NBER WORKING PAPER SERIES

ABOVE A SWAMP: A THEORY OF HIGH-QUALITY SCIENTIFIC PRODUCTION

Bralind Kiri
Nicola Lacetera
Lorenzo Zirulia

Working Paper 21143
http://www.nber.org/papers/w21143

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2015

We are grateful to the pe participants to several seminars (University of Bologna, University of Florence, TOBB University) and conferences (2014 Izmir Workshop on Scientific Misconduct and Research Ethics in Economics, 2014 EARIE Meetings, 2014 BRICK Conference, 2014 DRUID Conference, 2015 Barcelona GSE Summer Forum, 2015 Italian Economic Association Conference, 2015 REER Conference, and 2016 Annual Meeting of the French Economic Association) for their feedback. Astrid Marinoni provided outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Above a Swamp: A Theory of High-Quality Scientific Production
Bralind Kiri, Nicola Lacetera, and Lorenzo Zirulia
NBER Working Paper No. 21143
May 2015, Revised August 2016
JEL No. L31,O31

ABSTRACT

We elaborate a model of the incentives of scientists to perform activities of control and criticism when these activities, just like the production of novel findings, are costly, and we study the strategic interaction between these incentives. We then use the model to assess policies meant to enhance the reliability of scientific knowledge. We show that a certain fraction of low-quality science characterizes all the equilibria in the basic model. In fact, the absence of detected low-quality research can be interpreted as the lack of verification activities and thus as a potential limitation to the reliability of a field. Incentivizing incremental research and verification activities improves the expected quality of research; this effect, however, is contrasted by the incentives to free ride on performing verification if many scientists are involved, and may discourage scientists to undertake new research in the first place. Finally, softening incentives to publish does not enhance quality, although it increases the fraction of detected low-quality papers. We also advance empirical predictions and discuss the insights for firms and investors as they "scout" the scientific landscape.

Bralind Kiri
TOBB University of Economics and Technology
Söütözü Mh., Söütözü Cd No:43
06520 Söütözü/Ankara
Turkey
bkiri@etu.edu.tr

Nicola Lacetera
University of Toronto
Institute for Management and Innovation
3359 Mississauga Road, Room KN 235
Mississauga, ON L5L 1C6
CANADA
and NBER
nicola.lacetera@utoronto.ca

Lorenzo Zirulia
University of Bologna
Department of Economics
Strada Maggiore, 45
Bologna, Italy
lorenzo.zirulia@unibo.it

> 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 production of reliable and high-quality scientific research is valuable not only within the ivory tower of academia. Firms and investors, for example, assess opportunities also on the basis of the science underlying a new product, process or service, and "scout" the scientific landscape in search for discoveries that are scientifically sound and commercially promising (Baum and Silverman, 2004; Merck, 2015; Pfizer, 2015; Ryan, 2013). More broadly, scientific knowledge is a powerful engine of economic growth and social welfare (Romer, 1990; Stephan, 2012).

For this reason, the debate about the reliability of research involves not only the scientific community, but also firms, policymakers and the public opinion. According to several accounts, science is currently undergoing a "reproducibility crisis" (Allison et al., 2016; The Economist, 2013). In psychology, for example, a project attempting to replicate 100 studies succeeded only in 39 cases (Open Science Collaboration, 2015). Begley and Ellis (2012) reported that they could replicate only 6 out of 53 fundamental studies in oncology and haematology, and in a meta-analysis of genetic associations studies, Ioannidis et al. (2001) found that the results of the first study, often suggesting a stronger genetic effect, correlated only weakly with subsequent research. The social and economic costs of this lack of reliability may be substantial; according to Freedman et al. (2015), for example, every year 28 billion dollars are spent in the US on preclinical research that is not reproducible.

Science may "go wrong" for outright fraud or mistakes that, if major and detected, lead to retraction from publication (Azoulay et al., 2015a-b; Broad and Wade, 1982; Lacetera and Zirulia, 2011; Lu et al., 2013). Incentives prevailing in scientific communities, such as the "publish or perish" imperative (Abelson, 1990; Giles, 2007), are often blamed for inducing to frauds or grave inaccuracies. In a less pessimistic view, however, flaws, limitations and mistakes in a study just occur as "natural" steps toward better theories and findings (Aschwanden, 2015). 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 includes findings that are limited or flawed in some ways, with corrections and improvements occurring as long as new results, confirming or falsifying the original ones, are accumulated (Howson and Urbach, 1989). Building upon previous research and potentially identifying its limitations thus appears essential for a healthy working of the scientific community (Carpi and Eggers, 2011).

The history of science provides many examples of how subsequent research challenged accepted findings. In some cases, improvements and corrections (or sometimes full-blown controversies) led to a better understanding of a given phenomenon. For instance, the Copernican revolution benefited from and was refined by critiques to some of its aspects, even if those critiques were based on wrong theories, such as Tycho Brahe's observations about inconsistencies in the heliocentric view (Sherwood, 2011). In other cases, such as the research on HIV and AIDS, advances occurred through progressive criticisms and falsifications of earlier results, for example obtained with less reliable empirical strategies (Holmberg, 2008). In climatology, there is increasing agreement about the anthropogenic nature of climate change. However, counterarguments and evidence of scholars who are more skeptic are contributing to improve the overall reliability of research in this area (Sherwood, 2011). Critical views may be particularly valuable when scientific results attract media attention, as was recently the case in paleoanthropology following the discovery of Homo Naledi (Lents, 2016).

In other instances, research that built on previous work led to discarding that earlier work entirely; examples include polywater and cold nuclear fusion (Rousseau and Porto, 1970; Taubes, 1993). Livio (2014) describes "blunders" by some great scientists. Darwin's theory of evolution, for example, presented in its initial versions some flaws that were pointed out by Fleeming Jenkin, a Scottish engineer, with this critique containing, in turn, some limitations as subsequently reported by Arthur Sladen Davis; the contributions of Linus Pauling, the Nobel Laureate for Chemistry in 1954 (and for Peace in 1962), to the definition of the DNA structure were soon identified as flawed by Crick and Watson. Catalini et al. (2015) find that articles in immunology receiving "negative" citations (i.e. citations that criticize or limit the validity of a study) tend to be highly cited overall and therefore more prominent and relevant; in turns, papers making negative citations are not marginal (again as measured by overall citations).[1]

Based on these premises, this paper proposes a game-theoretic analysis of the interplay between the incentives to exert scientific effort and provide accurate results on the one hand, and the incentives to verify the validity of previous findings on the other hand. With our model, we address the following positive and normative questions:

- What are the incentives of scientists to perform research on existing, established topics to potentially exert control and criticism?

[1] A more limited analysis that we conducted on 1,037 articles on climate change published in *Nature* (between 1975 and early 2015) and *Nature Climate Change* (2007 - early 2015) shows about 215 cases in which some papers were negatively cited.

- Will these activities always improve upon or correct previous findings, or shall we expect some degree of imprecision at any given time?
- How do these incentives interact with those to produce novel, high-quality findings?
- What factors determine the incidence of imperfect science?
- Which policy interventions could improve the reliability of science? Which policies, instead, would be ineffective or even counterproductive?

A first key result, derived from the basic version of the model described in Section 2 and solved in Section 3, is that a certain fraction of low-quality scientific knowledge characterizes all the equilibria of the game. Incentives to verify findings may be too low, thus reducing also incentives to perform high effort to produce reliable research; or they may be high enough to lead to verification with positive probability, and in turn, to the production of higher-quality research on average. An implication of this result is that never observing low-quality research in a scientific field may be due to a lack of verification activities and, as such, can be a source of concern rather than a signal of the solidity of a body of knowledge. Therefore, fields that display controversies and where flaws are pointed out may indicate greater health and promise than fields where no such activities are observed.

Although our result suggests that observing a certain fraction of flawed science may indeed be considered a natural and desirable feature, the identification of those characteristics associated to higher reliability allows to assess different policies meant to increase the overall reliability of science, as well as the ability to sort scientific results of different quality. We do this in Section 4, where we perform comparative statics exercises on the basic version of the model and we extend it in several directions.

We show, first, that 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. Conversely, reducing the costs (or increasing the benefits) for scientists to verify the results of others increases the overall expected quality of research. This finding highlights an important role for incremental research aimed at reinforcing, limiting, or even just confirming previous findings. We also identify, however, a few countervailing effects of enhancing verification activities. For example, less costly (or too frequent) verification activities may lead a scientist to not undertake a new, potentially socially valuable research project in the first place. Thus some level of "protection" of one's research (e.g. concerning policies for data sharing) might be desirable in certain cases.

Additional results concern the impact of the size of the scientific community. The performance of verification activities by a high number of scientists may lead to the overall reduction of these activities and of the expected quality of research if individual rewards from scrutiny are lower because they are shared among colleagues. Also, in scientific communities where interactions are repeated and frequent, scientists may "collude," i.e. avoid to verify each other's research and save on the investment required by expensive experimental procedures.

A final set of policies that we consider in the model regard the direct involvement of more prestigious journals in certifying research reliability. We derive that this may reduce reliability by crowding out the scientists' incentives to perform verification activities; in contrast, attention of journals to other aspects of quality would generally improve reliability.

In an effort to further bring the theoretical analysis to actual applications, in the concluding Section 5 we outline insights for companies and investors interested in exploring the scientific landscape for business opportunities, we propose a few empirical tests based on the model's prediction, and we extend the application of the framework to other contexts beyond the working of the scientific community.

Related literature Our paper is related to a few streams of literature. Two early contributions that analyze replication activities formally are Mirowski and Skivas (1991) and Wible (1998). Mirowski and Skivas analyze the interaction between an originator of knowledge and a potential replication, 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 indistinguishable from the latter but more costly to produce. In general some non-replicable research will be produced in equilibrium. Although in different ways, both studies 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 effects of the various interventions that have been proposed to increase the quality and reliability of research.

The model also shares some features with Lacetera and Zirulia (2011), who analyze the incentives to commit and detect fraudulent research, and derive the likelihood for fraudulent articles to be submitted, published, and not be caught. For instance, in the basic version of their model there is always a positive probability of fraudulent papers, as we also find here for low-quality papers. However, Lacetera and Zirulia assumed that the project's probability of success was exogenous; in case of an unsuccessful project, the scientist can nevertheless submit a paper, thus committing a fraud, at no additional cost. Here the probability of a paper being of high quality is endogenous, because it depends on the scientist's (costly) effort. This different assumption has an impact on the nature of the game and on the results.[2]

More broadly, this paper contributes to the stream of economic analyses of the operating of academia and the scientific community that has focused on such issues as scientists' motivations, the allocation of research projects between universities and companies, the choice between basic and applied research, the commercialization of science and the allocation of authority within universities (see for example Aghion et al., 2008; Banal-Estañol and Macho-Stadler, 2010; Dasgupta and David, 1994; Häussler et al., 2014; Jensen and Thursby, 2001; Lacetera and Zirulia, 2012; Macho-Stadler and Pérez-Castrillo, 2010; Masten, 2006; Mialon, 2010; Stern, 2004).

