

# Questionable Research Practices and Credit in Academic Careers\*

Remco Heesen<sup>†‡</sup>

November 16, 2020

## Abstract

This paper investigates how the credit incentive to engage in questionable research practices (up to and including fraud) interacts with cumulative advantage, the process whereby high-status academics more easily increase their status than low-status academics. I use a mathematical model to highlight two dynamics that have not yet received much attention. First, due to cumulative advantage, questionable research practices may pay off over the course of an academic career even if they do not appear attractive at the level of individual publications. Second, because of the role of bottleneck moments in academic careers, questionable research practices may be selected for even if they do not provide a benefit in expectation. I also observe that, within the model, the most successful academics are the most likely to have benefited from fraud.

---

\*Thanks to Liam Bright for valuable discussion. This work was supported by the Netherlands Organisation for Scientific Research (NWO) under grant 016.Veni.195.141.

<sup>†</sup>Department of Philosophy, School of Humanities, University of Western Australia, Crawley, WA 6009, Australia. Email: [remco.heesen@uwa.edu.au](mailto:remco.heesen@uwa.edu.au).

<sup>‡</sup>Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, the Netherlands.

# 1 Introduction

Trust in academic science consists at least partially in trust in academics. It is a cause of concern, then, when that trust appears to have been misplaced, as happens when cases of fraud are revealed.

Cases of data fabrication, plagiarism, and other forms of outright fraud attract a lot of attention when they are uncovered, but these are perceived by many observers as being quite rare. In contrast, so-called questionable research practices (e.g., p-hacking, salami publishing) are perceived as less bad but also much more widespread.

Here I focus on what fraud and questionable research practices have in common. Practices that fall under either of these labels are ways of enhancing an individual academic's productivity and prestige (which mutually reinforce one another in a process known as cumulative advantage) at some epistemic cost. As a result, the same sorts of reasons may attract academics to either of them.

This paper investigates these reasons. I consider the consequences for an academic's productivity and her career of engaging in these epistemically costly practices, focusing on the role of cumulative advantage and the trade-off between short-term benefits in terms of productivity and prestige and long-term costs. I highlight two dynamics in particular. First, even when the chance of being caught and associated penalties are sufficiently high that epistemically dubious research outputs do not individually confer a net benefit, academics may gain a career advantage from them. Second, even when these practices are not rewarded on average, they may spread in academic communities. I will provide some tentative reasons to think that the most successful members of an academic community are relatively more likely to have benefited from fraud or questionable research practices.

The paper proceeds as follows. Section 2 provides a more careful introduction to the key concepts of cumulative advantage, the credit economy, fraud, and questionable research practice. In section 3 I develop a simplified

model of cumulative advantage in which an academic's productivity is represented as a non-homogeneous Poisson process. Section 4 adds a key downside of questionable research practices to the model, namely the possibility of being exposed. Section 5 shows how the two dynamics mentioned above may arise in the model, while section 6 considers more systematically when this happens by varying the parameters of the model. Section 7 concludes.

## 2 Cumulative Advantage and Fraud

The academic world is strongly hierarchical (Cole and Cole 1973). There is a small group of professors who seem to have it all: a chair at a prestigious university, plenty of time and money for research, lots of graduate students, publications that appear in highly regarded journals and are frequently cited, prizes, media appearances, and so on. In contrast there is a much larger group with few or none of these status markers, ranging from lesser-known but tenured or tenurable professors to the large group of academics without secure employment (including postdocs, adjunct professors, lab technicians, and graduate students). These differences in status are often keenly felt, as revealed for example in the phenomenon of *prestige bias* in hiring or publishing, where status markers such as one's individual reputation, institutional affiliation, or publishing track record influence one's chances of being hired (Clauset et al. 2015, de Cruz 2018) or navigating a paper through peer review (Tomkins et al. 2017, Lee et al. 2013, p. 7).

Since the institution where one is hired and the journals one publishes in are themselves status markers, prestige bias become a self-reinforcing effect. Those who manage to obtain some of these markers of prestige, especially early in their academic career, will then have an easier time being hired into a prestigious job, acquiring research grants, and more generally increasing their status. In contrast, those who struggle a bit more at the start of their career and fall behind in the prestige hierarchy will find it that much more

difficult to catch up. This general pattern, where early success begets more success, is known as *cumulative advantage* (DiPrete and Eirich 2006). In the academic context, it is known as the *Matthew effect* (Merton 1968).

Given the central importance of prestige, one might wonder what academics might be willing to do to acquire it. This amounts to asking about the *incentive structure of academic science*, also known as the *credit economy*. Academics receive credit first and foremost for making (and publishing) academic contributions, with originality being particularly prized (Merton 1957, Strevens 2003). As indicated above, this form of credit (or prestige; I will use these as synonyms) interacts in a mutually reinforcing way with other forms, such as citations, prizes, and prestigious appointments and grants.

Of the various forms of credit, the only one that academics have significant individual control over is the production of academic contributions and the submission of papers describing these contributions to prestigious journals. Given the importance of credit in establishing one's place in the academic hierarchy, this leads to an intense pressure to produce and publish research output, the so-called *publish or perish* culture (Fanelli 2010, Brischoux and Angelier 2015).

Academics facing this pressure might look to take shortcuts to increase their productivity. I will use the term *questionable research practices* (QRPs) for such shortcuts. For my purposes here I will regard all of the following behaviors as QRPs (some of which aim to improve productivity directly, while others aim to increase the impact of publications). Fabricating data and other forms of outright fraud (Bright 2017). Using multiple model specifications but only reporting those in which a result is statistically significant, i.e., p-hacking (Simmons et al. 2011). Hypothesizing after the results are known (Kerr 1998). Distributing findings from a single study over multiple papers, i.e., salami publishing (Abraham 2000). General sloppiness due to the desire to complete projects quickly, i.e., rushing into print (Heesen 2018). Being named as author on work where one has made no substantial intellectual

contribution, i.e., honorary coauthorship (Flanagin et al. 1998).

I do not mean to suggest that all QRPs are equally bad. Most would agree that outright fraud is the worst one, and for some of the others the jury is still out on whether they should be regarded as bad at all. But when it comes to the incentive to engage in them provided by the credit economy, these QRPs may be treated equally. That is, whenever this paper identifies scenarios in which academics have an incentive to engage in QRPs, this applies to any and all of the foregoing behaviors. In order to emphasize the worst-case outcomes that my argument supports, I will at times summarize my findings in terms of an incentive to commit fraud.

### **3 Modeling Cumulative Advantage**

I will now construct a relatively simple model of the credit economy with cumulative advantage built in at its core. I take the simplified nature of the model to be a strength rather than a weakness, as it helps to focus the attention on a small number of features of the credit economy and their consequences. Combining this with what we know empirically about academics' incentives and behaviors may yield a reasonable degree of confidence that the patterns of incentives identified here also operate in the real world.

The model provided here differs in a number of ways from other models of the credit economy in the literature. First of all, it is explicitly dynamic. In this respect it differs from early models of the credit economy, which were static in nature (Kitcher 1990, Dasgupta and David 1994, Strevens 2003, Zollman 2018), although by now plenty of dynamic models exist as well (Smaldino and McElreath 2016, O'Connor and Bruner 2019, O'Connor 2019, Zollman 2019). Second, it uses continuous time rather than discrete time units. Where the issue has come up at all, previous models have tended to use discrete time (e.g., Boyer 2014, Zollman 2019). Third, rather than assuming that academics maximize expected credit as is commonly done, I

take academics' aim to be to satisfy relative to particular credit thresholds. I will motivate this modeling choice below. To my knowledge, this makes my model unique among those that look at academic incentives in a rational choice model (as opposed to an evolutionary model, where an analogous move has been made by Smaldino and McElreath 2016, O'Connor 2019).

As already noted, *publications* play a central role in the academic world and in the way credit is distributed. The basis of the present model is a stochastic counting process that keeps track of the publications of a given academic over a period of time. The thought here is that an academic's productivity (both how many publications are produced and how they are distributed over time) has both a random component and a systematic component. The random component stands in for all factors affecting productivity that are not explicitly modeled, such as extraneous circumstances in the academic's life, the difficulty of the particular scientific problem she is working on, etc. The systematic component consists of the academic's talent and skill, as well as the amount of time and resources she has available. This includes in particular the cumulative advantage effect: a scientist who has already been productive is more likely to be given time and resources that help her become even more productive.

To model the random component I use a *Poisson process* (see any textbook on stochastic models, e.g., Tijms 2003, Norris 1998). In a Poisson process the time between two publications is assumed to be an exponentially distributed random variable. Moreover, the time between any two publications is probabilistically independent of anything that happened in the process before the first of these two publications, i.e., the 'interarrival' times of publications are independent and identically distributed. Under these assumptions, the total number of publications over a given time interval follows a Poisson distribution.

