102320161311.
- Grimes, David Robert, Chris T Bauch, and John P A Ioannidis. 2018. “Modelling science trustworthiness under publish or perish pressure.” *Royal Society Open Science* 5:171511.

- Heesen, Remco. 2016. “Academic superstars: competent or lucky?” *Synthese*, no. September 2015: 1–20. ISSN: 15730964. doi:10.1007/s11229-016-1146-5.
- Heesen, Remco. 2018. “Why the Reward Structure of Science Makes Reproducibility Problems Inevitable.” *Journal of Philosophy* 115 (12): 661–674.
- John, Leslie K., George F Loewenstein, and Drazen Prelec. 2012. “Measuring the Prevalence of Questionable Research Practices With Incentives for Truth Telling.” *Psychological Science* 23 (5): 524–532. ISSN: 14679280. doi:10.1177/0956797611430953.
- Kotlikoff, Michael I. 2018. *Statement of Cornell University Provost Michael I. Kotlikoff*. Accessed April 9, 2019. <https://statements.cornell.edu/2018/20180920-statement-provost-michael-kotlikoff.cfm>.
- Latour, Bruno, and Steve Woolgar. 1979. *Laboratory Life: The Construction of Scientific Facts*. Beverly Hills: Sage Publications.
- Loewenstein, George, and Nachum Sicherman. 2002. “Do Workers Prefer Increasing Wage Profiles?” *Journal of Labor Economics* 9 (1): 67–84. ISSN: 0734-306X. doi:10.1086/298259.
- Lu, Susan Feng, Ginger Zhe Jin, Brian Uzzi, and Benjamin Jones. 2013. “The Retraction Penalty: Evidence from the Web of Science.” *Scientific Reports* 3 (1): 3146. ISSN: 2045-2322. doi:10.1038/srep03146.
- Nissen, Silas B., Tali Magidson, Kevin Gross, and Carl T. Bergstrom. 2016. “Publication bias and the canonization of false facts.” *eLife* 5:1–19. doi:10.7554/eLife.21451. arXiv: 1609.00494.
- Nosek, Brian A., Jeffrey R. Spies, and Matt Motyl. 2012. “Scientific Utopia: II. Restructuring Incentives and Practices to Promote Truth Over Publishability.” *Perspectives on Psychological Science* 7 (6): 615–631. ISSN: 1745-6916. doi:10.1177/1745691612459058.
- Price, Derek De Solla. 1976. “A general theory of bibliometric and other cumulative advantage processes.” *Journal of the American Society for Information Science* 27 (5): 292–306. ISSN: 10974571. doi:10.1002/asi.4630270505. arXiv: arXiv:1011.1669v3.
- Romero, Felipe. 2017. “Novelty vs Replicability : Virtues and Vices in the Reward System of Science.” *Philosophy of Science* 84 (5): 1–14.

- Stern, Andrew M, Arturo Casadevall, R Grant Steen, and Ferric C Fang. 2014. "Financial costs and personal consequences of research misconduct resulting in retracted publications." *eLife* 3:1–10. doi:10.7554/eliflife.02956.
- Tsetlin, Ilia, Anil Gaba, and Robert L. Winkler. 2004. "Strategic choice of variability in multi-round contests and contests with handicaps." *Journal of Risk and Uncertainty* 29 (2): 143–158. ISSN: 08955646. doi:10.1023/B:RISK.0000038941.44379.82.
- van 't Veer, Anna Elisabeth, and Roger Giner-Sorolla. 2016. "Pre-registration in social psychology: A discussion and suggested template." *Journal of Experimental Social Psychology* 67:2–12. ISSN: 10960465. doi:10.1016/j.jesp.2016.03.004.
- Wansink, Brian. 2016. *The Graduate Student Who Never Said "No"*. Accessed April 9, 2019. <https://web.archive.org/web/20170312041524/http://www.brianwansink.com/phd-advice/the-grad-student-who-never-said-no>.
- Zee, Tim van der. 2017. *The Wansink Dossier: An Overview*. Accessed April 9, 2019. <http://www.timvanderzee.com/the-wansink-dossier-an-overview/>.
- Zee, Tim van der, Jordan Anaya, and Nicholas J. L. Brown. 2017. "Statistical heartburn: an attempt to digest four pizza publications from the Cornell Food and Brand Lab." *BMC Nutrition* 3 (54).
- Zollman, Kevin J.S. 2018. "The credit economy and the economic rationality of science." *Journal of Philosophy* 115 (1).
- Zuckerman, Harriet. 1998. "Accumulation of advantage and disadvantage: The theory and its intellectual biography." In *Robert K. Merton and contemporary sociology*, edited by Carlo Mongardini and Simonetta Tabonni. New Brunswick: Transaction Publishers.