

#### Paper to be presented at the

DRUID Society Conference 2014, CBS, Copenhagen, June 16-18

## Above a Swamp: A Theory of High-Quality Scientific Production

#### **Bralind Kiri**

University of Bologna
Department of Economics
bralind.kiri2@unibo.it

#### Nicola Lacetera

University of Toronto
Rotman School of Management
nicola.lacetera@utoronto.ca

#### Lorenzo Zirulia

University of Bologna
Department of Economics
lorenzo.zirulia@unibo.it

#### Abstract

Building upon previous research and potentially pointing out its limitations is essential for a healthy working of the scientific community, as it allows science to self-correct and evolve. In this paper, we propose a game-theoretic model to investigate the incentives of scientists to perform these activities of control and criticism. The two-player game includes a scientist producing a new scientific result, and another scientist who can verify that result. The quality of the new result depends on the amount of effort chosen by a scientist. A high-quality paper would stand against any future work building on it, whereas a low-quality paper presents limitations that are detected only if verification through additional work occurs. We show that a certain fraction of low-quality scientific knowledge characterizes all the equilibria in the basic version of model. As a consequence, the absence of low-quality research in a field must interpreted as the lack of verification activity and then as a potential source of concern. We suggest that increasing the benefit that a scientist derives from building and potentially qualifying previous research should be considered as a primary way to improve scientific research reliability. By contrast, softening overall incentives to publish per se does not have an impact on research quality, although it increases the fraction of low-quality papers that are identified.

## Above a Swamp: A Theory of High-Quality Scientific Production

February 28, 2014

#### Abstract

Building upon previous research and potentially pointing out its limitations is essential for a healthy working of the scientific community, as it allows science to self-correct and evolve. In this paper, we propose a game-theoretic model to investigate the incentives of scientists to perform these activities of control and criticism. The two-player game includes a scientist producing a new scientific result, and another scientist who can verify that result. The quality of the new result depends on the amount of effort chosen by a scientist. A high-quality paper would stand against any future work building on it, whereas a low-quality paper presents limitations that are detected only if verification through additional work occurs. We show that a certain fraction of low-quality scientific knowledge characterizes all the equilibria in the basic version of model. As a consequence, the absence of low-quality research in a field must interpreted as the lack of verification activity and then as a potential source of concern. We suggest that increasing the benefit that a scientist derives from building and potentially qualifying previous research should be considered as a primary way to improve scientific research reliability. By contrast, softening overall incentives to publish per se does not have an impact on research quality, although it increases the fraction of low-quality papers that are identified.

Science does not rest upon solid bedrock. The bold structure of its theories rises, as it were, above a swamp. It is like a building erected on piles. The piles are driven down from above into the swamp, but not down to any natural or 'given' base; and when we cease our attempts to drive our piles into a deeper layer, it is not because we have reached firm ground. We simply stop when we are satisfied that they are firm enough to carry the structure, at least for the time being.

Karl R. Popper, The Logic of Scientific Discovery (1959, p. 111).

#### 1 Introduction

The importance of science for economic growth and social welfare is widely recognized (Stephan, 2012). When referring to scientific knowledge, one has typically in mind knowledge that is "true." However, not all the outcomes of scientific activities have this characteristic. An interest toward understanding whether and why science can "go wrong", while originally confined within the boundaries of specialists in the scientific community, has more recently been displayed also by policymakers and the public opinion concerned about the economic and social consequences of these phenomena, as witnessed, for example, by a long report in a recent issue of the *Economist* (October 19th, 2013).

Scientific fraud is an extreme case of production of false knowledge (Broad and Wade, 1982; Lacetera and Zirulia, 2011). More commonly and frequently, however, the scientific results that are reported in journals might have not been checked against all potential robustness tests, may rely on questionable methods of inquiry, or may contain honest mistakes. In some cases flawed or erroneous results, just like fraudulent ones, are retracted from publications (Azoulay et al., 2013; Lu et al., 2013); most often, they are just seen as a "natural" step toward better theories and findings. Karl Popper's view of science, for example, holds that a finding or theory can be defined as scientific to the extent that it is falsifiable (Popper, 1959). Therefore, at each given time, the body of scientific knowledge will include findings that are limited or flawed in some ways. Over time, the reliability of a scientific proposition

will have less and less uncertainty as long as new results, confirming or falsifying the original ones, are accumulated (Howson and Urbach, 1989).