Why do I model the random component in this way? The Poisson process has the following important feature: looking backwards the publications

generated by it will appear to be randomly scattered in time. More precisely, if a Poisson process produced a specific number of publications over a given time interval, without more specific information the conditional probability distribution for when each of these publications arrived is the uniform distribution (Tijms 2003, theorem 1.1.5). So there is a precise sense in which this is a truly random model. As Norris (1998, p. 73) puts it, “a Poisson process is the natural probabilistic model for any uncoordinated stream of discrete events in continuous time”. That said, this feature will no longer hold once I add the systematic component to the model.

A Poisson process has one parameter, usually denoted  $\lambda$ . It is interpreted as the rate of publication, i.e., the expected number of publications per unit time. This can be used to add a systematic component to the model. For example, the publication output of two academics might be modeled using two Poisson processes with parameters  $\lambda_1$  and  $\lambda_2$ , with  $\lambda_1 > \lambda_2$  to indicate that the first academic has more time and resources and so is expected to be more productive.

Using the parameter  $\lambda$  in this way allows me to model persistent productivity differences between academics (cf. Heesen 2017b). However, I also want to capture cumulative advantage, i.e., the effect of earlier publication output on later productivity. This requires the systematic component to vary dynamically and endogenously. For this purpose I use a *non-homogeneous Poisson process* (also known as a non-stationary Poisson process, see Tijms 2003, section 1.3), which is like a regular Poisson process except that the rate of publication  $\lambda(t)$  is a function of time.

To get the cumulative advantage effect going, I will assume that publications generate credit over time (it takes time for a new publication to have its influence as word spreads, other academics start citing it, etc.). The credit accumulated by an academic may be turned into time and resources for research (Latour and Woolgar 1986, chapter 5). This is a complex and multi-faceted process involving, among other things, higher chances of win-

ning research grants, a job at a more prestigious university, or more graduate students. Rather than attempt to model this explicitly, I assume that credit buys time and resources directly.

Suppose that each publication generates  $c_g$  units of credit per unit time (the subscript  $g$  stands for ‘good’; a notion of ‘bad’ credit will be introduced below). For an academic with  $p_g$  publications this means she accumulates credit at a rate of  $c_g \cdot p_g$  per unit time. I will use  $c(t)$  to denote the credit accumulation rate at a given time  $t$  and  $C(t)$  for the total amount of credit accumulated, i.e.,

$$c(t) = c_g \cdot p_g \quad \text{and} \quad C(t) = \int_0^t c(u) du. \quad (1)$$

By turning this into time and resources the academic increases her rate of publication. I will use the following formula to capture this effect:

$$\frac{d}{dt} \lambda(t) = \log(c(t) + 1). \quad (2)$$

That is, the rate of publication increases (over time) proportionally to the logarithm of the credit rate (the +1 makes sure that if  $c(t) = 0$  then  $\frac{d}{dt} \lambda(t) = 0$ ). The logarithm reflects a type of decreasing marginal returns: if the credit accumulation rate is already high then the effect (on the publication rate) of increasing it further is smaller than if the credit accumulation rate is low. The underlying idea is that large amounts of credit are harder to effectively turn into resources than moderate amounts.

Note that the credit rate  $c(t)$  is a step function: it is constant between publications and then jumps instantaneously from  $c_g \cdot p_g$  to  $c_g(p_g + 1)$  when a new publication appears. As a result the rate of publication changes linearly between publications: if  $T_0$  and  $T_1$  are consecutive publication arrival times then for  $T_0 < t < T_1$  the rate of publication is

$$\lambda(t) = \lambda(T_0) + \log(c(t) + 1)(t - T_0). \quad (3)$$



An example of how this may develop is shown in figure 1. The rate of pub-




Figure 1: The rate of publication increases faster as more publications (indicated by dots) come in (parameters:  $c_g = 1$  and  $\lambda(0) = 1$ ).

lication stays flat until the first publication comes in just before  $t = 1$ , then starts increasing in steepness as more publications arrive. Meanwhile, publications arrive closer and closer together as the rate of publication increases. Thus publications and the rate of publication mutually reinforce one another, producing the cumulative advantage effect.

In this example the academic accumulates  $C(5) = 22.74$  units of credit (one unit per publication from the moment of publication until the end of the simulation at  $t = 5$ ). This may vary due to the stochastic nature of the process, but repeated simulation runs give a sense for the typical outcomes. In this case average credit at  $t = 5$  is  $\mathbb{E}[C(5)] \approx 27.93$ . There is significant variation though: in my 10,000 runs the minimum credit was 0, the

maximum 112.77, and the standard deviation 15.74.

## 4 Replications, Exposures, and Negative Credit

The model described above assumes that once a publication occurs it generates credit indefinitely. There are a number of respects in which this may be unrealistic. First, the impact of most publications fades over time. Since my aim is to model relatively short intervals of time, e.g., from being hired to going up for tenure, or from starting graduate school to going on the job market, I will ignore this factor (though the model I present here is an instance of a more general model known as a Hawkes process or self-exciting process which can incorporate this factor). Second, and more immediately relevant, fraudulent or shoddy work may be exposed, and even research of the highest standard may fail to replicate. As recent studies have shown, significant proportions of published results in various empirical sciences fail to replicate (Open Science Collaboration 2015, Nosek and Errington 2017, Camerer et al. 2018). When this happens, it changes how the original work and its author(s) are perceived, i.e., it changes the credit associated with that publication.

To incorporate this in the model, assume that for each publication there is a chance of it being ‘exposed’. This may mean a failure to replicate, a discovery that the work was fraudulent, or any other event with significant negative impact on the perception of the work. In particular, I introduce a new parameter, the publication exposure rate  $\mu$ , and assume that for each publication the time between it being published and it being exposed is exponentially distributed with rate  $\mu$ . Note that the probability of a publication never being exposed equals the probability that this exponential distribution fails to trigger in the time window under consideration. As a result of these assumptions the stochastic process that counts exposure events for a given academic is a non-homogeneous Poisson process, with a total exposure rate

at any given time of  $p_g \cdot \mu$  (the number of publications available to be exposed times the publication exposure rate).

Once a publication is exposed it is removed from the set of publications available to be exposed ( $p_g$  is reduced by one) and stops generating  $c_g$  units of credit per unit time. Instead it is added to the set of exposed ('bad') publications ( $p_b$  is increased by one) and starts losing  $c_b$  units of credit per unit time. The idea is that as news of the exposure spreads through the academic community, credit is taken away from the academic whose paper has been exposed, adjusting the total amount of credit generated by this publication downwards. The function  $c$  that keeps track of the credit accumulation rate is adjusted to reflect this:

$$c(t) = c_g \cdot p_g - c_b \cdot p_b. \quad (4)$$

Unlike before, the credit accumulation rate may now decrease or even become negative. This requires adjusting the formula for changes in the rate of publication as well, since the logarithm is not defined for negative numbers:

$$\frac{d}{dt} \lambda(t) = \begin{cases} \log(c(t) + 1) & \text{if } c(t) \geq 0, \\ c(t) & \text{otherwise.} \end{cases} \quad (5)$$

Figure 2 illustrates this increased range of possibilities. Shown is the development of the rate of publication over time for two simulation runs. In the former the occasional exposure can be seen to slow down the rate of publication, but on the whole publications come in fast enough so the rate of publication continues to increase. In the latter there are fewer publications and more exposures, with the publication rate eventually dropping down to zero. In the former the academic accumulates  $C(5) = 19.21$  units of credit, whereas in the latter the academic ends up with negative credit ( $C(5) = -2.11$ ). With these parameters the former outcome is more typical: after 10,000 runs average credit is  $\mathbb{E}[C(5)] \approx 14.65$  (standard deviation 10.50),




Figure 2: Two simulation runs (one in which publications outrun exposures and one in which they do not) with parameters  $c_g = 1$ ,  $c_b = 1/2$ ,  $\lambda(0) = 1$ , and  $\mu = 1/4$ . Publication events are marked as open dots and exposure events as closed dots.

with only 6.16% of runs accumulating zero credit or less.

## 5 What are Academics' Incentives?

So far, questionable research practices have not taken center stage. As mentioned, such practices allow the academic to work more quickly and will lead to higher impact publications as the academic is able to achieve flashier, more surprising, or more newsworthy results. Under what circumstances might academics have an incentive to engage in QRPs?

Suppose an early-career academic faces the choice between either engag-

ing in fraud or QRPs or not doing so. Regular (non-fraudulent or ‘honest’) publications tend to yield credit at a rate of  $c_g^* = 1$  whereas publications obtained using fraud or QRPs tend to yield more:  $c_g^\dagger = 1.25$ . This captures the fact that the latter tend to have higher impact. Indirectly, it also captures the fact that QRPs allow the academic to work more quickly, as the higher credit obtained is converted into an increased rate of publication as described above.