Our model, finally, is related to the literature on information search in sender-receiver games such as Henry (2009) and Henry and Ottaviani (2014). These papers adopt a principal-agent framework in which diverging preferences about the true state of the world is a key element, differently from our analysis.

2 The model

2.1 The basic game

There are two risk-neutral players: a scientist (S - he) and a colleague (C - she). The scientist S is the originator of a new scientific result, which we assume to be published. The colleague C decides whether to undertake activities to verify the quality of S's work. The quality of a paper can be high or low. A paper is of high quality if, when scrutinized by C, it does not show errors or significant lack of robustness. Otherwise, quality is low. Through his choice of effort, S affects the quality of the knowledge that he produces. Absent C's verification, high-quality and low-quality papers provide the same benefit to the players. The players do not observe each other's effort choice. Therefore, this is a game of imperfect information, with Nash equilibrium as solution concept. The payoff matrix in normal form is in Table 1.

S chooses between high effort (e_H) and low effort (e_L). If S chooses e_H, the paper is of

[2] Our model also relates to the game-theoretic treatments of other types of fraudulent behavior, such as tax evasion (Erard and Feinstein, 1994; Graetz et al., 1986; Reinganum and Wilde, 1986). A difference from these studies is that our model reflects some specificities of the scientific community. In the tax evasion literature, for instance, benefits and costs are pecuniary and the law enforcement agency maximizes total government revenues (net of audit costs), which are negatively related to the payoffs of tax payers. In science, both types of agents, the "producers" of new knowledge and the "auditors" belong to the same community, acting both as complementors and competitors, and their benefits are mostly nonpecuniary (reputational).

		C	
		v	nv
S	e_H	$B_S - e_H;\ 0 - \beta\Delta e$	$B_S - e_H;\ 0$
	e_L	$pB_S - e_L;\ (1-p)B_C - \beta\Delta e$	$B_S - e_L;\ 0$

Table 1: Payoff matrix of the basic game in normal form. In each cell, the first payoff is of player S, and the second payoff is of player C.

high quality with probability 1; if S chooses e_L, 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 chooses between verifying the quality of the results by S (action v), or not verify (action nv). If C chooses v, then she bears a cost $\beta\Delta e = \beta(e_H - e_L) \geq 0$; thus the verification cost for C is proportional to the additional cost for S to produce a high-quality paper. We will assume that β is not too low, i.e. $\beta \geq \frac{B_C}{B_S}$. Following the performance of v, the uncertainty concerning the quality of the paper is 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.

2.2 A discussion of the model's assumptions

Before we solve the basic game and explore its implications, it is worth discussing the key assumptions that we made and how they relate to the working of the scientific community.

First, the model assumes that S produces a high-quality paper 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 to effort.[3] In turn, effort is affected by the prevailing incentives in the scientific community. 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 S both if the paper is of high quality, and if it is of low quality and not identified. Thus, we do not consider any intrinsic reward from high quality that S may receive, although we could express intrinsic motivations through a lower cost e_H. In our interpretation, the value of B_S can be seen as primarily influenced by

[3] To refer back to an historical example mentioned above, the flaws in Pauling's approach to define the structure of DNA, as described by Livio (2014), were attributed by his collaborators to the fact that he just did not try hard enough and spent only little time on the problem.

the prestige of the journal where the research is published, by the institutional context (such as the "publish or perish" culture), or by characteristics of S (such as his career stage).

Third, the value of 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 relevant in some ways and that could lead to more cited publications, or appear in more prestigious journals (Ellison, 2002). In our model, higher quality is associated with a characteristic of research, i.e. its reliability. High effort by S can be interpreted as "internal replication" (Hamermesh, 2007) and for this reason we will refer to Δe as verification costs. Δe can be expected to be large in those fields, such as biomedical research or psychology, in which internal replication requires the repetition of the experiment, whereas it is likely lower in those cases where, for example, it is mostly performed through "robustness" analyses on the same data.

With regard to the modeling of player C, the notion of verification that we use to denote her action is to be intended broadly. First, it includes direct replication. Second, verification also occurs through design replication, whereby an alternative research design is used to answer the same questions (Muma, 1993). Third, the action v applies also to conceptual or scientific replication, where a different experiment or analysis is conducted, but in a way that might inform about the solidity of the original result (Hamermesh, 2007; Wible, 1998).[4] Finally, and more broadly, any form of "incremental" research, i.e. research that heavily builds on existing findings by offering only small advances, can be considered as a form that action v takes. What these activities have in common is that they tend to guarantee a reward to the replicator if they negatively affect the validity or applicability of the original research, thus potentially affecting (to some degree) the benefits of the author of the original work. Direct replication is rarely observed, often because the same exact conditions cannot be re-created or the original data are proprietary or too costly to be collected again; design and conceptual replications are more common, with the latter being often in the form of incremental research.

The parameter β measures the relative magnitude of verification costs for C with respect to internal verification by S. Modeling the verification costs of C with Δe (up to a multiplicative factor) might come across as a strong restriction and simplification. However, note that one could always re-parametrize verification costs in terms of the difference between the effort costs of research for S and a proportional factor. Furthermore, it makes sense to establish a simple comparison between the internal verification costs by S and the verification by an external peer. Values of β greater than 1 indicate, for example, the existence of some private information or tacit knowledge about the project that make it easier for S to perform additional checks (Collins, 1985). Values of $\beta < 1$ may occur in the case of theoretical results,

[4] Camerer et al. (2016), reporting the results of the "Experimental Economics Replication Project" (http://experimentaleconreplications.com/) is an example of this type of replications.

which can be invalidated by a single counterexample or by identifying a logical error in the proof. β can also depend on the rules of the scientific community. For instance, policies that favor the access to the original data (if feasible) have the effect of reducing β.

As for the benefit of discovering a low-quality paper (B_C), it may come from publication and visibility as well as from a direct utility from sustaining the quality of science. Note also that assuming that the quality of the original paper is known with certainty after C's verification excludes uncertainty around the success of the verification activities.[5]

3 The equilibria of the game, and the inherent presence of low-quality research

The game has a unique Nash equilibrium in either pure or mixed strategies according to different parameter values. The pure-strategy equilibrium displays low effort and no verification, whereas in the mixed-strategy equilibrium there is a positive probability of performing high effort and of verifying a paper. This is formalized in the following proposition and represented graphically in Figure 1. The proofs to this and all of the following propositions are in the Appendix.

Proposition 1 *The game has a unique Nash equilibrium. i) If* $\Delta e > \frac{(1-p)B_C}{\beta}$*, then the pure-strategy Nash equilibrium is* $(e_L; nv)$*; 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* $q^* = 1 - \frac{\beta \Delta e}{(1-p)B_C}$*, and* C *playing* v *with probability* $r^* = \frac{\Delta e}{(1-p)B_S}$.

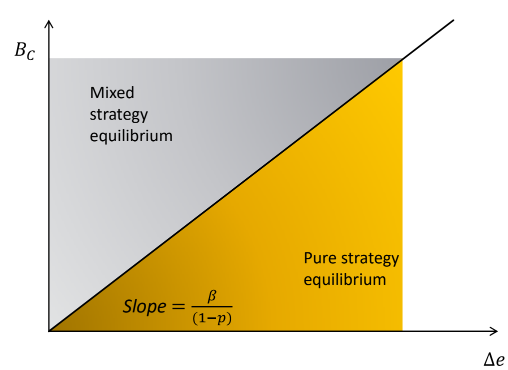

Figure 1: The equlibria of the basic game in the $(\Delta e, B_C)$ space.

[5] The exclusion of uncertainty about verification is obviously a simplification – also replication activities may be scrutinized. Recently, for example, a study questioning the reliability of studies in psychology mentioned in the Introduction (Open Science Collaboration, 2015), received, in turn, criticisms (Gilbert et al., 2016).

Proposition 2 characterizes the likelihoods of two events that will be the subjects of our comparative exercises below: i) a paper being of high quality, and ii) a paper being of low quality and being identified as such. The probability that a paper is of high quality, or the share of high-quality papers, can be taken as a measure of the actual reliability of scientific knowledge (independently of what is observed). As we will point out below, although the probability that a paper is of low quality and is so identified can be seen as a measure of the limits of scientific knowledge or a field of inquiry, it may also represent a signal of "healthy" verification activities in science.

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

$$\Pr(high\ quality) = \begin{cases} 1 - \frac{\beta \triangle e}{B_C} \ if \ \triangle e \leq \frac{(1-p)B_c}{\beta} \\ p \ if \ \triangle e > \frac{(1-p)B_c}{\beta} \end{cases}. \tag{1}$$

The probability that a low-quality paper is identified as such is:

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

Figures 2 and 3 below report $P = \Pr(high\ quality)$ and $P = \Pr(low\ quality\ and\ identified)$ as a function of $\triangle e$, for different values of B_C and $\beta = 1$. The probability that a paper is of high quality is non-increasing in the verification cost $\triangle e$ (Figure 2). If $\triangle e$ is large (relative to C's expected gain from verification), then no verification occurs (pure-strategy equilibrium), and the fraction of high-quality papers is given by 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 that verification activities by C, although being less costly, are also less frequent in this case because the probability to find a low-quality paper is smaller.

The probability that a paper is of low quality and is identified as such is non-monotone and discontinuous as a function of $\triangle e$ (Figure 3); it is positive and increasing in $\triangle e$ when $\triangle e$ is "low" and zero when $\triangle e$ is "high," because we enter the pure-strategy equilibrium region.

Combining Figures 2 and 3, the probability of high-quality papers being produced is 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 (flawed, more limited, or less relevant than initially believed), rather than a signal of the *absence* of these types of papers, indicates the lack of any verification activities activity and, as such, may be cause of concern about the reliability of the whole field.