The history of science provides numerous examples of how accepted findings have been challenged by subsequent research, and of how improvements and corrections in some cases, and full-blown controversies in others, have conducted to a better understanding of a given phenomenon. The Copernican revolution benefitted from and was refined also by critiques to some of its aspects, even if coming from scholars who were, overall, claiming wrong theories; for example, Tycho Brahe's observations about some inconsistencies in the Copernican view led to important improvements of it (Sherwood, 2011). More recently, for example, the body of research that led to identifying the causes and transimission mechanisms of HIV, its connection to AIDS, and to device treatments for the disease, proceeded through progressive criticisms and falsifications of earlier results, for example obtained with less reliable empirical strategies (from the analysis of individual occurrencies, to case-control studies, to randomized clinical trials; Holmberg, 2008). Also, theories and evidence on global warming, which are gaining more (though not definitive) consensus, are improved and refined also thanks to the counterarguments and evidence of scholars who are more skeptic about the anthropogenic nature of climate change (Eggers and Carpi, 2011; Sherwood, 2011). Building upon previous research and potentially identifying its limitations can therefore be seen as essential for a healthy working of the scientific community, as it allows science to self-correct and evolve.

But what are the incentives of scientists to perform these activities of controls and criticism? Will these activities always improve upon or correct previous findings, or, conversely, shall we expect some degree of imprecision at any given time? And what determines the incidence of imperfect science and of the effort to improve upon it and produce higher-quality research? In this paper, we propose a game-theoretic model to address these questions.

We study the interplay between the incentives to provide accurate, or reliable, results on the one hand, and the incentives to verify the validity of previous findings by peers. In the game, there are two players, a scientist producing a new scientific result, and another scientist who could verify that result. The quality of the new result depends on the amount of (costly) effort chosen by a scientist. A high-quality paper would stand against any scrutiny, whereas a low-quality paper presents flaws that are detected only if verification occurs; absent verification, high and low-quality paper are indistinguishable to the scientific community.

We show that a certain fraction of low-quality scientific knowledge characterizes all the equilibria of the basic version of model. In particular, if the benefits from verification are low (or the costs are high), verification does not occur and consequently low-quality papers are not identified. Conversely, if the incentives towards verification are high enough, then these activities are performed with positive probability, and the expected quality of research is higher. An implication of this result is that never observing low-quality research in a scientific occurs because of the lack of any verification activity and, as such, can be a source of concern. In the region of the parameters where verification occurs with positive probability, an increase in the identification of low-quality research goes together with an increase in the expected quality of research. In terms of normative implications, we suggest that increasing the benefit from verification (which may be obtained also giving more recognition to incremental research) should be considered as a primary way to improve scientific research reliability. Along these lines, in an extension of the model where scientists can obtain positive benefits also from confirmatory results (i.e. verifying high-quality papers), we show that low-quality research can be eliminated. Similarly, reducing the costs that scientists incur when verifying the results of others increases research quality. This highlights an important role for "incremental" research aimed at reinforcing, limiting, or even just confirm previous findings. In contrast, reducing the value of a publication for the knowledge originator, as some scholars have suggested (for example by softening the "publish or perish" paradigm), does not have an impact on research quality, although it increases the fraction of low-quality papers that are identified.

Our paper contributes to the theoretical literature in the economics of science that focuses on understanding the working of the scientific community with a particular attention to scientists' motivations (Dasgupta and David, 1994; Aghion et al., 2008). Two early contributions that analyze replication activities through the lens of a formal model are Mirowski and Skivas (1991) and Wible (1998). Mirowski and Skivas analyze the interaction between an originator of knowledge and a potential replicator, plus a set of potential extenders. In their model, (exact) replication never occurs unless editors require the originator to reveal a high enough level of information about their work, whereas extensions are more likely to occur in equilibrium. Wible shows an application of Becker's consumption-production theory to the time allocation of a scientist into genuinely replicable articles and seemingly replicable articles, the former being undistinguishable from the latter but more costly to produce. In general some non-replicable research will be produced in equilibrium. Although in a different way, both Mirowski and Skivas and Wible make the extent to which research is replicable endogenous. With respect to these papers, our work makes a contribution in two directions. First, we allow that the scientist himself may be ex-ante uncertain about the quality of his work, while at the same time controlling (in part) the quality level by the choice of effort level. In this way, we enrich the nature of the strategic interaction among the scientists playing different roles in the scientific community. Second, we perform an explicit analysis of the determinants of research quality, which allows us also to investigate the likely effects of the various interventions that have been proposed to increase the quality and reliability of research.

The model shares some features with Lacetera and Zirulia (2011), who analyze the incentives for committing and detecting fraudulent research, and derive the likelihood for fraudulent articles to be submitted, published, and not be caught. In that paper it is assumed that the project's probability of success is exogenous: in case of an unsuccessful project, the scientist can nevertheless submit a paper, thus committing a fraud. Here the probability of a paper being of high quality that is endogenous, because it depends on the scientist's effort. This different assumption significantly changes the nature of the game, as well as the results.

Finally, our paper relates to the broader literature on information search in sender-receiver

games. For instance, Henry (2009) and Henry and Ottaviani (2013) consider models where an agent can collect information, and then use it to try to influence the principal's decision. In Henry (2009), the quantity of research performed by the agent determines the amount of informative signals that she obtains, which then can be selectively reported to the principal. Henry and Ottaviani consider a dynamic model in which the agent decides the amount of search, and the principal fixes an approval standard. The main application of their framework concerns the case of a pharmaceutical firm collecting information about the effectiveness of a drug, with a regulatory authority as a principal. In both papers, principal-agent models are analyzed, in which diverging preferences about the true state of the world is a key element, differently from our analysis.