The use of QRPs comes at a cost, however. Publications acquired in this way are more likely to be exposed, i.e.,  $\mu$  will be higher. Once they are exposed, they are also punished more harshly as it will be recognized not only that the published result was wrong, but that bad methods were used to obtain it, i.e.,  $c_b$  will be higher. There are important nuances here: sometimes academics are falsely accused of fraud (e.g., Fisher’s accusations against Mendel, or Newton’s accusations against Leibniz) and sometimes fraudulent work is recognized as irreproducible but not as fraud. Moreover, the increasing prevalence of large collaborations makes it harder both to detect fraud and to adjust individual credit in response (Huebner and Bright 2020, p. 364, see also Wray 2017, p. 129). Still, an academic is more likely to be exposed as fraudulent if they are in fact fraudulent, so the assumption that  $c_b$  will be higher if QRPs are used seems apt. To capture this with (fairly arbitrary) numbers, suppose  $\mu^* = 1/6$  whereas  $\mu^\dagger = 1/4$ , and  $c_b^* = 0$  whereas  $c_b^\dagger = 1/2$ .

These choices of parameter values look like they favor honest academic work over using QRPs: a relatively modest (25%) increase in the credit gained from publications would seem to be more than offset by the increased exposure rate and the negative credit associated with exposed publications. To substantiate this, I have worked out the credit that is expected to accrue to a single publication up to a given time  $t_1$  (I will continue to use  $t_1 = 5$  in all examples). For an academic with only a single publication at time  $t_0$  (i.e., the first publication arrives at  $T_0 = t_0$  and the second publication arrival

time  $T_1$  is after  $t_1$ ) it is not hard to show that

$$\mathbb{E}[C(t_1) \mid T_0 = t_0, T_1 > t_1] = \frac{c_g + c_b}{\mu} (1 - e^{-\mu(t_1 - t_0)}) - c_b(t_1 - t_0). \quad (6)$$

With the parameters as chosen above and a publication at time  $t_0 = 0$  the honest academic expects to get more credit than the fraudulent one:

$$\mathbb{E}[C^*(5) \mid T_0 = 0, T_1 > 5] \approx 3.39 > 2.49 \approx \mathbb{E}[C^\dagger(5) \mid T_0 = 0, T_1 > 5].$$

Similarly, if the publication arrives at  $t_0 = 1$  the honest academic expects 2.92 units of credit from it and the fraudulent one merely 2.42. So from the perspective of expected credit per publication it appears to be better to be honest (although this result eventually flips when the publication arrives closer to  $t_1$  as such publications are less likely to be exposed in the remaining time; on the other hand such publications contribute less overall as they are around less long).

One reason why looking at the expected credit for a single publication is misleading, however, is that this credit is not evenly distributed over time. Whereas the honest academic expects a steady stream of credit that lasts for a while, the fraudulent academic expects a bigger stream of credit initially, but also for a shorter period, followed by a period of negative credit. While the fraudulent academic ends up with less credit at time  $t_1$ , there is an initial period where she gets more than the honest academic, and since she can use this early credit to increase her rate of publication she might be able to offset the later negative effects by producing (many) more publications overall. In this sense, the pattern of credit accumulation by the fraudulent academic resembles a Ponzi scheme (Zollman 2019).

So by taking into account the cumulative advantage effect, it might be that QRPs pay off when the expected credit of all publications is considered, even though honesty is better from the perspective of credit per publication. Once again I estimate expected credit by simulating the process 10,000 times.

In the running example this yields

$$\mathbb{E}[C^*(5)] \approx 21.52 \quad \text{and} \quad \mathbb{E}[C^\dagger(5)] \approx 21.28.$$

So despite the cumulative advantage effect, with these parameters it is still slightly better to be an honest academic from the perspective of maximizing expected credit.

In models like this one it is typically assumed that academics' goal is to maximize expected credit. This makes a certain amount of sense, given the close analogy between credit (in motivating academics) and utility (in motivating arbitrary rational agents), and the role of expected utility theory as the conventional model of rational choice. But for most academics it will arguably be more important to meet specific credit thresholds. The competitive aspects of academic life are felt most keenly at a few pivotal moments in an academic's career, such as when she is on the job market or going up for tenure.

On the job market, the credit the academic has accumulated is likely to play an important role in her prospects. At such a time, what matters is whether the academic has accumulated enough credit: enough to be competitive, enough to land that dream job, enough to achieve whatever goal the academic has set for herself. Simplifying significantly (and ignoring many other factors such as teaching competence, personality, or how fashionable her research area is), the academic's goal may be formulated as a target amount of credit, such that meeting or exceeding this target is considered success, and falling short failure.

Going up for tenure is similar. This process is typically not competitive. Rather, the academic is given a set of (possibly vague) criteria she is expected to meet. At least as far as the research component of these criteria goes, the most important factor, if not the only one, will be the reputation the academic has built based on her publications, i.e., her accumulated credit. So in order to get tenure she needs to meet or exceed a certain credit threshold.

If getting a job or tenure is far more important to the academic than any other aspects of her career that depend on credit, then it would be rational for such an academic to aim to maximize the probability of meeting whatever credit threshold is relevant for her. Moreover, note that the job market and the tenure process are important points at which it is decided who stays in academia and who leaves (other points may be just as important in determining one's place in the academic hierarchy, but having any such place at all is only possible if one stays in academia). If those who meet credit thresholds stay and those who fail to meet them leave, then academia as a whole selects those academics who maximize their chances of meeting thresholds (cf. Smaldino and McElreath 2016, O'Connor 2019).

The strategy that maximizes the probability of meeting a threshold need not be the same as the strategy that maximizes expected credit. In the present model an academic who chooses to use QRPs introduces more random variation in the amount of credit she accumulates as compared to doing honest academic work. With the particular parameter values used above, using QRPs decreases expected credit over five time units, but it increases the variance. In these circumstances, it is possible that using QRPs increases the probability of meeting the threshold (despite the lower mean), provided the threshold is relatively high.

And this is exactly what happens. If the threshold the academic aims to meet is, say, 25, then it is better to use QRPs:

$$\Pr(C^*(5) \geq 25) \approx 0.3655 \quad \text{and} \quad \Pr(C^\dagger(5) \geq 25) \approx 0.3708,$$

i.e., doing honest academic work gives her a chance of meeting the threshold of about 36.5%, but using QRPs her chances are just over 37%. The difference increases if the threshold is raised: with a threshold of 30 the honest academic's chances are 24.2%, but using QRPs raises this to 26.6%. With a threshold of 50 the respective chances are 2.1% and 3.5%.

A high threshold corresponds to a low probability of success. This may



be a feature of the competitive process (a job market where only a fraction of academics gets a job, or a tenure process in which less than half of the candidates are given tenure) or a feature of the specific academic (given her own dispositions, background, and training, she is from the outset relatively unlikely to get a job or get tenure), or both.

## 6 How Common is the Incentive to Commit Fraud?

Above I considered a single academic facing the choice to commit fraud (or more generally use QRPs) or do honest research. I highlighted two phenomena. First, the possibility that choosing fraud may be rational for the academic even when the expected credit of individual honest publications is higher, due to cumulative advantage and the way credit accumulates over time. Second, the possibility that choosing fraud may be rational when the expected credit of honest research is higher (even after taking cumulative advantage into account), if the academic's goals require beating a relatively high credit threshold. The previous section may be interpreted as a proof of possibility: it suggests these phenomena may occur, but says nothing about how often they do.

In order to say something a bit stronger, I will now investigate these phenomena a little more systematically as they arise or fail to arise in my model under different parameter settings. If they turn out to arise robustly in the model, this is not sufficient evidence to conclude they commonly arise outside of the model as well, but it is suggestive, especially if one has been persuaded by the preceding sections that the model captures important qualitative features of the incentive structure of academic science.

The parameter settings considered in this section are as follows. First I fix the time scale by ending all simulations at  $t = 5$ . For the intended interpretation of a graduate school education or a tenure clock, this suggests

that one time unit is roughly equal to a year. This is a harmless assumption, as I could set the simulations to end, e.g., at  $t = 60$ , interpret time units as months, rescale the other parameters appropriately, and get exactly the same results. Then I pick a range for the other variables. For the initial publication rate  $\lambda(0)$  this runs from 0.5 to 2 in increments of 0.5 (which can be interpreted as assuming academics vary in their initial average productivity between half a paper and two papers per year). For honest academics, the credit accumulation rate for non-exposed papers  $c_g^*$  is set to either 1 or 2, the exposure rate  $\mu^*$  ranges from 0 to 0.25, and the negative credit for exposed papers  $c_b^*$  from 0 to 0.2.