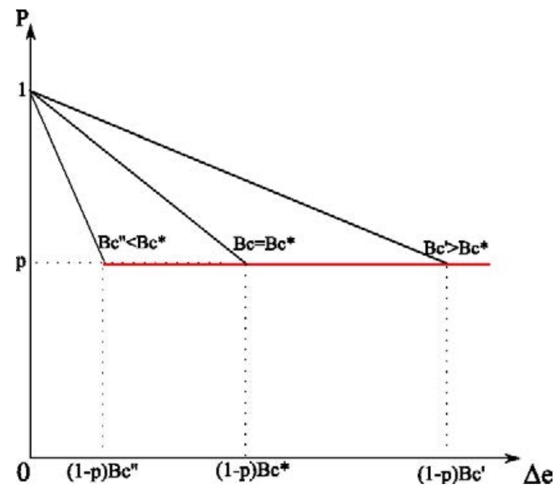

Figure 2: Probability P of a high-quality paper (P) for different values of Δe and B_C (e.g. for three hypotetical levels B''_C, B^*_C and B'_C). P is equal to p for $\Delta e > (1-p)B_C$.

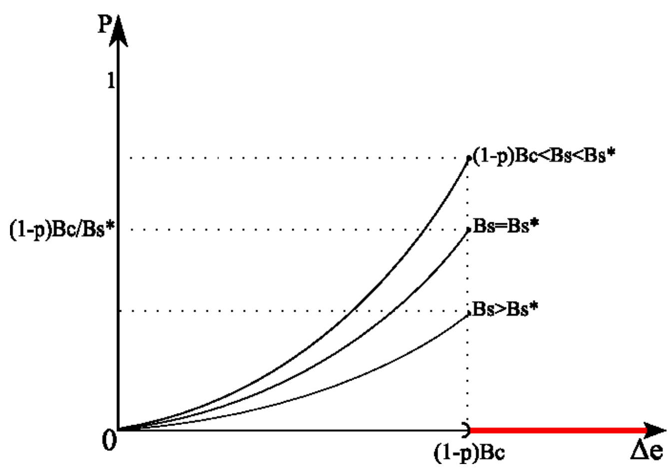

Figure 3: Probability P that a low-quality paper is identified, for different values of Δe and B_S. P drops to zero for $\Delta e > (1-p)B_C$.

4 Assessing and improving the quality of scientific production: implications from the model

The model delivers several results and comparative statics, and is amenable to extensions that allow to analyze the factors that affect the quality of scientific production and to assess costs and benefits of interventions proposed to improve the activities of the scientific community. In this Section we discuss the implications obtained by (1) modifying the incentives to produce novel, radical results; (2) enhancing motivation to engage in incremental research and replication by scientists; (3) accounting in a stylized way for the communitarian activity of science; and (4) adding certification activities to the function of scientific journals.

4.1 Incentives for new findings and the quality of knowledge

Just a few weeks after receiving the Nobel Prize in Medicine, in December 2013 Randy Schekman announced that he would not send his papers to some of the major scientific journals; he claimed that there is too much pressure in the scientific community to produce "novel," "newsworthy" findings, that the policies of the most prestigious journals are further fueling this, and that such a pressure comes at the expense of rigor and depth of inquiry (Schekman, 2013). This view resonates with several other positions on how the "publish or perish" paradigm in science may backfire and lead to worse knowledge production (Abelson, 1990; Giles, 2007). One way to contain incentives for new findings could be, for instance, to weaken the link between funding and publication record. Some authors (Gillies, 2014; Ioannidis, 2011) have suggested egalitarian or random fund allocation rules to overcome the supposedly negative of performance-based research funding systems that give too much importance to quantitative measures such as publication counts (Hicks, 2012).

The following three applications of our model help to assess these views.

Softening incentives to produce novel results A first, natural comparative static concerns the parameter B_S, i.e. the benefits for S to publish a novel result. Note first, from Proposition 2, that the probability that the paper is of high quality does not depend on B_S. 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 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 paper is of low quality and is identified as such. Thus high effort, and consequently, high-quality papers become more attractive because the cost of losing the publication value is larger. However, as a consequence C responds to the increase in B_S by lowering the intensity of her (costly) verification activity, making e_L more attractive up to the point at which S is again indifferent between high and low effort.

Moreover, the probability of identifying low-quality papers decreases with higher values of 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 (Figure 2). This is because the opportunity cost of low effort is also higher when publications are more valuable; as a consequence, C may save on verification activities while leaving S indifferent between high and low effort. A reduction in the publish-or-perish attitude can be interpreted as a reduction of the *relative* value of path-breaking research with respect to more incremental research, causing a simultaneous decrease in B_S and increase in B_C. As we will show below, this would simultaneously increase research quality and the identification of low quality research; thus, although the expected quality of research would be unaffected by softening publication incentives alone, a broader change in the structure of incentives for different types of research may indeed achieve this goal.

A value for low-quality research Let us now extend the basic model and assume that S obtains a positive benefit also from low-quality research. For example, scientists may benefit from the quantity of publications per se. Hamermesh and Pfann (2012) find that quantity has impact on the salaries of academic economists, irrespective of quality. Also, even a study not fully corroborated in subsequent research may nevertheless maintain some validity; the author can still obtain recognition for having opened a new line of research, or having contributed in some other way to the improvement of a scientific theory, for instance identifying weaknesses of an otherwise valid theory (as the examples in Introduction suggested). Finally, the lack of confirmation following the verification from C might depend on factors such as the design of the experiment or the environment where it took place, so that the non-confirmatory result cannot be taken as a definitive proof of the unreliability of the original research. The payoff matrix is in Table 2, where the benefit from high-quality (low-quality) research is B_S^H (B_S^L), with $B_S^L \leq B_S^H$. The proposition that follows presents the solution to this game.

		C: v	C: nv
S	e_H	$B_S^H - e_H; -\beta\Delta e$	$B_S^H - e_H; 0$
	e_L	$pB_S^H + (1-p)B_S^L - e_L; (1-p)B_C - \beta\Delta e$	$B_S^H - e_L; 0$

Table 2: Payoff matrix of the normal form of the game extended to the case of positive rewards to S from a low-quality paper if it is scrutinized. In each cell, the first payoff is for player S, and the second payoff is for player C.

Proposition 3 *i) If* $\Delta e > \frac{(1-p)B_C}{\beta}$*, the Nash equilibrium is* $(e_L; nv)$*; ii) if* $(1-p)(B_S^H - B_S^L) < \Delta e < \frac{(1-p)B_C}{\beta}$*, the Nash equilibrium is* $(e_L; v)$*; iii) if* $\Delta e \leq \min\left[(1-p)[B_S^H - B_S^L]; \frac{(1-p)B_C}{\beta}\right]$, *the mixed-strategy equilibrium has* $r^* = \frac{\Delta e}{(1-p)[B_S^H - B_S^L]}$ *and* $q^* = 1 - \frac{\beta \Delta e}{(1-p)B_C}$.

When B_S^L is sufficiently small (i.e. $(1-p)(B_S^H - B_S^L) > \frac{(1-p)B_C}{\beta}$), then the set of equilibria corresponds to the one of the basic version of the game. If instead B_S^L is closer to B_S^H, the set of equilibria expands by having $(e_L; v)$ as a pure-strategy Nash equilibrium for intermediate values of Δe. In this case, the verification costs are low enough to induce action v by C, but too high to induce high effort by S. By comparing this extension with the basic case of no value for low-quality research we note that, for $(1-p)(B_S^H - B_S^L) < \Delta e < \frac{(1-p)B_C}{\beta}$, the expected quality of research is reduced ($p < 1 - \frac{\beta \Delta e}{B_C}$ because $\Delta e < \frac{(1-p)B_C}{\beta}$), whereas the fraction of low-quality research that is identified as such is higher (because $\frac{\Delta e^2}{(1-p)[B_S^H - B_S^L]B_c} > \frac{(\triangle e)^2}{(1-p)B_S B_c}$ for $\Delta e < \frac{(1-p)B_C}{\beta}$, assuming $B_S^H = B_S$). Within the region of the mixed-strategy equilibrium, the expected quality is unaffected, but low-quality research is identified more frequently. This happens because verification must occur more often to reduce the incentives to exert low effort when low-quality research is positively valued.

To summarize, reducing the reward gap between high and low quality papers has a negative (or nil) impact on the incentives to produce high quality, but it has positive (or nil) impact on the incentives to scrutinize others' work.

A value for confirmed results A further way to alter the incentives to produce novel and reliable research is to provide additional rewards if a study is of high quality and is confirmed, or replicated as such. For example, replications that confirm the original findings, or incremental research based on a given study may also positively affect the originator of that study. Starting from the basic game, we modify the payoffs by allowing S to obtain a higher reward when his research is of high quality and is verified. Let B_S^v denote the benefit for S if research is of high quality and is verified, and B_S^{nv} the benefit of unverified research, with $B_S^v \geq B_S^{nv}$.[6] Table 3 presents the game in normal form and Proposition 4 provides the solution.

		C: v	C: nv
S	e_H	$B_S^v - e_H; 0 - \beta\Delta e$	$B_S^{nv} - e_H; 0$
	e_L	$pB_S^v - e_L; (1-p)B_C - \beta\Delta e$	$B_S^{nv} - e_L; 0$

Table 3: Payoff matrix of the normal form of the game extended to the case of higher rewards to S from a high-quality paper if it is scrutinized than if it is not scrutinized. In each cell, the first payoff is for player S, and the second payoff is for player C.

[6] Like in the basic model, we assume $\frac{B_C}{\beta} \leq B_S^v$.

Proposition 4 *The game has a unique Nash equilibrium. i) If* $\Delta e > \frac{(1-p)B_C}{\beta}$*, then the pure-strategy Nash equilibrium is* $(e_L; nv)$*; 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* $q^* = 1 - \frac{\beta \Delta e}{(1-p)B_C}$*, and C playing v with probability* $r^* = \frac{\Delta e}{(1-p)B_S^v}$.

Propositions 1 and 4 coincide for $B_S^v = B_S$. By distinguishing between B_S^v and B_S^{nv}, Proposition 4 shows that the probability that C scrutinizes depends on the benefit for confirmed high quality, rather than on high quality per se. If this value increases, the return from e_H increases, and then r must be reduced to leave S indifferent between e_H and e_L. Notice that increasing the value for confirmed results, which may come from the publication of systematic replication exercises such as, in economics, the early contribution by Dewald et al. (1986), would not have an impact on quality, but would increase the fraction of low quality papers that are identified.

4.2 Motivating incremental research and replication

The concerns that the excessive incentives toward producing compromise depth and reliability are often expressed together with a recommendation to provide more recognition for incremental and replicative work (Ioannidis, 2014; Schekman, 2013). Several initiatives have been undertaken lately to enhance further reviews and research on existing findings. An increasing number of journals (as well as public funding agencies) have data-sharing policies. The platform PubPeer (https://pubpeer.com/), allows scientists to review, comment and potentially propose corrections to published papers. Journals of the Public Library of Science, especially *PLOS ONE* (http://www.plosone.org/), and in sociology, such as *Sociological Science*, encourage post-publication comments whereas other journals, such as the *New England Journal of Medicine*, include a "Journal Watch" section on their website to stimulate the collection and discussion of interesting published findings (http://www.jwatch.org/). Finally, the "Experimental Economics Replication Project" (http://experimentaleconreplications.com/) is an initiative to promote replication of published lab experimental studies in economics, associated to a "prediction market" where people can bet on what studies will be replicated (see also Hanson, 1990 for a proposal of a formal betting market in science).