The paper is organized as follows. In the next section we present the basic model and discuss its assumptions. Section 3 reports the equilibria of the game and their consequences in terms of the quality of the produced research. Extensions of the model are proposed in Section 4, followed by a discussion of the model normative implications in Section 5. Section 6 concludes by advancing further applications of our framework as well as avenues for future research. All proofs are in the Appendix.

#### 2 The model

#### 2.1 The basic game

There are two players: the scientist (S) and the colleague (C). The scientist S (he) is the originator of a new scientific result, i.e. the producer of some new knowledge, which, for simplicity, we assume to be published as a journal article. The colleague C (she) decides whether to undertake activities to verify the quality of S's work. Through his choice of effort, S affects the quality of the knowledge that he produces, which can be high or low. A high-quality paper is a paper that, if scrutizined by C, does not show errors or significant lack of robustness. Otherwise, the paper is low quality. Absent C's attempt of verification, high quality and low quality papers are indistinguishable both for S and C, and thererefore

provide the same benefit to S and C.

More formally, S can choose between high effort  $(e_H)$  and low effort  $(e_L)$ . If S chooses  $e_H$ , then the paper is of high quality with probability 1; if S chooses  $e_L$ , then the paper is of high quality with probability  $p \in [0,1)$ .  $e_H$  and  $e_L$  denote both the feasible actions for S, and their associated costs, with  $e_H \geq e_L \geq 0$ . C can choose between verifying the quality of the results by S, which we denote as action v, or not verify (action nv). If C chooses v, then it bears a cost  $\beta \Delta e = \beta(e_H - e_L) \geq 0$ , where  $\beta \geq 1$  is a parameter that indicates that the verification cost for C will be generally higher than S's additional cost to produce a high-quality paper, i.e.  $\Delta e$ . Following v, the uncertainty concerning the quality of the paper is fully resolved. For S, the benefit obtained when C plays nv, or when she plays v and the paper is of high quality, is  $B_S$ ; the benefit is 0 when C plays v and the paper turns out to be of low quality. C obtains a positive benefit  $B_C \leq B_S$  when she plays v and the paper is low quality, and 0 otherwise. Both players are risk-neutral.

Throughtout the paper, we shall assume that the effort choice by S is not observed by C. Therefore, this is a game of imperfect information akin to a simultaneous-move game, with Nash equilibrium as a solution concept. The game is presented in normal form in Table 1.

#### 2.2 A discussion of the assumptions

In this section we provide a discussion of the key assumptions of the model, and the interpretation of the key parameters to which we will refer when commenting upon the results.

First, the model assumes that S can obtain a high-quality paper (and the corresponding benefit) with certainty if he exerts high effort. That high effort excludes low quality papers just simplifies the analysis by allowing us to focus on our main point, i.e. that the reliability of a scientific result is endogenous, depending on the prevailing incentives in the scientific community. More importantly, the model represents a view of science as a process of search for the "true state of the world", in which high (low) effort yields a perfect (imperfect) signal and S and C are indifferent with respect to the true state. In other words we exclude bias, both of S and C, in favor or against a specific scientific result, e.g. a positive result confirming

a theory or a negative result rejecting it.

Second,  $B_S$  is the value of a publication for which low quality was not identified. An implication of this assumption is that S does not obtain any additional benefit if his paper is scrutinized and turns out to be of high quality. For example, if this is the case, the work of C is less likely to be published, and does not reduce the impact of S's paper. A second implication is that S does not take into account the expected quality of the paper, which he knows given the effort exerted. Thus, any intrinsic reward from high quality that S can get is not considered here, although the same effect may be captured by a lower value of the cost  $e_H$ . In our interpretation, the value of  $B_S$  can be seen as primarily influenced by the prestige of the journal where the research is published, by the institutional context (such as "the publish or perish" attitude) or by personal characteristics of S, such as at the career stage.

Third, notice that  $B_S$  does not depend on effort. Therefore, higher effort does not lead to "better" scientific results, e.g. results that are more general or more relevant in some dimensions and which could lead to more cited publications, or appearing in more prestigious journals (Ellison, 2002). In our model, higher quality is associated to a characteristic of research, i.e. its reliability, which becomes evident only when the paper is scrutinized. High effort by S, then, should be rather been interpreted as "internal replication" (Hamermesh, 2007) and for that reason we will refer to  $\Delta e$  (which appears in both S and C payoffs) as verification costs.