As in the previous section, the fraudulent academic expects to get more credit from her papers in the short run, but they are more likely to be exposed and accrue more negative credit when this happens. The former is implemented by increasing the credit accumulation rate  $c_g^\dagger$  by a percentage (ranging from 10 % to 60 %) relative to the honest academic's rate  $c_g^*$ . The exposure rate  $\mu^\dagger$  is set to be between 0.05 and 0.25 higher than that of the honest academic. And the negative credit for exposed papers  $c_b^\dagger$  varies from 0.3 to 0.5 (this is not set relative to the corresponding parameter for the honest academic  $c_b^*$ , but all possible values for the fraudulent academic are higher than all possible values for the honest academic). All of this is summarized in table 1.

I focus on the effect of the (extra) credit for non-exposed fraudulent papers, as captured in the parameter  $c_g^\dagger$ . The first result is that even when this parameter is at its lowest setting ( $c_g^\dagger = 1.1 \cdot c_g^*$ , i.e., a 10 % credit premium for fraud), there is a non-negligible range of values of the other parameters for which fraud is a better strategy than honesty in expectation. To state this a bit more precisely, note that if we fix  $c_g^\dagger = 1.1 \cdot c_g^*$  there are 864 possible combinations of values of the other parameters listed in table 1 (though this does involve some double counting because when  $\mu^* = 0$  the value of  $c_b^*$  has no effect). Of these combinations there are 147 (about 17 %) for which

parameter	values
$\lambda(0)$	$\{0.5, 1, 1.5, 2\}$
$c_g^*$	$\{1, 2\}$
$c_b^*$	$\{0, 0.1, 0.2\}$
$\mu^*$	$\{0, 0.05, 0.15, 0.25\}$
$c_g^\dagger$	$c_g^* \cdot \{1.1, 1.2, 1.3, 1.4, 1.5, 1.6\}$
$c_b^\dagger$	$\{0.3, 0.4, 0.5\}$
$\mu^\dagger$	$\mu^* + \{0.05, 0.15, 0.25\}$

Table 1: Parameter values used in this section.

$\mathbb{E}[C^\dagger(5)] > \mathbb{E}[C^*(5)]$ , i.e., fraud is the best strategy in expectation. In the remaining 717 cases, the honest strategy is better in expectation when  $c_g^\dagger$  is at its lowest value.

As we increase the value of  $c_g^\dagger$ , fraud becomes more attractive: the number of parameter settings for which the fraudulent academic expects higher credit than the honest academic gradually increases from 17% when  $c_g^\dagger = 1.1 \cdot c_g^*$  to over 99% (857 out of 864 cases) when  $c_g^\dagger = 1.6 \cdot c_g^*$ . This is shown by the solid lines and black dots in figure 3. Similarly, the number of parameter settings where the fraudulent strategy has a greater probability of exceeding a credit threshold of 30 increases as  $c_g^\dagger$  increases (dashed lines and gray dots in figure 3), as it does when the credit threshold is 50 (dot-dashed lines and white dots in figure 3). In particular, when the credit threshold is 50 and  $c_g^\dagger = 1.6 \cdot c_g^*$ , the fraudulent strategy has a greater probability of exceeding the threshold than the honest strategy in all 864 cases.

This idea can be shown to generalize. That is, regardless of the value of the other parameters, if the credit premium for non-exposed fraudulent papers is large enough, the fraudulent academic is better off than the honest academic. This holds regardless of whether ‘better off’ is cashed out in terms




Figure 3: Percentage of parameter settings for which the fraudulent strategy is preferred over the honest strategy, as a function of the ratio  $c_g^\dagger/c_g^*$ . Solid lines and black dots indicate the percentage of cases where  $\mathbb{E}[C^\dagger(5)] > \mathbb{E}[C^*(5)]$ ; dashed lines and gray dots indicate the percentage where  $\Pr(C^\dagger(5) \geq 30) > \Pr(C^*(5) \geq 30)$ ; dot-dashed lines and white dots indicate the percentage where  $\Pr(C^\dagger(5) \geq 50) > \Pr(C^*(5) \geq 50)$ .

of expected credit or in terms of the probability of meeting a credit threshold.

**Theorem 1.** *Let  $\lambda(0) > 0$  and  $t_1 > 0$ . For all values of  $c_g^*$ ,  $c_b^*$ ,  $\mu^*$ ,  $c_b^\dagger$ , and  $\mu^\dagger$ , there exist values of  $c_g^\dagger$  large enough such that  $\mathbb{E}[C^\dagger(t_1)] > \mathbb{E}[C^*(t_1)]$ . Moreover, for any credit threshold  $\theta > 0$  there exist values of  $c_g^\dagger$  large enough such that  $\Pr(C^\dagger(t_1) \geq \theta) > \Pr(C^*(t_1) \geq \theta)$ .*

I take these simulation results and the theorem to be quite suggestive. They show that, at least within this particular model, the credit incentive for fraud is not an isolated phenomenon. Rather, such an incentive arises system-

atically whenever the credit premium for non-exposed fraudulent papers is large enough (with the theorem showing this in principle, and the simulation results showing that the values of  $c_g^\dagger$  required are not always unrealistically high).

The second phenomenon I highlighted was the possibility that fraud can be incentivized even when the expected credit of honest research is higher, as career success may require beating a credit threshold. While the theorem does not speak to this directly, the simulation results provide some support for this idea. In particular, across 5184 combinations of parameter settings studied, there were 233 instances (about 4.5%) where  $\Pr(C^\dagger(5) \geq 50) > \Pr(C^*(5) \geq 50)$  even though  $\mathbb{E}[C^*(5)] > \mathbb{E}[C^\dagger(5)]$ .

What does this tell us outside the model? There is first the general question whether the model captures the right sort of dynamics to have any relevance to real academics. I have attempted to motivate this throughout the construction of the model and will not say more on this general question here. But there is a second, more specific question to be asked: what should we make of the condition that the credit premium for non-exposed fraudulent papers is ‘large enough’?

My claim is that in a given community, there will be at least some academics for whom there is a credit incentive to use QRPs or fraud, because the parameters of the model will be different for each academic. While some relevant factors are largely fixed within an academic community (e.g., the chance that a fraudulent paper is exposed, the amount of credit one needs to accumulate to have a chance at a job or tenure), others depend on the skills, dispositions, and luck of specific academics (e.g., an academic’s productivity with a fixed level of resources, as captured in the initial publication rate). Importantly, the credit premium for non-exposed fraudulent papers has aspects of both. It partially depends on the academic community (e.g., how much value this community assigns to ‘surprisingness’ or ‘flashiness’), but it also depends on the academic’s ability to dress up shoddy work, adver-

tise its virtues, hide its weaknesses, and quickly convert this into productive resources.

If there is variety in the parameter values experienced by different members of the academic community, there will be some academics within the community for whom the credit premium for non-exposed fraudulent papers is relatively high. The simulation results and the theorem suggest that these academics in particular may have an incentive to use QRPs or fraud.

Whereas the previous section provided (merely) a proof of possibility, this section provides some tentative evidence for a stronger claim. The claim is that in most (if not all) academic communities, there will be some academics who have an incentive to use QRPs or fraud as a direct result of the need to accumulate credit to get a job or tenure.

As noted in the previous section, using QRPs tends to increase the variance in how much credit a given academic accumulates. Speaking somewhat loosely, this means that academics using QRPs are more likely to do either very well or very poorly. This suggests that, in those communities where at least some academics have an incentive to use QRPs, academics using QRPs will be overrepresented at the bottom and at the top of the credit distribution. So in these types of communities, the most famous academics are likely to be the ones using QRPs.

The simulation results discussed in this section illustrate this phenomenon. Suppose that the parameter ranges given in table 1 describe the variety among individual members of a given academic community. Then the most famous academics in that community will be those who accumulate the largest amount of credit. We can get a sense for who this might be by looking at the highest credit totals realized across all simulation runs (since there are 5184 parameter settings with 10,000 simulation runs for the fraudulent strategy and 10,000 runs for the honest strategy this involves more than a hundred million data points).

This maximum of 458.0 units of credit is realized under the fraudulent

strategy, with  $\lambda(0) = 2$ ,  $c_g^\dagger = 3.2$ ,  $c_b^\dagger = 0.5$ , and  $\mu^\dagger = 0.05$ . In contrast, the highest credit achieved across all simulation runs using the honest strategy is 302.1 units. This is arguably not surprising given that the fraudulent strategy tends to do better than the honest strategy by most measures when  $c_g^\dagger = 1.6 \cdot c_g^*$ . But the following is a bit more surprising: even if we restrict ourselves to parameter settings where  $c_g^\dagger = 1.1 \cdot c_g^*$ , we can find a simulation run that achieves 327.9 units of credit.

So, with these parameter ranges, even if the credit premium for non-exposed fraudulent publications is restricted to 10%, we should expect the most famous academics to be fraudulent. And these extremely successful academics will be exactly those who have gotten lucky in that few or none of their papers have been exposed (Heesen 2017a reaches a similar conclusion in a different model with a different notion of luck).