Our model provides insights about these views and initiatives. In particular, we show that acting on the incentives to produce novel findings and affecting the incentives to review existing work do not necessarily lead to the same or symmetric effects.

Increasing the benefits from detecting low-quality research In the previous section we showed that a decrease in the payoffs from publishing fully novel results does not lead, per se, to enhancing the expected quality of research. This is a result of the interplay of incentives in mixed strategy equilibria. We now show that there is no symmetry of results when, instead, the payoff from detecting a low-quality study increases. Note from Proposition 2 that in the mixed-strategy equilibrium region, larger benefits B_C from identifying a low-quality paper (or lower costs via a reduction in β) do increase the fraction of high-quality papers through a strategic effect on S; because verification is more rewarding (or relatively less costly), S increases his effort in order to reduce C's incentives to verify.

The effect of increasing incentives for verification can therefore be different from just softening incentives for entirely novel findings, as shown above. In a broader sense, a reduction in the publish-or-perish attitude can be interpreted as a reduction of the relative value of path-breaking research with respect to more incremental research, or a simultaneous decrease in B_S and increase in B_C. This, as demonstrated above, would increase both research quality and the identification of low quality research. Not surprisingly, a reduction in verification costs (reduction of Δe and β) has a similar effect to an increase in B_C. In that respect, we observe that a reduction in publish-or-perish attitude would also reduce the opportunity cost of scrutinizing others' work, which includes effort and time that the scientist could devote to her own research. As for a reduction in β, which may come from policies favouring data sharing, it turns out to have a direct positive effect on scientific quality. However, in this section we will show how this conclusion may not hold when we take into account the indirect effects on the incentives of scientists to start new projects.

A value for confirmatory results An direct way to reward incremental and replicative studies is to reward them even if they do not include any new findings (including results that limits the validity for the original study), and just confirm the work by a scientists to be reliable. For instance, editors could explicitly solicit replication attempts, promising publication in exchange (Hamermesh, 2007; Wagenmakers and Forstman, 2014). Suppose therefore that C obtains a positive payoff even when verifying a paper that is of high quality, as in the game represented in Table 4.

		C: v	C: nv
S	e_H	$B_S - e_H;\ B_C^H - \beta\Delta e$	$B_S - e_H;\ 0$
	e_L	$pB_S - e_L;\ pB_C^H + (1-p)B_C^L - \beta\Delta e$	$B_S - e_L;\ 0$

Table 4: Payoff matrix of the normal form of the game extended to the case of positive rewards to C from verifying a high-quality paper. In each cell, the first payoff is for player S, and the second payoff is for player C.

B_C^H (B_C^L) corresponds to the value for C of discovering a high (low)-quality publication. We assume that $B_C^H \leq B_C^L$; although positive, the benefit from confirming a high-quality result is no greater than spotting a lower-quality study. The payoff structure above implies that, in this case, it is possible to sustain an equilibrium where no low-quality papers are produced. Proposition 5 formalizes the solution.

Proposition 5 *i) If* $\Delta e > \frac{pB_C^H+(1-p)B_C^L}{\beta}$*, then the Nash equilibrium is* (e_L, nv)*; ii) if* $(1-p)B_S < \Delta e < \frac{pB_C^H+(1-p)B_C^L}{\beta}$*, the Nash equilibrium is* (e_L, v) *; iii) if* $\Delta e < \min\left\{(1-p)B_S; \frac{B_C^H}{\beta}\right\}$*, the equilibrium is* $(e_H,\ v)$ *; iv) finally, if* $\frac{B_C^H}{\beta} \leq \Delta e \leq \min\left\{\frac{pB_C^H+(1-p)B_C^L}{\beta}; (1-p)B_S\right\}$*, 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 Δe is small enough, C verifies even if she expects the research by S to be of high quality with probability 1. Thus, if confirmatory results are positively valued by the scientific community, it is possible that low-quality papers are not produced.

Note however that for intermediate values of Δe (i.e. $(1-p)B_S < \Delta e < \frac{pB_C^H+(1-p)B_C^L}{\beta}$), 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 in the model. In other words, verification is a necessary, but not a sufficient condition for eliciting high effort. If verification is incentivized by increasing the benefit for C per se, then a sufficiently high benefit for S will be required to lead to high quality.

Incentives for incremental research and the decision to start new projects A further way to not only consider the direct effects of enhancing replication incentives, but also the indirect, strategic effects, is to consider how replication and incremental research might affect the decision itself of a scientists to start a novel project, something that we took for granted to happen thus far. A scientist can always decide to not start a project at all; in our model, if this outside option has a payoff of zero, then S will exert effort at all only if he expects a non-negative payoff as an equilibrium of the game. An extended version of the basic game introduced in Section 2.1, which includes the entry decision is represented in extensive form in Figure 4 below, and is solved by backward induction.

The expected payoff of S depends on which type of equilibrium prevails in the subgame following the entry decision. If $\Delta e \geq \frac{(1-p)B_C}{\beta}$, the pure-strategy equilibrium is $(e_L; nv)$, which corresponds to a payoff of $B_S - e_L$. Therefore S will start the project as long as $B_S \geq e_L$. If $\Delta e < \frac{(1-p)B_C}{\beta}$, a mixed-strategy equilibrium prevails, and S obtains an expected payoff equal to $B_S - e_H$ (this is the payoff that S obtains by playing e_H, and in a mixed-strategy equilibrium S must be indifferent between any of his possible strategies). Therefore, S will

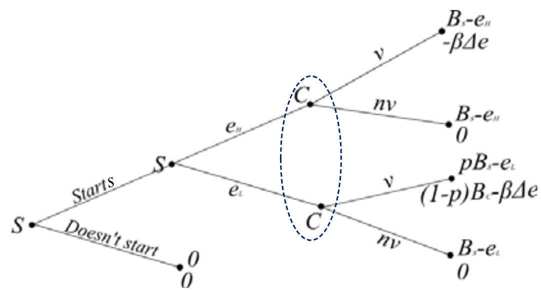

Figure 4: The basic game with the entry decision: extensive form representation. The top payoff at each end node refers to S, and the bottom payoff to C. The dotted circle indicates nodes that are part of the same information set.

start the project if $B_S \geq e_H$. We re-write $B_S - e_L$ as $\Delta e + B_S - e_H$ and $B_S - e_H$ as $-\Delta e + B_S - e_L$ and represent graphically the decision to start a project in Figure 5.

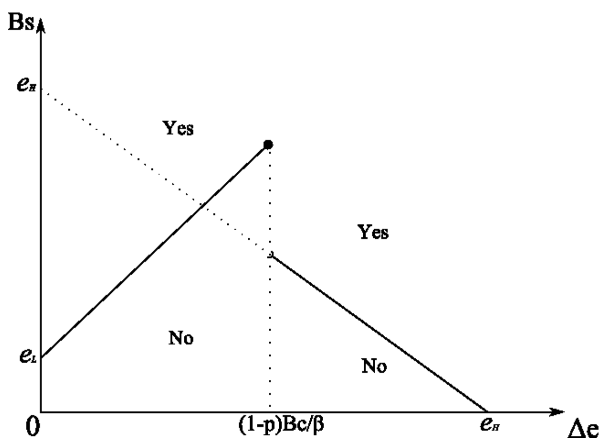

Figure 5: The decision to start a project. S starts a project in the regions indicated with "Yes", and does not start a project in regions denoted with "No".

Although an increase in the benefit from publication B_S always raises the likelihood that a project is started, a reduction in Δe, i.e. of the extra cost of producing high-quality research, does not have an unambiguous effect. A reduction in the cost differential between low and high effort increases the likelihood of starting a project only in the mixed-strategy region (low values of Δe); in the pure-strategy parameter space a reduction in Δe can lower the likelihood that the project is started if B_S is low.

Consider also the effect of changes in the parameter β, which affects C's verification costs. Suppose that β is lowered, say because of policies favoring the sharing of data and methods within a field, or is lower in certain fields. Then, the vertical line corresponding to $\frac{(1-p)B_C}{\beta}$ moves right, enlarging the parameter space associated to a mixed-strategy equilibrium. For intermediate values of B_S, moving from a pure to a mixed-strategy equilibrium may lead S to prefer to not start the project, whenever the positive probability of verification causes a negative expected payoff for S. These foregone projects may be socially valuable, because the positive externalities from research may misalign social incentives and the private incentives of S and C. If that is the case, a policy that would have unambiguous positive effects in the basic game (a reduction in Δe) could instead backfire if the decision of S to start a project is considered.[7]

To conclude, although policies to enhance incremental or replication research may enhance the quality of research overall and the ability to sort quality, making incremental or replication research less costly or more diffused may also have some opposite effects; it may depress motivations to undertake certain research in the first place, and also might crowd out incentives to perform verification activities at all. These countervailing (and somewhat counterintuitive) effects of verification activities need to be considered in devising optimal policies. For example, some degree of control or protection over one's data (maybe temporary) might then help keeping in balance the trade-off between producing novel findings and incremental or replicative research. We also showed that a higher number of potential verifiers may crowd out incentives for scrutiny. Recent debates among psychology scholars concern whether too much attention to dissecting existing studies might come at the expense of more innovative exploration (Bartlett, 2014); this appears to highlight the tradeoffs derived in the analysis above.

4.3 Science as a community and scientific quality

The analysis so far considered a simple, one-shot interaction between two scientists, one playing the role of knowledge originator and the other one the role of potential scrutinizer. This excludes some important communitarian aspects of science.

First, multiple scientists can engage in verification activities simultaneously. In this section we consider the simplest extension in this direction, i.e. the presence of two researchers who choose simultaneously whether to verify or not. *Prima facie*, a larger community of scientists active in research that builds upon and verifies new findings may be considered an alternative means to guarantee higher reliability of science.[8] Our results below define boundary conditions to this view.

Second, in a scientific community scientists are active both as generators of new knowl-

[7] Mueller-Langer and Andreoli Versbach (2014) propose a model where mandatory (and immediate) data disclosure policies might inhibit researchers to undertake research in the first place.