As for the modeling of player C, we first notice that the notion of verification that we use to denote her action should be intended broadly. First, it includes direct replication. Second, it may take the form of design replication, whereby an alternative research design is used to answer the same questions (Muma, 1993). Finally, it applies also to conceptual (Wible, 1998) or scientific (Hamermesh, 2007) replication, where a different experiment, or analysis, is conducted, but in a way that might inform on the solidity of the original result. What these forms of verification have in common is that they tend to guarantee a reward to the

replicator only if they negatively affect the validity or applicability of the original research, thus potentially affecting (to some degree) the benefits of author of the original work. Direct replication is rarely observed; design and conceptual replications are more common, with the latter being often in the form of incremental research, i.e. additional inquiry on a given topic as opposed to the undertaking of an entirely new subject of analysis. As a second remark on C, the parameter  $\beta$  plays a particularly important role, as it measures the extent to which verification is more costly for C than the "internal" verification by S. Values of  $\beta$  greater than 1 can be justified, for example, by the esistence of some private information about the project that make it simpler for S to perform additional checks (Collins, 1985). At the same time,  $\beta$  is a parameter that can be at least partially affected by the rules of the scientific community. For instance, journal policies that favor the access to the original data have the effect of reducing  $\beta$ , although the "manipulability" of this parameter may be limited by constraints such as the proprietary nature of data. As for  $B_C$ , the benefit of discovering a low-quality paper may come from publication and visibility. Finally, assuming that the paper quality is known with certainty after C's replication excludes from the analysis the fact that replication activities themselves are subject to uncertainty with respect to their reliability.

#### 3 Results

#### 3.1 The Nash equilibria of the game

For all values of the parameters, the game has a unique Nash equilibrium, which is either in pure strategies or mixed strategies according to different parameter values. The pure strategy equilibrium displayes low effort and no verification, whereas in the mixed-strategy equilibrium there is a positive probablity of performing high effort and of verifying a paper. This is formalized in the following proposition.

**Proposition 1** The game has a unique Nash equilibrium. i) If  $\Delta e > \frac{(1-p)B_C}{\beta}$ , then the unique, pure-strategy Nash equilibrium is  $(e_L; v)$ ; ii) if  $\Delta e \leq \frac{(1-p)B_C}{\beta}$ , then the Nash Equilibrium is in mixed strategies, with S playing  $e_H$  with probability  $1 - \frac{\beta \Delta e}{(1-p)B_C}$ , and C playing

check with probability  $\frac{\Delta e}{(1-p)B_S}$ .

## 3.2 Implications: the inherent presence of low-quality science, and the share of low and high-quality work

A first implication of Proposition 1 is that a situation in which low-quality papers have zero probability to be produced is never an equilibrium. In other words, errors, flaws, limitations and and other forms of unreliable or incomplete results are a distinctive feature of the scientific endeavor as captured by our stylized model. If the incentives towards verification are low  $(\Delta e > \frac{(1-p)B_C}{\beta})$ , low-quality papers are not identified (because verification does not occur), and constitute a fraction (1-p) of all papers. Conversely, if the incentives towards verification are high enough  $(\Delta e \leq \frac{(1-p)B_C}{\beta})$ , then verification activities are performed with positive probability: in turn, this leads S to exert high effort with positive probability.

From Proposition 1, we can also determine the likelihoods of two events that will constitute the subject of our comparative exercises below: i) the probability that a paper is of high quality, which is a measure of actual reliability of scientific knowledge (independently of what is observed); and ii) the probability that papers are of low quality and they are identified as such. In determining such probabilities we impose  $\beta = 1$ ; this is without loss of generality because what is relevant here is only the ratio  $\frac{B_C}{\beta}$ .

**Proposition 2** The probability that the paper is of high quality is:

$$\Pr(high\ quality) = \begin{cases} 1 - \frac{\triangle e}{B_C} \ if \ 0 \le \triangle e \le (1 - p)B_c \\ p \ if \ \triangle e > (1 - p)B_c \end{cases}$$
 (1)

The probability that a low quality paper is identified is:

$$\Pr(low\ quality\ and\ identified) = \begin{cases} \frac{(\triangle e)^2}{(1-p)B_SB_c} & if\ 0 \le \triangle e \le (1-p)B_c \\ 0 & if\ \triangle e > (1-p)B_c \end{cases} . \tag{2}$$

Figures 1 and 2 report Pr("high quality") and Pr(low quality and identified) as a function of  $\triangle e$ , for different values of  $(1-p)B_c$ . The probability that a paper is of high quality is non-increasing in the verification cost  $\triangle e$ . If  $\triangle e$  is large (with respect to C's

expected gain from verification), then no verification occurs, and the fraction of high-quality papers only depends on the exogenous probability p. If  $\triangle e$  is low (with respect to C's expected gain from verification), then the lower  $\triangle e$ , the larger the fraction of high quality papers because exerting higher effort is less costly for S. Note, however, that verification activities by C, although being less costly, are less frequent because the probability to find a low-quality paper is smaller.