## 7 Conclusion

I have highlighted two phenomena that favor the use of QRPs that become apparent in a dynamic model of the credit economy. First, cumulative advantage may allow a fraudulent academic to be successful even if the fraud does not appear to be paying off at the level of individual publications (cf. Zollman 2019). Second, selection events at particular times in an academic's career may lead fraudulent academics to be successful even when fraud does not pay off in expectation.

The two lessons I want to emphasize in this conclusion are the following. First, the two highlighted phenomena suggest that the incentive to engage in QRPs or fraud may be a bit stronger than it appeared based on previous models of the credit economy (e.g., Bruner 2013, Bright 2017, Heesen 2018).

Second, within my model the most successful academics tend to be the most unscrupulous ones: those who are willing to gamble on fraud and manage to get away with it. If the model accurately captures the dynamics and

incentives related to cumulative advantage and fraud, it raises the worry that those academics who are praised as being the best actually work according to worse than average epistemic standards.

## A Proof of Theorem 1

Let  $\underline{C}(t)$  denote the credit directly associated with a single publication, i.e., the credit that would be accrued up to time  $t$  if we assumed that the publication rate  $\lambda$  drops to zero and stays there immediately after the first publication arrives. Let  $T \sim \text{Exp}(\lambda(0))$  be the arrival time for that publication. Let  $X \sim \text{Exp}(\mu)$  be the waiting time until the publication is exposed (so that exposure occurs at time  $T + X$ ).

We first consider the expected credit conditional on  $T = t < t_1$ . The density function of  $X$  is given by  $f_X(x) = \mu e^{-\mu x}$ . If  $X \geq t_1 - t$  the credit accrued is  $c_g(t_1 - t)$ , otherwise it is  $c_g X - c_b(t_1 - t - X)$ . So

$$\begin{aligned} \mathbb{E}[\underline{C}(t_1) | T = t] &= \int_0^\infty (c_g \min\{t_1 - t, x\} - c_b \max\{t_1 - t - x, 0\}) f_X(x) dx \\ &= \frac{c_g + c_b}{\mu} (1 - e^{-\mu(t_1 - t)}) - c_b(t_1 - t). \end{aligned}$$

A few observations. First, this justifies equation (6). Second,  $\mathbb{E}[\underline{C}(t_1) | T = t]$  is a linearly increasing function of  $c_g$ . Consequently we can guarantee that the expectation is positive by choosing  $c_g$  large enough:

$$\mathbb{E}[\underline{C}(t_1) | T = t] > 0 \quad \text{if and only if} \quad c_g > c_b \frac{\mu(t_1 - t)}{1 - e^{-\mu(t_1 - t)}} - c_b. \quad (7)$$

Third, if  $\mathbb{E}[\underline{C}(t_1) | T = t_0] > 0$  then for all  $t \in (t_0, t_1)$  also  $\mathbb{E}[\underline{C}(t_1) | T = t] > 0$ . This can be seen from (7) by noting that  $x/(1 - e^{-x})$  is a strictly increasing function for all  $x$ .

Now we consider  $\mathbb{E}[\underline{C}(t_1)]$ . The density function of  $T$  is given by  $f_T(t) = \lambda(0)e^{-\lambda(0)t}$ . If  $T \geq t_1$  the credit accrued is zero, otherwise the expected credit



is  $\mathbb{E}[\underline{C}(t_1) \mid T]$ . So

$$\begin{aligned}\mathbb{E}[\underline{C}(t_1)] &= \int_0^{t_1} \mathbb{E}[\underline{C}(t_1) \mid T = t] f_T(t) dt \\ &= \frac{\mu(c_g \lambda(0) + c_b \mu)(1 - e^{-\lambda(0)t_1}) - \lambda(0)^2(c_g + c_b)(1 - e^{-\mu t_1})}{\lambda(0)\mu(\mu - \lambda(0))} - c_b t_1.\end{aligned}$$

We are particularly interested in how  $\mathbb{E}[\underline{C}(t_1)]$  varies with  $c_g$ :

$$\frac{d}{dc_g} \mathbb{E}[\underline{C}(t_1)] = \frac{\mu(1 - e^{-\lambda(0)t_1}) - \lambda(0)(1 - e^{-\mu t_1})}{\mu(\mu - \lambda(0))}.$$

So  $\mathbb{E}[\underline{C}(t_1)]$  is also a linear function of  $c_g$ . To see that it is an increasing function of  $c_g$  it suffices to show that the derivative is positive. One way to see this is by interchanging the derivative and the integral (which is legitimate because the functions involved are differentiable and bounded):

$$\begin{aligned}\frac{d}{dc_g} \mathbb{E}[\underline{C}(t_1)] &= \int_0^{t_1} \frac{d}{dc_g} \mathbb{E}[\underline{C}(t_1) \mid T = t] f_T(t) dt \\ &= \int_0^{t_1} \frac{1}{\mu} (1 - e^{-\mu(t_1-t)}) \lambda(0) e^{-\lambda(0)t} dt > 0\end{aligned}$$

because the integrand is strictly positive for all  $0 \leq t < t_1$ .

How does  $\mathbb{E}[\underline{C}(t_1)]$  relate to  $\mathbb{E}[C(t_1)]$ ? Recall that  $\underline{C}(t)$  only counts the credit associated with a single publication, whereas  $C(t)$  tracks the credit for all publications combined. The probability distribution for the arrival time of each subsequent publication is (by design) quite complicated, as it depends on the number of previous publications, the number of previous exposures, and the precise arrival times of each of these. But conditional on its arrival time  $t$ , we know that the contribution each publication makes to the expected credit is equal to  $\mathbb{E}[\underline{C}(t_1) \mid T = t]$ . Moreover we know from the third observation above that if  $\mathbb{E}[\underline{C}(t_1) \mid T = 0] > 0$  then  $\mathbb{E}[\underline{C}(t_1) \mid T = t] > 0$  for all  $t < t_1$ . So it follows that if  $\mathbb{E}[\underline{C}(t_1) \mid T = 0] > 0$  (which we can make sure is true by choosing  $c_g$  sufficiently high) then the contribution to the

expected credit of each publication beyond the first is positive, and therefore  $\mathbb{E}[C(t_1)] > \mathbb{E}[\underline{C}(t_1)]$ .

Since  $\mathbb{E}[\underline{C}(t_1)]$  increases linearly with  $c_g$ , we can make  $\mathbb{E}[\underline{C}(t_1)]$  arbitrarily high by setting  $c_g$  sufficiently high. In particular, if  $\lambda(0)$ ,  $c_g^*$ ,  $c_b^*$ , and  $\mu^*$  are fixed then  $\mathbb{E}[C^*(t_1)]$  is thereby fixed as well. For arbitrary (fixed) values of  $c_b^\dagger$  and  $\mu^\dagger$  we can then choose  $c_g^\dagger$  high enough such that  $\mathbb{E}[\underline{C}^\dagger(t_1)] > \mathbb{E}[C^*(t_1)]$ . If we also choose  $c_g^\dagger$  high enough so the inequality from the previous paragraph obtains, we get the desired result regarding expectation:

$$\mathbb{E}[C^\dagger(t_1)] > \mathbb{E}[\underline{C}^\dagger(t_1)] > \mathbb{E}[C^*(t_1)].$$

Next we consider the probability of exceeding a credit threshold  $\theta > 0$ . We proceed similarly by investigating  $\underline{C}(t_1)$ . We know that

$$\Pr(\underline{C}(t_1) = 0) = \Pr(T > t_1) = e^{-\lambda(0)t_1}.$$

If  $T < t_1$  then  $\underline{C}(t_1) = \min\{c_g X - c_b(t_1 - T - X), c_g(t_1 - T)\}$ . Therefore  $\Pr(\underline{C}(t_1) > \theta \mid T > t_1 - \theta/c_g) = 0$ . Whereas if  $t < t_1 - \theta/c_g$  then  $c_g(t_1 - t) > \theta$  and hence

$$\begin{aligned} \Pr(\underline{C}(t_1) > \theta \mid T = t) &= \Pr(c_g X - c_b(t_1 - t - X) > \theta) \\ &= \Pr\left(X > \frac{\theta + c_b(t_1 - t)}{c_g + c_b}\right) \\ &= \exp\left\{-\mu \frac{\theta + c_b(t_1 - t)}{c_g + c_b}\right\}, \end{aligned}$$

where  $\exp\{x\} = e^x$ . We can now find an expression for the probability that

$\underline{C}(t_1)$  exceeds  $\theta$ :

$$\begin{aligned} \Pr(\underline{C}(t_1) > \theta) &= \int_0^{t_1 - \theta/c_g} \Pr(\underline{C}(t_1) > \theta \mid T = t) f_T(t) dt \\ &= \frac{1}{1 - \frac{\mu}{\lambda(0)} \frac{c_b}{c_g + c_b}} \left( \exp \left\{ -\mu \frac{\theta + c_b t_1}{c_g + c_b} \right\} - e^{-\lambda(0)t_1 - (\mu - \lambda(0))\theta/c_g} \right). \end{aligned}$$