[8] Especially in smaller, peripheral countries, the size of the local scientific community may be considered a policy variable, as long as governments can intervene to increase the degree of international openness of national journals (Sambunjak et al., 2009) or incentivize the participation of the country's scholars to the global scientific community.

edge and replicators/verifiers of existing findings. We address this with a somewhat major modification of the basic game to allow for repeated interactions between two scientists. Each player is in the position of the originator of new research as well as of the scrutinizer, and this happens multiple times.

A larger community of verifiers Let us denote with C_1 and C_2 the two researchers who choose simultaneously whether to verify S's research or not. We assume that when both C_1 and C_2 verify, they equally share the benefit of discovering a low-quality paper, and that low-quality is ascertained if at least one colleague verifies. We can interpret this assumption also in terms of a "winner takes all" race where only the first to discover obtains recognition, and that, ex ante, all players have the same likelihood of being first.

In the case of mixed-strategy equilibria, we focus our attention on symmetric equilibria with respect to the behavior of C_1 and C_2.[9] The solution is summarized in the following Proposition.

Proposition 6 *The equilibria of the game are as follows: i) if* $\Delta e > \frac{(1-p)B_C}{\beta}$*, then the unique, pure-strategy Nash equilibrium is* $(e_L; nv; nv)$*; ii) if* $\Delta e \leq \frac{(1-p)B_C\left[\frac{1}{2}+\frac{1}{2}\left(1-\sqrt{1-\frac{\beta\Delta e}{(1-p)B_S}}\right)\right]}{\beta}$*, then there exists a symmetric Nash equilibrium in mixed strategies, with* S *playing* e_H *with probability* $q^* = 1 - \frac{\beta\Delta e}{(1-p)B_C\left[\frac{1}{2}+\frac{1}{2}\left(1-\sqrt{1-\frac{\beta\Delta e}{(1-p)B_S}}\right)\right]}$*, and* C_1 *and* C_2 *playing* v *with probability* $r^* = 1 - \sqrt{1 - \frac{\Delta e}{(1-p)B_S}}$.

The proposition conveys a number of insights. First, the parameter space in which a pure-strategy equilibrium prevails is not affected by the number of potential scrutinizers, because this area is determined by the condition that no researcher verifies. In other words, the existence of multiple potential verifying colleagues is not, per se, a sufficient condition to expect some verification activities to occur.

Second, in the mixed-strategy equilibrium region the comparison of q^* and r^* with the case of a single C shows that the probabilities that high effort is exerted and that verification activities occur are lower (see the Appendix for the proof). Therefore, a larger set of potential scrutinizers reduces both the expected quality of research and the probability that low quality is detected. The result hinges upon the lower reward from scrutiny that each colleague obtains due to sharing the credit in the case simultaneous scrutiny; this, in turn, lowers S's incentives to provide high effort.

[9] For $\frac{(1-p)B_C\left[\frac{1}{2}+\frac{1}{2}\left(1-\sqrt{1-\frac{\beta\Delta e}{(1-p)B_S}}\right)\right]}{\beta} < \Delta e \leq \frac{(1-p)B_C}{\beta}$, there is no symmetric mixed-strategy equilibrium. However, for $\Delta e \leq \frac{(1-p)B_C}{\beta}$, there always exist asymmetric equilibria whereby one colleague plays the pure strategy nv, whereas the other colleague and S play mixed strategies with probabilities as in Proposition 1.

Proposition 6 thus shows crowding out of incentives for scrutiny in larger communities. However, we must recognize that the peer recognition for having detected a low quality paper may be *higher* in larger (and possibly, more visible) communities. In that case, B_C would be higher in the case of two colleagues, counteracting the negative effect on q^*.

Repeated interactions Consider two researchers, 1 and 2, interacting repeatedly for $t = 1, 2, .., \infty$. The two researchers take in turn the role of S and C. In particular, 1(2) plays the role of S in odd (even) periods. Let δ be the common discount factor. In order to avoid the problems associated to the presence of mixed-strategy equilibria in repeated games, let us assume that 1 and 2 play in each period the game that allows for equilibria in pure strategies where verification occurs (see page 15, Proposition 5 in particular). Specifically, if $(1-p)B_S < \Delta e < \frac{pB_C^H+(1-p)B_C^L}{\beta}$, then the unique Nash equilibrium is (e_L, v), whereas if $\Delta e < \min\left\{(1-p)B_S; \frac{B_C^H}{\beta}\right\}$, the equilibrium is $(e_H,\, v)$.

We are interested, in particular, in investigating the possibility of "collusion" between 1 and 2, i.e. the sustainability of an equilibrium where each player, when acting as C, refrains from verifying, expecting the other researcher to do the same in the future. In other words, we ask whether (e_L, nv) can be the action pair played in every period of the repeated game, when it would not be an equilibrium in a one-shot or finitely repeated game. If this is the case, the repetition of the game may thus reduce the expected quality of research or the fraction of low-quality research that is discovered.

We assume that the researchers play trigger strategies in which (e_L, nv) is played at $t = 1$ and in any subsequent period as long as players acting as C have always played nv, turning to Nash equilibrium otherwise. The following Proposition holds.

Proposition 7 *(i) Suppose that* $\Delta e < \min\left\{(1-p)B_S;\ \frac{B_C^H}{\beta}\right\}$. *Then,* $(e_L,\ nv)$ *is sustainable as outcome for each* t *if* $\delta \geq \frac{-\Delta e+\sqrt{(\Delta e)^2+4(1-p)(B_C^L-B_C^H)\left[pB_C^H+(1-p)B_C^L-\beta\Delta e\right]}}{2(1-p)(B_C^L-B_C^H)}$. *(ii) Suppose that* $(1-p)B_S < \Delta e < \frac{pB_C^H+(1-p)B_C^L}{\beta}$. *Then* (e_L, nv) *is sustainable as outcome for each* t *if* $\delta \geq \frac{pB_C^H+(1-p)B_C^L-\beta\Delta e}{(1-p)B_S}$.

Proposition 7 implies that independently from the pure-strategy equilibrium in the stage game, players who are sufficiently patient may prefer to save on costly verification in the current period, expecting to receive the same treatment in the future when acting as knowledge originator. As usual, δ can be interpreted as related to the frequency of interactions between 1 and 2; in our context, it may represent how likely it is for each researcher to meet the other again, with exchanged roles. Because this probability is higher in smaller or more specialized communities, collusion, and then the suppression of verification activities, is more likely in this case.

An implication of these results is that collusion may be more likely in smaller scientific communities, where researchers involved both in the production of new knowledge and in incremental, replicative work are more likely to interact with each other multiple times. Proposition 6 above, however, identifies, a drawback in the ability of *larger* communities to provide incentives for incremental and replicative work, because of "free-rider" or competition effects.

4.4 Journals as quality certifiers

Scientific journals enjoy different levels of recognition; to the extent that a journal's status is related to the reliability of the studies that it publishes, top-rated journals, in particular, may at least partially serve as a certification for the quality of research. The last set of policies that we propose considers this case.

Assume that there are two journals: journal A, which publishes high-quality papers and gives a scientist a reward B_S^A, and journal B, where both high and low-quality works are published and from which a scientist derives a reward $B_S^B \leq B_S^A$. For simplicity, suppose that the evaluation from the journals is instantaneous (there is no cost of time). Thus, regardless of the level of effort chosen by S, his optimal submission strategy is always i) submitting first to journal A, and ii) submitting to journal B if rejected by A. A further implication is that C would never verify a paper published on journal A. Therefore, her relevant choice is between v and nv only after observing a paper published in journal B. Furthermore, if journal A publishes all the high-quality papers (as assumption that we will relax below), then C, upon observing a paper in journal B, can correctly predict that the paper is of low quality with probability 1. In other words, the two journals deliver perfect discrimination between high and low quality papers.

Denote with B_C^B the credit that C obtains from verifying a low-quality paper.[10] Based on the previous considerations, the strategic interaction can be represented as in Table 5 below, whereas the solution is laid out in Proposition 8.

		C: v	C: nv
S	e_H	$B_S^A - e_H;\ 0 - \beta\Delta e$	$B_S^A - e_H;\ 0$
	e_L	$pB_S^A + (1-p)B_S^B - e_L;\ B_C^B - \beta\Delta e$	$pB_S^A + (1-p)B_S^B - e_L;\ 0$

Table 5: Payoff matrix of the normal form of the game extended to the case of two journals. In each cell, the first payoff is for player S, and the second payoff is for player C.

[10] One could argue that the low quality of papers published in journal B is common knowledge in the scientific community; however, action v is needed to identify the specific weaknesses and limitations of the paper.

Proposition 8 *The game has a unique Nash equilibrium. i) If* $\Delta e \geq \max\left\{(1-p)(B_S^A - B_S^B); \frac{B_C^B}{\beta}\right\}$, *the pure-strategy Nash equilibrium is* $(e_L; nv)$*; ii) if* $(1-p)(B_S^A - B_S^B) \leq \Delta e \leq \frac{B_C^B}{\beta}$, *then the pure-strategy Nash equilibrium is* $(e_L; v)$*; iii) if* $\Delta e \leq (1-p)(B_S^A - B_S^B)$, *the pure-strategy Nash equilibrium is* $(e_H; nv)$.

What Proposition 8 shows is that if journals can perfectly sort out the quality of papers, then the equilibrium in which a scientist exerts high effort (thus quality is high of probability 1) and no verification activities occur is sustainable, with the parameter region in which this occurs being larger the larger is $(B_S^A - B_S^B)$. Note that this does not occur when low-quality research provides positive credit but there are no multiple journals, as in the model in Section 4.1. Another implication is that if $(1-p)(B_S^A - B_S^B) \geq \frac{B_C^B}{\beta}$ (a condition more likely to hold when the credit for discovery a limitation in journal B is low), then verification activities are fully crowded out by the certification activity of the journal. If we compare this extension with the basic version of the model, such crowding out may be even detrimental to overall quality if it extends the region where low effort and no verification activity are chosen.

Although high quality is a necessary condition to be published in a high-status journal, it may not be sufficient, for example because of a journal's capacity constraints. In a further extension (whose details are in the Appendix), we introduce a probability x that a high-quality paper is accepted in journal A. As expected, we find that the smaller is x, the smaller is the region where $(e_H; nv)$ is sustainable as an equilibrium, because the incentives towards high effort are lower in this case.

Another interesting case is when the prestige of a journal is unrelated to the reliability of a paper. This is akin to the analysis in Ellison (2002) where quality has two dimensions, one affecting the probability to publish on a high-status journal, and the other associated to the reliability of the underlying research. Let us assume that in this case journals differ only for the credit that they assign to scientists ($B_S^A \geq B_S^B$) and to colleagues ($B_C^A \geq B_C^B$). If we consider the basic set-up of the model (as in Table 1), we can derive the following two the implications: first, the pure-strategy region in which no verification activities and low effort occur is larger for journal B, and second, in the mixed-strategy regions, journal A is characterized by higher quality and less verification activities.