We also notice that, in the mixed-strategy equilibrium region, larger benefits  $B_C$  from identifying a low-quality paper (or a lower cost via a reduction in  $\beta$ ) increase the fraction of high quality papers: because verification is more rewarding (or less costly), S must increase his effort in order to reduce C's incentives to verify. However, the probability that the paper is of high quality does not depend on  $B_S$ , i.e. the value of a publication. To understand the intuition for this, consider that for a given intensity of C's control (i.e. for given value of r), the marginal effect of  $B_S$  on the S's payoff is 1 when he exerts high effort, and rp+(1-r)<1 when he exerts low effort, because in this case S must take into account that the value of publication is lost if the papers turns out to be of low quality and it is identified as such. In that respect, then, high effort, and consequently, high-quality papers become more attractive because the cost of losing the publication value is larger. However, as a consequence of this, C responds to the increase in  $B_S$  by lowering the intensity of her (costly) verification activity, making  $e_L$  more attractive until the point in which S is again indifferent between high and low effort.

We next look at the probability that a paper is of low quality and is identified as such. First, notice that this probability is positive when  $\triangle e$  is "low" and zero when  $\triangle e$  is "high". Also, the probability of high quality papers if unambigously higher when some low quality papers are identified than when no low quality papers are discovered. In other words, the absence in a field of scientific results that are found to be of low quality (false, floawed, limited, or less relevant that initially believed), rather than a signal of the absence of these types of papers, must interpreted as the lack of any verification activities activity and then

as a source of concern. As we will see in the extension of the model discussed in Section 4, this may not hold only if, somewhat contrary to the current structure of incentives in science, confirmatory results, i.e. coming from the verification of a high quality paper, provide a large benefit to C.

However, in the region of mixed-strategy equilibrium, lower verification costs imply both higher expected quality and lower rate of low-quality discovery; this occurs because both low quality papers and verification activities are less frequent. The same effect occurs for an increase in  $B_C$ , but in this case this is due to the increase effort by S. The probability of verification by C, indeed, does not depend on  $B_C$ : an increase in the reward from discovery a low quality paper operates as a threat that "discipline" S's behavior, leaving the intensity of the verification activity unaffected. As a consequence, observing a reduction of the frequency of low quality papers that are discovered (but not their disappearance!) is a good sign, because it is unambigously associated to higher expected quality of research.

Note, finally, that the probability of identifying low-quality papers is negatively affected by  $B_S$ . Therefore, an increase in the value of publication reduces the probability that low-quality papers are recognized as such, but without affecting the probability that such papers are produced. The intuition is that, when publications are more valuable, the opportunity cost of low effort is in fact higher: as a consequence, C may save on verification activities while leaving S indiffferent between high and low effort.

### 4 Extension: a value for confirmatory results

The model can extended in several directions. In this section we briefly discuss a case where C can obtain a positive reward even when verifying a paper that turns out to be of high quality. In this case, the normal form of the game is expressed by the payoff matrix in Table 2, where  $B_C^H$  ( $B_C^L$ ) corresponds to the value for C of discovering a high (low)-quality publication, with  $B_C^H \leq B_C^L$ . Nash equilibria of the game are summarized by the following proposition.

**Proposition 3** The equilibrium of the game is unique. In particular, i) if  $\Delta e > \frac{pB_C^F + (1-p)B_C^L}{\beta}$ ,

the Nash equilibrium is  $(e_L, nv)$ ; ii) if  $(1-p)B_S < \Delta e < \frac{pB_C^H + (1-p)B_C^L}{\beta}$ , then the Nash equilibrium is  $(e_L, v)$ ; iii) if  $\Delta e \leq \min\left\{(1-p)B_S; \frac{B_C^H}{\beta}\right\}$ , then the equilibrium is  $(e_H, v)$ ; iv) finally, if  $\frac{B_C^H}{\beta} < \Delta e < \min\left\{\frac{pB_C^H + (1-p)B_C^L}{\beta}; (1-p)B_S\right\}$ , then the mixed strategy equilibrium has  $r^* = \frac{\Delta e}{(1-p)B_S}$  and  $q^* = \frac{pB_C^H + (1-p)B_C^L}{(1-p)(B_C^L - B_C^H)} - \frac{\beta \Delta e}{(1-p)(B_C^L - B_C^H)}$ .

The main insight from this extension is that allowing C to gain utility from the verification of high-quality research enlarges the set of possible equilibria. In particular, when  $\Delta e$  is small enough, C can prefer to verify even if the research by S is of high quality with probability 1. Thus, if confirmatory results are positively valued by the scientific community, it is possible that low-quality papers cease to exist. Moreover, for intermediate values of  $\Delta e$  (i.e. if  $(1-p)B_S < \Delta e < \frac{pB_C^H + (1-p)B_C^L}{\beta}$ )), the verification activity of C does not deter S from exerting low effort. In these two cases the expected quality of papers radically differs, being respectively the highest and the lowest possible quality in the model. In other words, verification is a necessary, but not a sufficient condition for high-quality research.