We see that this probability is increasing with  $c_g$  by inspecting the integral: (1) the integrand is positive for all relevant values of  $t$ , (2) the range of integration increases with  $c_g$  because  $\frac{d}{dc_g}(t_1 - \theta/c_g) = \theta/c_g^2 > 0$ , and (3) the integrand increases with  $c_g$ :

$$\frac{d}{dc_g} \Pr(\underline{C}(t_1) > \theta \mid T = t) = \mu \frac{\theta + c_b(t_1 - t)}{(c_g + c_b)^2} \exp \left\{ -\mu \frac{\theta + c_b(t_1 - t)}{c_g + c_b} \right\} > 0.$$

Moreover, we see that as  $c_g$  increases, the probability that  $\underline{C}(t_1)$  exceeds  $\theta$  approaches the probability of any arrivals at all. Since for any constant  $a$ , the fraction  $a/(c_g + c_b)$  vanishes as  $c_g$  gets large, we have

$$\begin{aligned} \lim_{c_g \rightarrow \infty} \frac{1}{1 - \frac{\mu}{\lambda(0)} \frac{c_b}{c_g + c_b}} &= 1, \\ \lim_{c_g \rightarrow \infty} \exp \left\{ -\mu \frac{\theta + c_b t_1}{c_g + c_b} \right\} &= 1, \\ \lim_{c_g \rightarrow \infty} e^{-\lambda(0)t_1 - (\mu - \lambda(0))\theta/c_g} &= e^{-\lambda(0)t_1}, \quad \text{and therefore} \\ \lim_{c_g \rightarrow \infty} \Pr(\underline{C}(t_1) > \theta) &= 1 - e^{-\lambda(0)t_1} = \Pr(T < t_1). \end{aligned}$$

We now want to show that  $\Pr(C(t_1) > \theta) \geq \Pr(\underline{C}(t_1) > \theta)$  for sufficiently large values of  $c_g$ .

Let  $N$  (a random variable) be the total number of publications that occurs in the time interval  $[0, t_1]$ . The event  $N = 0$  occurs just in case  $T > t_1$  (which entails  $C(t_1) = \underline{C}(t_1) = 0$ ), so we can restrict attention to cases where  $N \geq 1$ . We now claim that if  $N \geq 1$ , then we can choose  $c_g$  high enough that (with

arbitrarily high probability)  $N$  is “large”.

In fact, the claim we will prove is a little bit stronger. Let  $S$  denote the sum over all publications of the time interval from publication until  $t_1$  (so ignoring whether they are exposed). That is, if  $T_0, T_1, \dots, T_{N-1}$  are the publication times, then  $S = \sum_{i=0}^{N-1} (t_1 - T_i)$ . We show that, with an appropriate choice of  $c_g$ , we can make  $S$  arbitrarily large with arbitrarily high probability (conditional on  $N \geq 1$ ). That is, not only are there arbitrarily many publications, but they do not all occur so close to  $t_1$  that the combined time they are in existence remains small.

Let  $\varepsilon > 0$ . We want to show for some arbitrary (large) constant  $k$  that

$$\Pr(S \geq k \mid N \geq 1) = \frac{\Pr(S \geq k, N \geq 1)}{\Pr(N \geq 1)} = \frac{\Pr(S \geq k, T < t_1)}{\Pr(T < t_1)} > 1 - \varepsilon.$$

Let

$$t^* = \min \left\{ \frac{1}{\lambda(0)} \log \left( 1 + \frac{\varepsilon}{3} (e^{\lambda(0)t_1} - 1) \right), \frac{1}{\mu} \log \left( \frac{1}{1 - \varepsilon/3} \right) \right\} > 0.$$

Then

$$\begin{aligned} \Pr(S \geq k \mid N \geq 1) &\geq \frac{\Pr(S \geq k, T < t_1 - t^*)}{\Pr(T < t_1)} \\ &= \Pr(S \geq k \mid T < t_1 - t^*) \frac{1 - e^{-\lambda(0)(t_1 - t^*)}}{1 - e^{-\lambda(0)t_1}} \\ &\geq (1 - \varepsilon/3) \Pr(S \geq k \mid T < t_1 - t^*). \end{aligned}$$

Let  $N^*$  denote the number of publications that occur in the time interval  $[T, T + \frac{1}{2}t^*]$ . Conditional on  $T < t_1 - t^*$ , we know that  $T + t^*/2 < t_1 - t^*/2$ , i.e., we have  $N^*$  publications occurring before  $t_1 - t^*/2$ , so  $S \geq N^* \cdot t^*/2$ . Thus if we let  $n = 2k/t^*$  we have

$$\Pr(S \geq k \mid T < t_1 - t^*) \geq \Pr(N^* \geq n \mid T < t_1 - t^*).$$

Let  $G_2$  denote the event that the first 2 publications fail to be exposed during the first  $t^*/2$  time units after their publication (assuming these publications occur at all). Note that

$$\Pr(G_2) \geq e^{-\mu t^*} \geq 1 - \varepsilon/3.$$

Putting this together, so far we have

$$\begin{aligned} \Pr(S \geq k \mid N \geq 1) &\geq (1 - \varepsilon/3) \Pr(S \geq k \mid T < t_1 - t^*) \\ &\geq (1 - \varepsilon/3) \Pr(N^* \geq n \mid T < t_1 - t^*) \\ &\geq (1 - \varepsilon/3) \Pr(N^* \geq n, G_2 \mid T < t_1 - t^*) \\ &\geq (1 - \varepsilon/3)^2 \Pr(N^* \geq n, \mid G_2, T < t_1 - t^*). \end{aligned}$$

We now consider the probability in the last line above, i.e., the probability of at least  $n$  publications in the time interval  $[T, T + \frac{1}{2}t^*]$  conditional on the first 2 publications not being exposed during that same time interval. Choose

$$c_g = \max \left\{ n \cdot c_b, \exp \left\{ \frac{16n}{(t^*)^2} + \frac{24}{\varepsilon(t^*)^2} \right\} \right\}.$$

For the purpose of determining the probability that  $N^* \geq n$ , we may assume that  $c(t) \geq c_g$  for all  $t \in [T, T + \frac{1}{2}t^*]$ . This is because (a) after the first publication at time  $T$ ,  $p_g = 1$  and  $p_b = 0$  so  $c_g p_g - c_b p_b = c_g$ , (b) due to the condition  $G_2$ , the second publication occurs before the first exposure, (c) after the second publication but before the  $n$ -th publication, due to condition  $G_2$ ,  $p_g \geq 2$  and  $p_b \leq n - 2$  and hence  $c_g p_g - c_b p_b \geq c_g + n c_b - (n - 2)c_b \geq c_g$ .

It follows from (3) that the rate of publication is

$$\lambda(t) \geq \lambda(0) + \log(c_g + 1)(t - T)$$

for all  $t \in [T, T + \frac{1}{2}t^*]$ . This means that (at least as long as  $N^* < n$ ) the

so-called average intensity during the time interval is

$$\Lambda\left(T, T + \frac{1}{2}t^*\right) = \int_T^{T+t^*/2} \lambda(t) dt \geq \frac{1}{2}\lambda(0)t^* + \frac{1}{8}\log(c_g + 1)(t^*)^2.$$

Let  $Y$  be a Poisson-distributed random variable with rate parameter  $\bar{\Lambda} = \lambda(0)t^*/2 + \log(c_g + 1)(t^*)^2/8$ . Then  $\Pr(N^* \geq n, | G_2, T < t_1 - t^*) \geq \Pr(Y \geq n)$ .