In other words, reputational mechanisms are such that they induce prestigious journals to produce also more reliable research. A sort of division of labor may emerge as optimal, where journal prestige is attached to dimensions of quality such as relevance, and scrutiny for reliability is delegated to the scientific community at large.

5 Discussion and applications

> "Jack, if you think you have a good idea, publish it! Don't be afraid to make a mistake. Mistakes do no harm because there are lots of smart people out there who will immediately spot a mistake and correct it. [...] If it happens to be a good idea, however, and you don't publish it, science may suffer a loss."
>
> Conversation between Linus Pauling and Jack Dunitz, as reported by Mario Livio in *Brilliant Blunders* (2014, p.143).

Our model conveys insights about the operating of the overall scientific endeavor, and clarifies 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, on the other hand, the incentives provide further work upon and, in the process, possibly question existing and established results. We also showed an asymmetry of effects between lowering incentives to produce new research, and increasing incentives to do additional work and verification on existing findings. Finally, we pointed out some countervailing effects of encouraging verification activities.

We conclude with an outline of some implications of our analysis for companies and investors, opportunities for empirical tests, applications beyond the scientific community, and limitations of our analysis.

Managerial insights The main message from our analysis for companies and investors exploring the scientific landscape in search for opportunities is that the absence of low-quality findings in a scientific area may not necessarily signal the promise of a given line of research; this can be the effect, instead, of the a lack of verification activities (additional incremental research, replications, etc.) and, as such, can be a source of concern about the reliability of the overall research. Therefore, scientific fields that display controversies and where flaws and limitations do emerge may be more solid and promising than fields where no such features are observed. Similarly, a field in which a lot of incremental research occurs may not necessarily be a mature or declining area of inquiry; incremental research could represent a source, again, of greater reliability. In contrast, the "popularity" of an area of research as represented by the size of a scientific community might not necessarily imply higher-quality science, if the incentives to perform verification and incremental work are diluted.

Ideas for empirical tests The managerial considerations just outlined also offer insights for testing our theory. Evidence of a positive correlation between a certain level of debates and critiques between scholars in a given scientific area, and the future success of that particular area of research would provide empirical support to the findings of the model. For example, debates and criticisms can be detected through the analysis of "negative citations" (Catalini et al., 2015), and their impact on the future development of a field can be measured through forward citations, breadth of applications (e.g. citations by studies in other fields), and how long papers published during "controversial" periods continue to be cited. Moreover, "shocks" such as the exogenous influx of scientists in a given discipline (Borjas and Doran, 2012; Moser et al., 2014) could be exploited to test how community size affects the scientific debate and the overall quality of the produced knowledge. Finally, we should expect findings to be less reliable in scientific communities that are more isolated and whose rewards comes from publications in "local" journals.

Beyond science This framework can also be applied to other environments characterized, like the scientific community, by the possibility of producing both new content and contributions to existing findings on the one hand, and 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 findings, and how different institutional arrangements affect these motivations, is of relevance in these settings too.

Limitations and avenues for further work Our model excluded from the analysis several factors that may have an impact on the functioning of scientific communities and their ability to produce reliable research.

First, we did not consider the choice of scientists between different types of research. In our context, this would imply the choice (in terms of effort and time allocation) between radical (or innovative) research and more incremental (replicative) research. Although it is commonly argued that the current structure of incentives in science is biased in favor of the former, a full-fledged analysis would help to substantiate this claim in terms of the social desirability of this outcome, and to assess the effects of policy interventions aimed at modifying the current incentive structure.

Second, as discussed in Section 2.2, we represent science as a process of search for truth, in which symmetry is imposed between positive results confirming theories and negative ones rejecting them. However, research has documented a publication bias against negative results several fields (Easterbrook et al., 1991; Ferguson, 2007; Stanley, 2005), and scientists, as most human beings, may suffer from confirmation bias (Nuzzo, 2015). A richer model would account for these aspects.

Finally, the model could be enhanced in its institutional realism by including a role for other actors such as reviewers and journals editors. Although Section 4.4 is a first step in this direction, modeling the objectives and behaviors of referees and editors may significantly enrich the analysis of the determinants of scientific quality.

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] Allison, D. B., Brown, A. W., George, B. J., & Kaiser, K. A. (2016). Reproducibility: A tragedy of errors. *Nature*, 530, 27-29.

[3] Aschwanden,C. (2015). Science Isn't Broken. It's just a hell of a lot harder than we give it credit for. *FiveThiryEight.* http ://fivethirtyeight.com/features/science-isnt-broken/

[4] 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.

[5] Azoulay, P., Furman, J. L., Krieger, J. L., & Murray, F. E. (2015a). Retractions. *Review of Economics and Statistics*, forthcoming.

[6] Azoulay, P., Bonatti, A. & Krieger, J. L. (2015b): The career effects of scandal: Evidence from scientific retractions. Working paper.

[7] Banal-Estañol, A., & Macho-Stadler, I. (2010). Scientific and commercial incentives in R&D: Research versus development?. *Journal of Economics and Management Strategy*, 19(1), 185-221.

[8] Baum, J. A., & Silverman, B. S. (2004). Picking winners or building them? Alliance, intellectual, and human capital as selection criteria in venture financing and performance of biotechnology startups. *Journal of Business Venturing*, 19(3), 411-436.

[9] Bartlett, T. (2014). Replication crisis in psychology research turns ugly and odd. *The Chronicle of Higher Education*, June 23.

[10] Begley, C. G., & Ellis, L. M. (2012). Drug development: Raise standards for preclinical cancer research. *Nature*, 483(7391), 531-533.

[11] Borjas, G.J. & Doran, K.B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *The Quarterly Journal of Economics*, 127(3), 1143-1203.

[12] Broad, W., & Wade, N. (1982). *Betrayers of the Truth. Fraud and Deceit in the Hall of Science.* Simon and Schuster, New York.

[13] Camerer *et al.* (2016), Evaluating replicability of laboratory experiments in economics. *Science*, forthcoming.

[14] Catalini, C., Lacetera, N., & Oettl, A. (2015). The incidence and role of negative citations in science. Proceedings of the National Academy of Sciences, 112(45), 13823-13826.

[15] Collins, H. (1992). *Changing Order: Replication and Induction in Scientific Practice*. Chicago: The University of Chicago Press.

[16] Dasgupta, P., & David, P. A. (1994). Toward a new economics of science. *Research Policy*, 23(5), 487-521.

[17] Dewald, W. G., Thursby,J.G., Anderson, R. G. (1986). Replication in empirical economics: the Journal of Money, Credit and Banking project. *American Economic Review*, 76(4), 587–603.

[18] Ellison, G. (2002). Evolving standards for academic publishing: a q-r theory. *Journal of Political Economy*, 110 (5),994-1034.

[19] Erard, B., & Feinstein, J. S. (1994). Honesty and evasion in the tax compliance game. *The RAND Journal of Economics*, 25(1), 1-19.

[20] Easterbrook, P. J., Gopalan, R., Berlin, J. A., & Matthews, D. R. (1991). Publication bias in clinical research. *The Lancet*, 337(8746), 867-872.

[21] Ferguson, C. J. (2007). Evidence for publication bias in video game violence effects literature: A meta-analytic review. *Aggression and Violent behavior*, 12(4), 470-482.

[22] Freedman, L. P., Cockburn, I. M., & Simcoe, T. S. (2015). The economics of reproducibility in preclinical research. *PLoS Biol*, 13(6), e1002165.

[23] Gilbert, D.T, King, G., Pettigrew, S. Timothy D., & Wilson, T.D. (2016). Comment on "Estimating the reproducibility of psychological science". *Science*, 351 (6277), 1037.

[24] Giles, J. (2007). Breeding cheats. *Nature*, 445(18): 242–243.

[25] Gillies, D. (2014). Selecting applications for funding: why random choice is better than peer review. *RT. A Journal on Research Policy and Evaluation*, 2(1).

[26] Graetz, M. J., Reinganum, J. F., & Wilde, L. L. (1986). The tax compliance game: Toward an interactive theory of law enforcement. *Journal of Law, Economics and Organization*, 2(1), 1-32.

[27] Hamermesh, D. S. (2007). Replication in economics. *Canadian Journal of Economics*, 40(3), 715-733.

[28] Hamermesh, D. S., & Pfann, G. A. (2012). Reputation and earnings: the roles of quality and quantity in academe. *Economic Inquiry*, 50(1), 1-16.

[29] Hanson, R. (1990). Could gambling save science? *Proceedings of the Eighth International Conference on Risk and Gambling*, London, July.

[30] Hartemink, A. E. (2000). Publish or Perish (3)–Fraud and ethics. *Bulletin of the International Union of Soil Sciences*, 97, 36-45.

[31] Häussler, C., Jiang, L., Thursby, M. & Thursby, J. (2014). Specific and general information sharing among competing academic researchers. *Research Policy*, 43(3), 465-75.

[32] Henry, E. (2009). Strategic disclosure of research results: the cost of proving your honesty. *The Economic Journal*, 119, 1036-1064.

[33] Henry, E., & Ottaviani, M. (2013). Research and the approval process. Working paper.

[34] Hicks, D. (2012). Performance-based university research funding systems. *Research Policy*, 41(2), 251-261.

[35] Holmberg, S. D. (2008). *Scientific Errors and Controversies in the HIV/AIDS Epidemic*. Greenwood Publishing Group.

[36] Howson, C., & Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach*. Open Court Publishing Co.

[37] 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.

[38] Ioannidis, J. P., Ntzani, E. E., Trikalinos, T. A., & Contopoulos-Ioannidis, D. G. (2001). Replication validity of genetic association studies. *Nature Genetics*, 29(3), 306-309.

[39] Ioannidis, J. P.A. (2011). More time for research: fund people not projects. *Nature*, 477(7366), 529-531.

[40] Ioannidis, J. P. A. (2014). How to make more published research true. *PLOS Medicine*, 11(10), e1001747.

[41] Jensen, R. & Thursby, M. (2001). Proofs and prototypes for sale: The licensing of university inventions. *American Economic Review*, 91(1), 240-259.

[42] Lacetera, N., & Zirulia, L. (2011). The economics of scientific misconduct. *Journal of Law, Economics, and Organization*, 27(3), 568-603.

[43] Lacetera, N., & Zirulia, L. (2012). Individual preferences, organization, and competition in a model of R&D incentive provision. *Journal of Economic Behavior & Organization*, 84(2), 550-570.