# 5 Improving the quality of scientific production: a normative perspective

In recent years, numerous discussions about the functioning of the scientific community have been undertaken, leading to various proposals for enhancing the quality and reliability of scientific research. Our model can be used to provide insights upon some of the proposals that have been advanced as well as about some of the current trends in the scientific community.

Several authors have identified in the lack of proper incentives towards replication, in its various forms as decribed above, including performing incremental research upon existing and "established" findings, a main driver of potential scientific unreliability. Recently, for example, the 2013 Medicine Nobel Laureate Randy Schekman announced that he would not send his papers to some of the major science journals, in particular because they excessively select "novel", "newsworhty" findings at the expense of rigor and depth fo inquiry, which is always given by additional work on an exisiting field (Schekman, 2013). Similar considerations

were expressed in the past by other prominent scientists. The results of our model are in line with these views. Increasing the expected reward of verification activities (via an increase in  $B_C$ ) or reducing its cost (via a reduction in  $\beta$ ) increases the expected quality of scientific knowledge. This does not happen because verification is indeed more frequent, but because knowledge originators exert higher effort in response. The analysis in Section 4 qualifies this claim, by showing that a positive utility from confirmatory results (which, to a large extent, does not characterize the current incentive structure of the scientific community) might lead to the reduction of low-quality papers, while keeping verification activities in place. In terms of discovery of low-quality research, our model suggests that we should expect an increase, if starting from a situation when verification activities were not performed; or a decrease otherwise.

Another frequent belief, sometimes considered equivalent to the one just discussed, is that the quality of scientific research may be negatively affected by too high-powered incentives to publish; proposals have therefore been advanced to soften the "publish-or-perish" paradigm (Abelson, 1985; Giles, 2007; Schekman, 2013). Our model shows, however, that acting directly on the incentives to publish may be different from increasing the incentives for incremental or confirmatory research, or other forms of verification. In particular, the expected quality of research would be unaffected by softening publication incentives alone (a reduction in  $B_S$ ), as argued in Section 3.1. However, the fraction of low-quality papers that could be recognized as such would increase when publication incentives are weaker because of a more intense control activity by C. This will be socially valuable as well, since it will reduce the uncertainty concerning scientific research quality. Also, a reduction in the publish-and-perish attitude can be interpreted as a reduction of the relative value of (supposedly) path-breaking research with respect to more incremental research, causing a simultaneous decrease in  $B_S$  and increase in  $B_C$ . From previous discussions, this would simultaneously increase research quality and the identification of low quality research.

#### 6 Conclusions

Although simple, our model conveys a number of relevant insights about the operating of the overall scientific endeavor, and to clarify how different rules and incentives affect the quality and reliability of scientific production. The basic mechanisms analyzed here, in particular, suggest that not only are scientific findings never complete or definitive and are always prone to improvement; but, also, that observing only apparently definitive or undisputed findings may be a sign of weakness of a scientific field rather than a proof of its solidity. Key driving forces in the model are the incentives to produce new research on the one hand, and the incentives provide further work upon and, in the process, possibly question existing and established results. We also show an interesting asimmetry of effects between lowering incentives to produce new research, and increasing incentives to do additional work and verification on existing findings.

This framework can also be applied to other environments characterized, like the scientific community, by the possibility of producing both new content and contribution to existing finding on the one hand, and by peer scrutiny on the other hand. One example is given by the news industry. Newsmakers are constantly in search of new facts and storied to report, however multiple reporting on a given story or fact-checking is considered essential to enhance the reliability and credibility of news. Another relevant example is the open source movement. Software developers in open source environments produce new code while building upon and checking existing programs; one of the frequently highlighted strengths of open-source software is that marginal improvements and corrections can be made more easily and quickly (Lakhani and Von Hippel, 2003). Understanding the incentives of different actors to produce new material versus work on existing one, and how different institutional arrangments affect them, is of relevance in these environments too.

There are a number of avenues for further extensions of the model. Within the current model structure, the relaxation fo some assumptions may lead to interesting developments. For instance, we could allow S to get a benefit from confirmation of his results by C. Similarly,

we could consider the case of observable effort, and then compare the results with those obtained in Section 3. Finally, we could introduce scientists' bias in favor or against certain hypotheses or findings.