Note that

$$\bar{\Lambda} \geq \frac{1}{8}\log(c_g)(t^*)^2 \geq 2n + 3/\varepsilon.$$

Since  $Y$  follows a Poisson distribution,  $\mathbb{E}[Y] = \text{Var}(Y) = \bar{\Lambda}$ . Since  $\mathbb{E}[Y] = \bar{\Lambda} > n$ , we can apply Cantelli's inequality (a one-sided version of the Chebyshev inequality) to get

$$\Pr(Y \geq n) = \Pr(Y - \mathbb{E}[Y] \geq n - \mathbb{E}[Y]) \geq 1 - \frac{\text{Var}(Y)}{\text{Var}(Y) + (n - \mathbb{E}[Y])^2}.$$

Thus

$$\begin{aligned} \Pr(Y \geq n) &\geq 1 - \frac{\text{Var}(Y)}{\text{Var}(Y) + (n - \mathbb{E}[Y])^2} = 1 - \frac{1}{\bar{\Lambda} + 1 - 2n + n^2/\bar{\Lambda}} \\ &\geq 1 - \frac{1}{\bar{\Lambda} - 2n} \geq 1 - \varepsilon/3. \end{aligned}$$

But now we are done:

$$\begin{aligned} \Pr(S \geq k | N \geq 1) &\geq (1 - \varepsilon/3)^2 \Pr(N^* \geq n, | G_2, T < t_1 - t^*) \\ &\geq (1 - \varepsilon/3)^2 \Pr(Y \geq n) \geq (1 - \varepsilon/3)^3 \\ &> 1 - \varepsilon. \end{aligned}$$

Returning to the big picture, recall that we are aiming to show that, for large  $c_g$ ,

$$\Pr(C(t_1) > \theta | N \geq 1) \geq \Pr(\underline{C}(t_1) > \theta | N \geq 1).$$

It suffices to show, for some  $\delta > 0$ , that

$$\lim_{c_g \rightarrow \infty} \Pr(C(t_1) - \underline{C}(t_1) \geq \delta \mid N \geq 1) = 1,$$

where  $C(t_1) - \underline{C}(t_1)$  is the credit accrued to all publications except the first. For each of these publications  $i = 1, \dots, N - 1$ , consider  $t_1 - T_i$ , i.e., the amount of time between publication and the cutoff time  $t_1$ . Suppose we form groups of publications such that for each group  $b$ , the sum of these amounts of time is at least  $t_1$  and at most  $2t_1$ , i.e.,

$$t_1 \leq \sum_{i \in b} (t_1 - T_i) \leq 2t_1.$$

Because  $t_1 - T_i \leq t_1$  for each  $i$ , we can be sure that we can form such groups without this sum exceeding  $2t_1$  for any of them. Moreover, since we have just shown that  $S = \sum_{i=0}^{N-1} (t_1 - T_i)$  becomes large for large values of  $c_g$ , we can guarantee that the number of such groups  $M$  will be large.

We now place the following lower bound on the amount of credit accrued by the combined publications in each group: if at least one of the publications in the group is exposed at any time, we assume all the publications in the group are exposed the entire time. Thus, for each group, the probability of at least one exposure is at most  $q = e^{-2\mu t_1}$ , and the credit associated with each group in which at least one exposure occurs is at least  $-2c_b t_1$ . Conversely, the probability of no exposures for a given group is at least  $1 - q$ , and the credit associated with such a group is at least  $c_g t_1$ .

Let  $Z$  be a random variable following a binomial distribution with parameters  $M$  (the number of trials) and  $1 - q$  (the success probability), i.e.,  $Z$  denotes the number of groups in which no exposures occur. Then according to the conservative way of estimating the accrued credit explained above,

$$C(t_1) - \underline{C}(t_1) \geq c_g t_1 Z - 2c_b t_1 (M - Z).$$

It follows that

$$\begin{aligned}
\Pr(C(t_1) - \underline{C}(t_1) \geq \delta \mid N \geq 1) &\geq \Pr(c_g t_1 Z - 2c_b t_1 (M - Z) \geq \delta \mid N \geq 1) \\
&= \Pr\left(Z \geq \frac{2c_b M + \delta}{c_g + 2c_b} \mid N \geq 1\right) \\
&= \sum_{m=0}^{\infty} \Pr\left(Z \geq \frac{2c_b m + \delta}{c_g + 2c_b} \mid M = m\right) \Pr(M = m \mid N \geq 1).
\end{aligned}$$

Let  $\varepsilon > 0$ . Choose  $c_g$  larger than  $\frac{2c_b + \delta}{(1-q)^2}$  but also large enough such that

$$\Pr\left(S \geq \frac{4t_1}{q(1-q)\varepsilon} \mid N \geq 1\right) \geq 1 - \varepsilon/2$$

(which we have previously shown is possible). Since  $Z$  is binomial, we know its mean is  $\mathbb{E}[Z \mid M = m] = (1-q)m$  and its variance  $\text{Var}(Z \mid M = m) = q(1-q)m$ . From  $c_g \geq \frac{2c_b + \delta}{(1-q)^2}$  it follows that for all  $m \geq 1$ ,

$$c_g(1-q)^2 m \geq (2c_b + \delta)m \geq 2c_b(1 - (1-q)^2)m + \delta.$$

Hence, for all  $m \geq 1$ ,  $(c_g + 2c_b)(1-q)^2 m \geq 2c_b m + \delta$  which entails  $\frac{2c_b m + \delta}{c_g + 2c_b} \leq (1-q)^2 m = (1-q)m - q(1-q)m$ . It follows that  $\frac{2c_b m + \delta}{c_g + 2c_b} - (1-q)m < -q(1-q)m < 0$ . So we can again use Cantelli's inequality to get

$$\begin{aligned}
\Pr\left(Z \geq \frac{2c_b m + \delta}{c_g + 2c_b} \mid M = m\right) &\geq 1 - \frac{q(1-q)m}{q(1-q)m + \left(\frac{2c_b m + \delta}{c_g + 2c_b} - (1-q)m\right)^2} \\
&\geq 1 - \frac{q(1-q)m}{q(1-q)m + (-q(1-q)m)^2} \\
&= 1 - \frac{1}{q(1-q)m + 1}.
\end{aligned}$$

From the above we get that whenever  $m \geq \frac{2}{q(1-q)\varepsilon}$ ,

$$\Pr\left(Z \geq \frac{2c_b m + \delta}{c_g + 2c_b} \mid M = m\right) \geq 1 - \varepsilon/2.$$



Thus

$$\begin{aligned}
& \Pr(C(t_1) - \underline{C}(t_1) \geq \delta \mid N \geq 1) \\
& \geq \sum_{m=\frac{2}{q(1-q)\varepsilon}}^{\infty} \Pr\left(Z \geq \frac{2c_b m + \delta}{c_g + 2c_b} \mid M = m\right) \Pr(M = m \mid N \geq 1) \\
& \geq (1 - \varepsilon/2) \Pr\left(M \geq \frac{2}{q(1-q)\varepsilon} \mid N \geq 1\right) \\
& \geq (1 - \varepsilon/2) \Pr\left(S \geq \frac{4t_1}{q(1-q)\varepsilon} \mid N \geq 1\right) \\
& \geq (1 - \varepsilon/2)^2 > 1 - \varepsilon.
\end{aligned}$$

Putting everything together, we have now shown that

$$\begin{aligned}
& \lim_{c_g \rightarrow \infty} \Pr(\underline{C}(t_1) > \theta) = \Pr(N \geq 1) = 1 - e^{-\lambda(0)t_1}, \\
& \lim_{c_g \rightarrow \infty} \Pr(C(t_1) - \underline{C}(t_1) \geq \delta \mid N \geq 1) = 1.
\end{aligned}$$

From this it follows that

$$\lim_{c_g \rightarrow \infty} \Pr(C(t_1) > \theta) = \Pr(N \geq 1) = 1 - e^{-\lambda(0)t_1}.$$

The desired claim follows straightforwardly from this. In particular, if  $\lambda(0)$ ,  $c_g^*$ ,  $c_b^*$ , and  $\mu^*$  are fixed then  $\Pr(C^*(t_1) > \theta)$  is thereby fixed as well. Moreover, we know that  $\Pr(C^*(t_1) > \theta) < 1 - e^{-\lambda(0)t_1}$  since  $\theta > 0$  and  $\Pr(C^*(t_1) = 0) = e^{\lambda(0)t_1}$  and  $\Pr(0 < C^*(t_1) < \theta) > 0$ . Thus  $\Pr(C^*(t_1) > \theta)$  is less than the limiting value by some positive amount. For arbitrary fixed values of  $c_b^\dagger$  and  $\mu^\dagger$ , the limiting result established above guarantees that we can choose  $c_g^\dagger$  large enough such that  $\Pr(C^\dagger(t_1) > \theta)$  is closer to  $1 - e^{-\lambda(0)t_1}$  than  $\Pr(C^*(t_1) > \theta)$ . Thus, in particular,

$$\Pr(C^\dagger(t_1) > \theta) > \Pr(C^*(t_1) > \theta).$$

## References

- P. Abraham. Duplicate and salami publications. *Journal of Postgraduate Medicine*, 46(2):67–69, 2000. URL <https://www.ncbi.nlm.nih.gov/pubmed/11013467>.
- Thomas Boyer. Is a bird in the hand worth two in the bush? Or, whether scientists should publish intermediate results. *Synthese*, 191(1): 17–35, 2014. ISSN 0039-7857. doi: 10.1007/s11229-012-0242-4. URL <https://doi.org/10.1007/s11229-012-0242-4>.
- Liam Kofi Bright. On fraud. *Philosophical Studies*, 174(2):291–310, 2017. ISSN 1573-0883. doi: 10.1007/s11098-016-0682-7. URL <https://doi.org/10.1007/s11098-016-0682-7>.
- François Brischoux and Frédéric Angelier. Academia’s never-ending selection for productivity. *Scientometrics*, 103(1):333–336, April 2015. ISSN 1588-2861. doi: 10.1007/s11192-015-1534-5. URL <https://doi.org/10.1007/s11192-015-1534-5>.
- Justin P. Bruner. Policing epistemic communities. *Episteme*, 10(4):403–416, December 2013. ISSN 1750-0117. doi: 10.1017/epi.2013.34. URL <https://doi.org/10.1017/epi.2013.34>.
- Colin F. Camerer, Anna Dreber, Felix Holzmeister, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Gideon Nave, Brian A. Nosek, Thomas Pfeiffer, Adam Altmejd, Nick Buttrick, Taizan Chan, Yiling Chen, Eskil Forsell, Anup Gampa, Emma Heikenstein, Lily Hummer, Taisuke Imai, Siri Isaksson, Dylan Manfredi, Julia Rose, Eric-Jan Wagenmakers, and Hang Wu. Evaluating the replicability of social science experiments in *Nature* and *Science* between 2010 and 2015. *Nature Human Behaviour*, 2(9):637–644, 2018. doi: 10.1038/s41562-018-0399-z. URL <https://doi.org/10.1038/s41562-018-0399-z>.
- Aaron Clauset, Samuel Arbesman, and Daniel B. Larremore. Systematic

- inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1):e1400005, 2015. doi: 10.1126/sciadv.1400005. URL <https://doi.org/10.1126/sciadv.1400005>.
- Jonathan R. Cole and Stephen Cole. *Social Stratification in Science*. The University of Chicago Press, Chicago, 1973. ISBN 0226113388.
- Partha Dasgupta and Paul A. David. Toward a new economics of science. *Research Policy*, 23(5):487–521, 1994. ISSN 0048-7333. doi: 10.1016/0048-7333(94)01002-1. URL <http://www.sciencedirect.com/science/article/pii/0048733394010021>.
- Helen de Cruz. Prestige bias: An obstacle to a just academic philosophy. *Ergo*, 5(10):259–287, 2018. doi: 10.3998/ergo.12405314.0005.010. URL <https://doi.org/10.3998/ergo.12405314.0005.010>.
- Thomas A. DiPrete and Gregory M. Eirich. Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments. *Annual Review of Sociology*, 32:271–297, 2006. ISSN 03600572. URL <http://www.jstor.org/stable/29737740>.
- Daniele Fanelli. Do pressures to publish increase scientists’ bias? An empirical support from US states data. *PLoS ONE*, 5(4):e10271, April 2010. doi: 10.1371/journal.pone.0010271. URL <https://doi.org/10.1371/journal.pone.0010271>.
- Annette Flanagin, Lisa A. Carey, Phil B. Fontanarosa, Stephanie G. Phillips, Brian P. Pace, George D. Lundberg, and Drummond Rennie. Prevalence of articles with honorary authors and ghost authors in peer-reviewed medical journals. *Journal of the American Medical Association*, 280(3):222–224, July 1998. ISSN 0098-7484. doi: 10.1001/jama.280.3.222. URL <https://doi.org/10.1001/jama.280.3.222>.
- Remco Heesen. Academic superstars: Competent or lucky? *Synthese*, 194(11):4499–4518, 2017a. ISSN 1573-0964. doi: 10.1007/s11229-016-1146-5.

URL <https://doi.org/10.1007/s11229-016-1146-5>.

Remco Heesen. Communism and the incentive to share in science. *Philosophy of Science*, 84(4):698–716, 2017b. ISSN 0031-8248. doi: 10.1086/693875. URL <https://doi.org/10.1086/693875>.

Remco Heesen. Why the reward structure of science makes reproducibility problems inevitable. *The Journal of Philosophy*, 115(12):661–674, 2018. ISSN 0022-362X. doi: 10.5840/jphil20181151239. URL <https://doi.org/10.5840/jphil20181151239>.

Bryce Huebner and Liam Kofi Bright. Collective responsibility and fraud in scientific communities. In Saba Bazargan-Forward and Deborah Tollefsen, editors, *The Routledge Handbook of Collective Responsibility*, chapter 25, pages 358–372. Routledge, New York, 2020.

Norbert L. Kerr. HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2(3):196–217, 1998. doi: 10.1207/s15327957pspr0203\_4. URL [https://doi.org/10.1207/s15327957pspr0203\\_4](https://doi.org/10.1207/s15327957pspr0203_4).

Philip Kitcher. The division of cognitive labor. *The Journal of Philosophy*, 87(1):5–22, 1990. ISSN 0022-362X. doi: 10.2307/2026796. URL <https://doi.org/10.2307/2026796>.

Bruno Latour and Steve Woolgar. *Laboratory Life: The Construction of Scientific Facts*. Princeton University Press, Princeton, second edition, 1986.

Carole J. Lee, Cassidy R. Sugimoto, Guo Zhang, and Blaise Cronin. Bias in peer review. *Journal of the American Society for Information Science and Technology*, 64(1):2–17, 2013. ISSN 1532-2890. doi: 10.1002/asi.22784. URL <https://doi.org/10.1002/asi.22784>.

Robert K. Merton. Priorities in scientific discovery: A chapter in the sociology of science. *American Sociological Review*, 22(6):635–659, 1957.

- ISSN 00031224. doi: 10.2307/2089193. URL <https://doi.org/10.2307/2089193>. Reprinted in Merton (1973, chapter 14).
- Robert K. Merton. The Matthew effect in science. *Science*, 159(3810):56–63, 1968. ISSN 00368075. URL <http://www.jstor.org/stable/1723414>. Reprinted in Merton (1973, chapter 20).
- Robert K. Merton. *The Sociology of Science: Theoretical and Empirical Investigations*. The University of Chicago Press, Chicago, 1973. ISBN 0226520919.
- James R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, 1998. doi: 10.1017/CBO9780511810633. URL <https://doi.org/10.1017/CBO9780511810633>.
- Brian A. Nosek and Timothy M. Errington. Making sense of replications. *eLife*, 6:e23383, Jan 2017. ISSN 2050-084X. doi: 10.7554/eLife.23383. URL <https://doi.org/10.7554/eLife.23383>.
- Cailin O’Connor. The natural selection of conservative science. *Studies in History and Philosophy of Science Part A*, 76:24–29, 2019. ISSN 0039-3681. doi: 10.1016/j.shpsa.2018.09.007. URL <https://doi.org/10.1016/j.shpsa.2018.09.007>.
- Cailin O’Connor and Justin Bruner. Dynamics and diversity in epistemic communities. *Erkenntnis*, 84(1):101–119, 2019. ISSN 1572-8420. doi: 10.1007/s10670-017-9950-y. URL <https://doi.org/10.1007/s10670-017-9950-y>.
- Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*, 349(6251):aac4716, Aug 2015. doi: 10.1126/science.aac4716. URL <http://www.sciencemag.org/content/349/6251/aac4716.abstract>.
- Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. False-positive psychology: Undisclosed flexibility in data collection and analysis allows pre-

- senting anything as significant. *Psychological Science*, 22(11):1359–1366, 2011. doi: 10.1177/0956797611417632. URL <http://pss.sagepub.com/content/22/11/1359.abstract>.
- Paul E. Smaldino and Richard McElreath. The natural selection of bad science. *Royal Society Open Science*, 3(9):160384, 2016. doi: 10.1098/rsos.160384. URL <http://rsos.royalsocietypublishing.org/content/3/9/160384>.
- Michael Strevens. The role of the priority rule in science. *The Journal of Philosophy*, 100(2):55–79, 2003. ISSN 0022-362X. doi: 10.5840/jphil2003100224. URL <https://doi.org/10.5840/jphil2003100224>.
- Henk C. Tijms. *A First Course in Stochastic Models*. John Wiley & Sons, Chichester, 2003. doi: 10.1002/047001363X. URL <https://doi.org/10.1002/047001363X>.
- Andrew Tomkins, Min Zhang, and William D. Heavlin. Reviewer bias in single- versus double-blind peer review. *Proceedings of the National Academy of Sciences*, 114(48):12708–12713, 2017. doi: 10.1073/pnas.1707323114. URL <https://doi.org/10.1073/pnas.1707323114>.
- K. Brad Wray. The impact of collaboration on the epistemic cultures of science. In Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg, editors, *Scientific Collaboration and Collective Knowledge*, chapter 6, pages 117–134. Oxford University Press, Oxford, 2017. doi: 10.1093/oso/9780190680534.003.0006. URL <https://doi.org/10.1093/oso/9780190680534.003.0006>.
- Kevin J. S. Zollman. The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1):5–33, 2018. ISSN 0022-362X. doi: 10.5840/jphil201811511. URL <https://doi.org/10.5840/jphil201811511>.
- Kevin J. S. Zollman. The scientific Ponzi scheme. Manuscript, 2019.