[44] Lakhani, K. R., & Von Hippel, E. (2003). How open source software works: "Free" user-to-user assistance. *Research Policy*, 32(6), 923-943.

[45] Lents, N.H. (2016). Paleoanthropology Wars: The discovery of Homo naledi has generated considerable controversy in this scientific discipline. *Skeptic*. http://www.skeptic.com/reading_room/paleoanthropology-wars-the-discovery-of-homo-naledi/

[46] Livio, M. (2014). *Brilliant Blunders. From Darwin to Einstein-Colossal Mistakes by Great Scientists that Changed our Understanding of Life and the Universe*. Simon and Schuster.

[47] Lu, S. F., Jin, G. Z., Uzzi, B., & Jones, B. (2013). The retraction penalty: Evidence from the Web of Science. Scientific Reports, 3, 3146.

[48] Macho-Stadler, I., & Pérez-Castrillo, D. (2010). Incentives in university technology transfers. *International Journal of Industrial Organization*, 28(4), 362-367.

[49] Masten, S. E. (2006). Authority and commitment: Why universities, like legislatures, are not organized as firms. *Journal of Economics and Management Strategy*, 15(3), 649-684.

[50] Merck (2015). Experts in your field, close to your location: http://www.merck.com/licensing/contact-us/home.html (last accessed 04/10/2015).

[51] Mialon, H. (2010). What's in a title? Predictors of the productivity of publications. Working paper.

[52] Mirowski, P., & Sklivas, S. (1991). Why econometricians don't replicate (although they do reproduce). *Review of Political Economy*, 3(2), 146-163.

[53] Moser, P., Voena, A. & Waldinger, F. (2014). German Jewish émigrés and US invention. *American Economic Review*, 104(10), 3222-55.

[54] Mueller-Langer, F. & Andreoli-Versbach, P. (2014) Open access to research data: Strategic delay and the ambiguous welfare effects of mandatory data disclosure. Munich discussion paper 2014-32.

[55] Muma, J. R. (1993). The need for replication. *Journal of Speech, Language and Hearing Research*, 36(5), 927.

[56] Nuzzo, R. (2015). Fooling ourselves. *Nature*, 526, 182-185.

[57] Pfizer (2015). Innovative partnering models:

http://www.pfizer.com/partnering/areas_of_interest/innovative_partnering_models (last accessed 04/10/2015).

[58] Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349.

[59] Popper, K. R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.

[60] Reinganum, J. F., & Wilde, L. L. (1986). Equilibrium verification and reporting policies in a model of tax compliance. *International Economic Review*, 27(3), 739-760.

[61] Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5 pt. 2), S71-S102.

[62] Rousseau, D. L., & Porto, S. P. S. (1970). Polywater: Polymer or artifact?. *Science*, 167(3926), 1715-1719.

[63] Ryan, P.K. (2013). *How Venture Capital Works*. Rosen Publishing.

[64] Sambunjak, D., Huić, M., Hren, D., Katić, M., Marušić, A., & Marušić, M. (2009). National vs. international journals: views of medical professionals in Croatia. *Learned Publishing*, 22(1), 57-70.

[65] Sauermann, H. & Stephan, P. (2013). Conflicting logics? A multidimensional view of industrial and academic science. *Organization Science*, 24(3), 889-909.

[66] Schekman, R. (2013). How journals like Nature, Cell and Science are damaging science. *The Guardian*.

http://www.theguardian.com/commentisfree/2013/dec/09/how-journals-nature-science-cell-damage-science

[67] Sherwood, S. (2011). Science controversies past and present. *Physics Today*, 64(10), 39-44.

[68] Stanley, T. D. (2005). Beyond publication bias. *Journal of Economic Surveys*, 19(3), 309-345

[69] Stephan, P. (2013), *How Economics Shapes Science*. Harvard University Press, Cambridge, MA, USA.

[70] Stern, S. (2004). Do scientists pay to be scientists?. *Management Science*, 50(6), 835-853.

[71] Taubes, G. (1993). *Bad Science: The Short Life and Weird Times of Cold Fusion*. London: Random House.

[72] The Economist (2013). Trouble at the Lab. http://www.economist.com/news/briefing/21588057-scientists-think-science-self-correcting-alarming-degree-it-nottrouble.

[73] Wagenmakers, E. J., & Forstman, B. U. (2014). Rewarding high-power replication research. *Cortex*, 51(10),105-106.

[74] Wible, J. R. (2004). *The Economics of Science: Methodology and Epistemology as if Economics Really Mattered*. Routledge.

A Proofs[11]

Proof or Proposition 1. To see that $(e_L; nv)$ can be a Nash equilibrium, note 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. $\Delta e > \frac{(1-p)B_C}{\beta}$. In contrast, 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 $\frac{B_C}{\beta} \leq B_S$, 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 nv, but S's best response to nv is e_L.

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

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

from which we obtain:

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

For C to be indifferent between v and nv, the following equality must hold:

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

therefore:

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

r^* is positive, and $r^* \leq 1$ if $\Delta e \leq (1-p)B_S$. Moreover, $q^* \leq 1$ for all parameter values, and it is positive if $\Delta e \leq \frac{(1-p)B_C}{\beta}$. Because $\frac{B_C}{\beta} \leq B_S$, 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, $\Pr(\textit{high quality}) = q^* + (1-q^*)p$. 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 $\Pr(\textit{low quality and verified}) = (1-q^*)(1-p)r^*$.

[11] The proofs reported here are detailed but do not include step-by-step derivations in all cases, for the sake of space. More detailed, step-to-step versions are available from the authors.

Proof of Proposition 3 We start, again, from the pure-strategy equilibria. If S chooses e_H then C will choose nv. If C chooses nv, S will choose e_L. Thus, e_H cannot be part of a Nash equilibrium. If S chooses e_L, C will choose v if $\Delta e < \frac{(1-p)B_C}{\beta}$, and nv otherwise. If C chooses v then S will choose e_L as long as $\Delta e > (1-p)(B_S^H - B_S^L)$, and e_H otherwise. Therefore $(e_L; nv)$ is a Nash equilibrium for $\Delta e < \frac{(1-p)B_C}{\beta}$ and $(e_L; v)$ is the equilibrium for $(1-p)(B_S^H - B_S^L) < \Delta e < \frac{(1-p)B_C}{\beta}$.

Moving to the mixed-strategy equilibrium, r is the probability that C will play v and q the probability that S will make a high effort e_H. The indifference condition for S is:

$$B_S^H - e_H = r(1-p)[B_S^L - B_S^H] + B_S^H - e_L,$$

therefore $r^* = \frac{\Delta e}{(1-p)[B_S^H - B_S^L]}$. For C to be indifferent between v and nv it must be:

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

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

For $r^*$and q^* to be within the unit interval, we need $\Delta e \leq \min\left[(1-p)[B_S^H - B_S^L]; \frac{(1-p)B_C}{\beta}\right]$.

Proof of Proposition 4 If S chooses e_H then C will choose nv.If C chooses nv, S will choose e_L. Thus, e_H cannot be part of a Nash equilibrium. If S chooses e_L, C will choose v if $\Delta e < \frac{(1-p)B_C}{\beta}$, and nv otherwise. If C chooses v then S will choose e_L as long as $\Delta e > (1-p)B_S^v$, and e_H otherwise. Since $\frac{B_C}{\beta} \leq B_S^v$, $(e_L; nv)$ is therefore a Nash equilibrium in pure strategies if $\Delta e > \frac{(1-p)B_C}{\beta}$. As for the mixed-strategy equilibrium, the following indifference conditions must hold:

$$r(B_S^v - e_H) + (1-r)(B_S^{nv} - e_L) = r(pB_S^v - e_L) + (1-r)(B_S^{nv} - e_L);$$

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

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

Proof of Proposition 5 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 verify if $pB_C^H + (1-p)B_C^L - \beta\Delta e > 0$; i.e. $\Delta e < \frac{pB_C^H + (1-p)B_C^L}{\beta}$, 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 choose e_L because $B_S - e_H < B_S - e_L$. 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, nv) if $\Delta e > \frac{pB_C^H + (1-p)B_C^L}{\beta}$, and (e_L, v) if $(1-p)B_S < \Delta e < \frac{pB_C^H + (1-p)B_C^L}{\beta}$.

Regarding the mixed-strategy equilibrium, again 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 i) $\Delta e \leq (1-p)B_S$ and ii) $\frac{B_C^H}{\beta} \leq \Delta e \leq \frac{pB_C^H+(1-p)B_C^L}{\beta}$, therefore the equilibrium exists for $\frac{B_C^H}{\beta} \leq \Delta e \leq \min\left\{\frac{pB_C^H+(1-p)B_C^L}{\beta}; (1-p)B_S\right\}$.

Proof of Proposition 6 For the same logic as in the case of a single colleague, no pure-strategy equilibrium exists involving e_H by S. Suppose that only C_1 plays v against e_L. For this to be an equilibrium, it should be $B_S - e_H < pB_S - e_L$, i.e. $(1-p)B_S < \Delta e$, for player S, and $(1-p)B_C - \beta\Delta e > 0$, i.e. $\Delta e < \frac{(1-p)B_C}{\beta}$, for player C_1, which are incompatible conditions because $\frac{B_C}{\beta} \leq B_S$. For a similar argument, both C_1 and C_2 playing v is never an equilibrium. Finally, for $\{e_L; nv_1; nv_2\}$ to be a Nash equilibrium, we need $(1-p)B_C - \beta\Delta e < 0$, i.e. $\Delta e > \frac{(1-p)B_C}{\beta}$.

In the mixed-strategy equilibrium, S is indifferent between e_H and e_L if:

$$B_S - e_H = [1-(1-r_1)(1-r_2)]\, pB_S + (1-r_1)(1-r_2)B_S - e_L.$$

Imposing symmetry between C_1 and C_2 , i.e. $r_1 = r_2 = r$, the condition above is equivalent to:

$$B_S - e_H = \left[1-(1-r)^2\right] pB_S + (1-r)^2 B_S - e_L,$$

which admits $r^* = 1 - \sqrt{1 - \frac{\Delta e}{(1-p)B_S}}$ as unique positive solution. Notice that this probability is always lower than the corresponding probability for the case of a single colleague, i.e. $\frac{\Delta e}{(1-p)B_S}$, because $1 - \frac{\Delta e}{(1-p)B_S} > \sqrt{1 - \frac{\Delta e}{(1-p)B_S}}$ for $\frac{\Delta e}{(1-p)B_S} < 1$.