#### References

- [1] Abelson, P. (1990). Mechanisms for evaluating scientific information and the role of peer review. Journal of the American Society for Information Science, 41(3), 216-222.
- [2] Aghion, P., Dewatripont, M., & Stein, J. C. (2008). Academic freedom, private-sector focus, and the process of innovation. *The RAND Journal of Economics*, 39(3), 617-635.
- [3] Azoulay, P., Furman, J. L., Krieger, J. L., & Murray, F. E. (2012). Retractions (No. w18499). National Bureau of Economic Research.
- [4] Broad, W., & Wade, N. (1982). Betrayers of the Truth. Fraud and Deceit in the Hall of Science. Simon and Schuster, New York
- [5] Carpi, A., & Egger, A. E. (2011). The Nature of Scientific Knowledge. Visionlearning.
- [6] Collins, H. (1992). Changing Order: Replication and Induction in Scientific Practice. Chicago: The University of Chicago Press.
- [7] Dasgupta, P., & David, P. A. (1994). Toward a new economics of science. Research Policy, 23(5), 487-521.
- [8] Ellison, G. (2002). Evolving standards for academic publishing: a q-r theory. *Journal of Political Economy*, 110 (5),994-1034.
- [9] Giles, J. (2007). Breeding cheats. *Nature*, 445(18): 242–243
- [10] Hamermesh, D. S. (2007). Viewpoint: Replication in economics. *Canadian Journal of Economics*, 40(3), 715-733.
- [11] Hartemink, A. E. (2000). Publish or Perish (3)—Fraud and ethics. *Bull Internat Union Soil Sci*, 97, 36-45.
- [12] Henry, E. (2009). Strategic disclosure of research results: the cost of proving your honesty. *The Economic Journal*, 119, 1036-1064.
- [13] Henry, E., & Ottaviani, M. (2013). Research and the approval process. Mimeo.
- [14] Holmberg, S. D. (2008). Scientific Errors and Controversies in the HIV/AIDS Epidemic. Greenwood Publishing Group.
- [15] Howson, C., & Urbach, P. (1989). Scientific Reasoning: The Bayesian Approach. Open Court Publishing Co.

- [16] Hull, L. D. (1988), Science as a Process: An Evolutionary Account of the Social and Conceptual Developments of Science. Chicago: The University of Chicago Press.
- [17] Lacetera, N., & Zirulia, L. (2011). The economics of scientific misconduct. *Journal of Law, Economics, and Organization*, 27(3), 568-603.
- [18] Lakhani, K. R., & Von Hippel, E. (2003). How open source software works: "free" user-to-user assistance. Research policy, 32(6), 923-943.
- [19] Lu, S. F., Jin, G. Z., Uzzi, B., & Jones, B. (2013). The retraction penalty: Evidence from the Web of Science. *Scientific Reports*, 3, 3146.
- [20] Mirowski, P., & Sklivas, S. (1991). Why econometricians don't replicate (although they do reproduce). Review of Political Economy, 3(2), 146-163.
- [21] Muma, J. R. (1993). The need for replication. *Journal of Speech, Language and Hearing Research*, 36(5), 927.
- [22] Popper, K. R. (1959). The Logic of Scientific Discovery. London: Hutchinson.
- [23] Schekman, R. (2013). How journals like Nature, Cell and Science are damaging science. *The Guardian*.
- [24] Sherwood, S. (2011). Science controversies past and present. Physics Today, 64(10), 39-44.
- [25] Stephan, P. (2013), How Economics Shapes Science. Harvard University Press, Cambridge, MA, USA.
- [26] Wible, J. R. (2004). The Economics of Science: Methodology and Epistemology as if Economics Really Mattered. Routledge.

$$S \quad e_{H} \quad \begin{matrix} v & nv \\ B_{S} - e_{H}; 0 - \beta \Delta e \\ e_{L} \end{matrix} \quad \begin{matrix} B_{S} - e_{H}; 0 \\ pB_{S} - e_{L}; (1 - p)B_{C} - \beta \Delta e \end{matrix} \quad \begin{matrix} B_{S} - e_{H}; 0 \\ B_{S} - e_{L}; 0 \end{matrix}$$

Table 1: The basic game in normal form

$$S = e_{H} \begin{bmatrix} v & nv \\ B_{S} - e_{H}; B_{C}^{H} - \beta \Delta e & B_{S} - e_{H}; 0 \\ e_{L} & B_{S} - e_{L}; pB_{C}^{H} + (1 - p)B_{C}^{L} - \beta \Delta e & B_{S} - e_{L}; 0 \end{bmatrix}$$

Table 2: Normal form of the game extended to the case of positive reward to C from verifying a high-quality paper

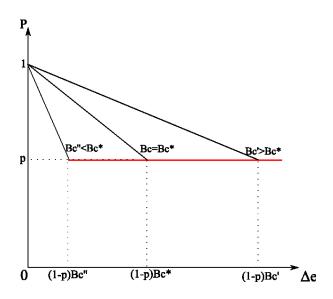


Figure 1: Probability of a high-quality paper

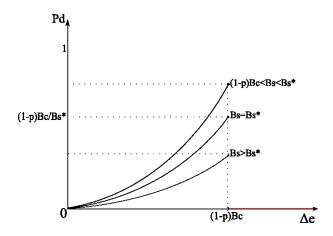


Figure 2: Probability that a low-quality paper is identified.

#### A Proofs

**Proof or proposition 1.** To see that  $(e_L; nv)$  can be a Nash equilibrium, notice that  $e_L$  is the best response to nv. If S chooses  $e_L$ , then C prefers nv if  $(1-p)B_C - \beta \Delta e < 0$ , i.e.  $\Delta e > \frac{(1-p)B_C}{\beta}$ . Note, in contrast, that the strategy pair  $(e_L, v)$  is never an equilibrium; if S chooses  $e_L$ , then C prefers v if  $(1-p)B_C - \beta \Delta e > 0$ , i.e.  $\Delta e < \frac{(1-p)B_C}{\beta}$ . However, in order for S to play  $e_L$  in response, the condition is that  $B_S - e_H < pB_S - e_L$  or  $\Delta e > (1-p)B_S$ . Because by assumption  $B_C \leq B_S$ , the two conditions cannot be simultaneously satisfied. Finally, to see that pure equilibria involving high effort do not exist, notice that C's best response to  $e_H$  is no check, but S's best response to no check is  $e_L$ .

As for the mixed strategy equilibrium, we denote with q the probability that S plays  $e_H$ , and r as the probability that C plays v. For S to be indifferent between  $e_H$  and  $e_L$  it must be:

$$B_S - e_H = r(pB_S - e_L) + (1 - r)(B_S - e_L)$$

from which we obtain:

$$r^* = \frac{\Delta e}{(1-p)B_S} \tag{3}$$

For C to be indifferent between v and nv it must be that:

$$-q\beta\Delta e + (1-q)((1-p)B_C - \beta\Delta e) = 0$$

from which we derive:

$$q^* = 1 - \frac{\beta \Delta e}{(1-p)B_C}. (4)$$

 $r^*$  is always positive, and  $r^* \leq 1$  if  $\Delta e \leq (1-p)B_S$ .  $q^* \leq 1$  for all parameter values, and it is strictly positive if  $\Delta e \leq \frac{(1-p)B_C}{\beta}$ . Because  $B_C \leq B_S$ , a mixed-strategy equilibrium exists if  $\Delta e \leq \frac{(1-p)B_C}{\beta}$ .

**Proof of Proposition 2** In equilibrium, a paper is of high quality with probability 1 if S exerts high effort, and with probability p if he exerts low effort. Therefore,  $Pr(high\ quality) =$ 

 $q^* + (1 - q^*)p$ . The discovery of low-quality papers occurs if i) S exerts low effort, ii) the paper is actually of low quality; and iii) C chooses to verify. The corresponding probability is  $\Pr(low\ quality\ and\ verified) = (1 - q^*)(1 - p)r^*$ .

**Proof of Proposition 3** First, we determine the existence of pure strategy equilibria. If S chooses  $e_H$ , C will play v as long as  $B_C^H - \beta \Delta e > 0$ , i.e.  $\Delta e < \frac{B_C^H}{\beta}$ , and nv otherwise C. If S plays  $e_L$ , C will choose to check if  $pB_C^H + (1-p)B_C^L - \beta \Delta e > 0$ ; i.e.  $\Delta e < \frac{pB_C^F + (1-p)B_C^L}{\beta}$ , and nv otherwise. If C plays v, S will choose  $e_H$  if  $B_S - e_H > pB_S - e_L$ , i.e.  $\Delta e < (1-p)B_S$ , and  $e_L$  otherwise. Finally if C does not verify, S will always choose  $e_L$  since  $B_S - e_H < B_S - e_L$ . Therefore the possible equilibria are  $(e_H, v)$  if  $\Delta e < \min\left\{(1-p)B_S; \frac{B_C^H}{\beta}\right\}$ ,  $(e_L, v)$  if  $\Delta e > \frac{pB_C^F + (1-p)B_C^L}{\beta}$ ,  $(e_L, v)$  if  $(1-p)B_S < \Delta e < \frac{pB_C^H + (1-p)B_C^L}{\beta}$ . As it concerns the mixed strategy equilibrium, we denote with r the probability for C to play v and with q the probability for S to play  $e_L$ . The indifference condition for S is:

$$B_S - e_H = r(pB_S - e_L) + (1 - r)(B_S - e_L)$$

from which we obtain:  $r^* = \frac{\Delta e}{(1-p)B_S}$ . For C to be indifferent between v and nv it must be:

$$q(B_C^H - \beta \Delta e) + (1 - q)(pB_C^H + (1 - p)B_C^L - \beta \Delta e) = 0$$

which yields:  $q^* = \frac{pB_C^H + (1-p)B_C^L}{(1-p)(B_C^L - B_C^H)} - \frac{\beta \Delta e}{(1-p)(B_C^L - B_C^H)}$ . For  $r^*$  and  $q^*$  to be within the unit interval we need respectively i)  $\Delta e \leq (1-p)B_S$  ii)  $\frac{B_C^H}{\beta} < \Delta e < \frac{pB_C^H + (1-p)B_C^L}{\beta}$ , therefore the equilibria exist for  $\frac{B_C^H}{\beta} < \Delta e < \min\left\{\frac{pB_C^H + (1-p)B_C^L}{\beta}; (1-p)B_S\right\}$ .