The indifference condition for C_1 (the condition for C_2 is symmetric) is:

$$(1-q)\left[r_2(1-p)\frac{B_c}{2\beta} + (1-r_2)(1-p)\frac{B_c}{\beta}\right] - \Delta e = 0.$$

Plugging $r_2 = r^* = 1 - \sqrt{1 - \frac{\Delta e}{(1-p)B_S}}$ and solving yields $q^* = 1 - \frac{\beta\Delta e}{(1-p)B_C\left[\frac{1}{2}+\frac{1}{2}\left(1-\sqrt{1-\frac{\beta\Delta e}{(1-p)B_S}}\right)\right]}$. This value is always lower than the corresponding value for the case of single colleague because $\frac{1}{2}+\frac{1}{2}\left(1-\sqrt{1-\frac{\beta\Delta e}{(1-p)B_S}}\right) < 1$. Finally, it is immediate to verify that the condition for r^* to lie in the unit interval is $\Delta e \leq (1-p)B_S$, while the condition for q^* is $\Delta e \leq \frac{(1-p)B_C\left[\frac{1}{2}+\frac{1}{2}\left(1-\sqrt{1-\frac{\beta\Delta e}{(1-p)B_S}}\right)\right]}{\beta}$, with the latter being stricter than the former.

Proof of Proposition 7 We first provide the proof part i) of the proposition. For $\Delta e \leq \min\left\{(1-p)B_S; \frac{B_C^H}{\beta}\right\}$, the unique Nash equilibrium in the stage game is (e_H, v).Without loss of generality, consider the possible deviation of 2 at $t=1$, playing v instead of nv. In the candidate equilibrium, 2 expects to make obtains $B_S - e_L$ in the periods when she plays S, and 0 (no verification) when he plays C:

$$0 + \delta\left(B_S - e_L\right) + 0 + \delta^3\left(B_S - e_L\right) + 0 + \delta^3\left(B_S - e_L\right) + ... = \frac{\delta}{1-\delta^2}\left(B_S - e_L\right);$$

by deviating, 2 obtains instead:

$$\left[pB_C^H + (1-p)B_C^L - \beta\Delta e\right] + \frac{\delta}{1-\delta^2}(B_S - e_H) + \frac{\delta^2}{1-\delta^2}(B_C^H - \beta\Delta e),$$

i.e. the payoff from playing v in $t=1$, and the discounted flow of Nash equilibrium payoffs afterwards. Thus the deviation is not profitable if

$$\frac{\delta}{1-\delta^2}\left(B_S - e_L\right) \geq \left[pB_C^H + (1-p)B_C^L - \beta\Delta e\right] + \frac{\delta}{1-\delta^2}(B_S - e_H) + \frac{\delta^2}{1-\delta^2}(B_C^H - \beta\Delta e),$$

which simplifies to:

$$\delta^2(1-p)(B_C^L - B_C^H) + \delta\Delta e - \left[pB_C^H + (1-p)B_C^L - \beta\Delta e\right] \geq 0$$

The solution for the inequality is $\delta \leq \underline{\delta}$ and $\delta \geq \overline{\delta}$,with $\underline{\delta}$ and $\overline{\delta}$ being respectively:

$$\underline{\delta} = \frac{-\Delta e - \sqrt{\Delta e^2 + 4(1-p)(B_C^L - B_C^H)\left[pB_C^H + (1-p)B_C^L - \beta\Delta e\right]}}{2(1-p)(B_C^L - B_C^H)};$$

$$\overline{\delta} = \frac{-\Delta e + \sqrt{\Delta e^2 + 4(1-p)(B_C^L - B_C^H)\left[pB_C^H + (1-p)B_C^L - \beta\Delta e\right]}}{2(1-p)(B_C^L - B_C^H)}.$$

It is immediate to check that $\underline{\delta}$ is negative (recall that $B_C^L \geq B_C^H$). As for $\overline{\delta}$, it is positive and smaller than 1 if $\Delta e > \frac{B_C^H}{1+\beta}$, which is a compatible condition in the parameter range that we are considering.

As for part ii), following the same logic as before the condition for "collusion" to be sustainable is

$$\frac{\delta}{1-\delta^2}\left(B_S - e_L\right) \geq \left[pB_C^H + (1-p)B_C^H - \beta\Delta e\right] + \frac{\delta}{1-\delta^2}(pB_S - e_L) + \frac{\delta^2}{1-\delta^2}(pB_C^H + (1-p)B_C^L - \beta\Delta e),$$

from which the condition reported in the Proposition is derived. To have $\frac{pB_C^H + (1-p)B_C^L - \beta\Delta e}{(1-p)B_S} < 1$ it must be $\Delta e > \frac{pB_C^H + (1-p)B_C^L}{\beta} - \frac{(1-p)B_S}{\beta}$, which is compatible with the parameter range that we are considering.

Proof of Proposition 8 Note first that $(e_H; v)$ is never an equilibrium because if S exerts high effort the probability to detect a low quality paper is 0. $(e_L; nv)$ is an equilibrium if $B_C^B - \beta\Delta e < 0$ and $pB_S^A + (1-p)B_S^B - e_L \geq B_S^A - e_H$.$(e_L; v)$ is an equilibrium if $pB_S^A + (1-p)B_S^B - e_L \geq B_S^A - e_H$ and $B_C^B - \beta\Delta e \geq 0$.Finally, $(e_H; nv)$ is an equilibrium if $B_S^A - e_H \geq pB_S^A + (1-p)B_S^B - e_L$.

We finally develop the extension to Proposition 8 mentioned in Section 4.4, with capacity constraints for high-quality journals. Define $x < 1$ as the probability that a high-quality paper is accepted in journal A. Moreover, define B_S^B as the credit obtained by S for publishing a high-quality or a non-verified paper in journal B, whereas we normalize to 0 the credit for a verified low-quality paper. The resulting payoff matrix and equilibria are summarized in the following Proposition.

Proposition 9 *The payoff matrix of this game is:*

		C: v	nv
S	e_H	$xB_S^A + (1-x)B_S^B - e_H;\ 0 - \beta\Delta e$	$xB_S^A + (1-x)B_S^B - e_H;\ 0$
	e_L	$p[xB_S^A + (1-x)B_S^B] - e_L;\ \frac{(1-p)}{(1-p)+(1-x)p}B_C^B - \beta\Delta e$	$xpB_S^A + (1-xp)B_S^B - e_L;\ 0$

The game has a unique Nash equilibrium. i) If $\Delta e \geq \max\left\{x(1-p)(B_S^A - B_S^B); \frac{(1-p)}{(1-p)+(1-x)p}\frac{B_C^B}{\beta}\right\}$, *then the pure-strategy Nash equilibrium is* $(e_L; nv)$*; ii) if* $(1-p)(xB_S^A + (1-x)B_S^B) \leq \Delta e \leq \frac{(1-p)}{(1-p)+(1-x)p}\frac{B_C^B}{\beta}$, *the pure-strategy Nash equilibrium is* $(e_L; v)$*; iii) If* $\Delta e \leq x(1-p)(B_S^A - B_S^B)$ *then the pure-strategy Nash equilibrium is* $(e_H; nv)$*;*

iv) if $x(1-p)(B_S^A - B_S^B) \leq \Delta e \leq \min\left\{(1-p)(xB_S^A + (1-x)B_S^B); \frac{(1-p)}{(1-p)+(1-x)p}\frac{B_C^B}{\beta}\right\}$, *the equilibrium is in mixed-strategy where* v *is played with probability* $r = \frac{\Delta e}{(1-p)B_S^B} - x\left(\frac{B_S^A - B_S^B}{B_S^B}\right)$, *and* e_H *is played with probability* $q = 1 - \frac{[(1-p)+(1-x)p]\beta\Delta e}{(1-p)B_C^B}$.

Proof. We first derive the payoffs in this modified version of the game. The likelihood of a publication in journal A is x and the likelihood of a high-quality paper is 1 if S plays e_H. Therefore, the expected payoff of S from playing e_H is $xB_S^A + (1-x)B_S^B - e_H$, regardless of the action of C. If S plays e_L, he receives a positive payoff only if the paper is of high quality, which occurs with probability p; moreover, the payoff depends on the type of journals where the article is published. Therefore S's expected payoff, when playing e_L and with C playing v, is $p[xB_S^A + (1-x)B_S^B]$. If C plays nv, the paper is published in journal A if it is of high quality (probability p) with probability x, and in journal B otherwise; the resulting expected payoff is $xpB_S^A + (1-xp)B_S^B - e_L$. The only relevant change to the expected payoff of C with respect to the basic game is in the case of S playing e_L and C playing v. Here C receives a positive payoff when detecting a low quality study; note that, if the paper is in journal A, then it is of high quality with probability 1; if, however, the paper is in journal B, it could be of either high or low quality. The probability of a low-quality paper, conditional on observing the paper in journal B when S plays e_L, can be expressed, applying the Bayes' rule, as:

$$prob(low\ quality,\ e_L|B\ journal) =$$

$$\frac{prob(B\ journal|low\ quality,\ e_L)*prob(low\ quality,\ e_L)}{prob(B\ journal|low\ quality,\ e_L)*prob(low\ quality,\ e_L)+prob(B\ journal|high\ quality,\ e_L)*prob(high\ quality,\ e_L)} =$$

$$\frac{1*(1-p)}{1*(1-p)+(1-x)*p} = \frac{(1-p)}{(1-p)+p(1-x)}.$$

As before, $(e_H; v)$ is never an equilibrium. $(e_L; nv)$ is an equilibrium if $\frac{(1-p)}{(1-x)p+(1-p)}B_C^B - \beta\Delta e < 0$ and $xpB_S^A + (1-px)B_S^B - e_L \geq xB_S^A + (1-x)B_S^A - e_H$. $(e_L; v)$ is an equilibrium if $xpB_S^A + p(1-x)B_S^B - e_L \geq xB_S^A + (1-x)B_S^A - e_H$ and $\frac{(1-p)}{(1-x)p+(1-p)}B_C^B - \beta\Delta e \geq 0$, and $(e_H; nv)$ is an equilibrium if $xB_S^A + (1-x)B_S^A - e_H \geq xpB_S^A + p(1-x)B_S^B - e_L$. Finally, the indifference conditions from which the probabilities of the mixed-strategy equilibrium are obtained are:

$$xB_S^A + (1-x)B_S^A - e_H = r\left[p(xB_S^A + (1-x)B_S^B) - e_L\right] + (1-r)\left[xpB_S^A + (1-px)B_S^B\right]$$

$$q(-\beta\Delta e) + (1-q)\frac{(1-p)}{(1-x)p+(1-p)}B_C^B - \beta\Delta e = 0.$$

r and q are in the unit interval for the relevant parameter space reported in the Proposition. ■