

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 those who attended our presentations at the University of Bologna, TOBB University, the 2014 Izmir Workshop on Scientific Misconduct and Research Ethics in Economics, the 2014 EARIE Meetings, the 2014 BRICK Conference and the 2014 DRUID Conference for their feedback. We also thank Astrid Marinoni for her 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Bralind Kiri, Nicola Lacetera, and Lorenzo Zirulia. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

ABSTRACT

Building on previous research to reinforce findings or point out limitations is essential for a healthy working of the scientific community because it allows science to self-correct and evolve, thus providing a more solid knowledge base to individuals, firms and societies. In this paper we propose a model to investigate the incentives of scientists to perform these activities of control and criticism when these activities, just like the production of high-quality research in the first place, are costly, and we study the strategic interaction among these incentives. We show that a certain fraction of low-quality scientific knowledge characterizes all the equilibria in the basic version of model. In fact, the absence of (detected) low-quality research in a field can be interpreted as the lack of verification activities and thus as a potential limitation to the reliability of that field. We also derive that facilitating incremental research and verification activities improves the expected quality of newly produced knowledge; this effect, however, is contrasted by the incentives to free ride on performing verification if many scientists are involved in it, and also might discourage scientists to undertake new research in the first place. Finally, the findings imply that softening overall incentives to publish does not enhance research quality, although it increases the fraction of low-quality papers that are identified. We also elaborate empirical predictions from the model and strategies to test them, and discuss the implications 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

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

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

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 has great value not only within the ivory tower of academia. Firms and investors evaluate business opportunities also on the basis of the science underlying a new product, process or service; in particular, venture capitalists and R&D-intensive companies regularly "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, it is widely accepted that scientific knowledge is a powerful engine of economic growth and social welfare (Romer, 1990; Stephan, 2012). It is therefore not surprising that debates about the reliability of research are not confined to the boundaries of specialists in the scientific community and have reached firms, policymakers and the public opinion (*Economist*, 2013).

Extreme cases in which science may "go wrong" include outright fraud or major mistakes that, if detected, lead to retraction from publication (Azoulay et al., 2015a-b; Broad and Wade, 1982; Lacetera and Zirulia, 2011; Lu et al., 2013). Most often, however, flaws, limitations and mistakes in a study just occur as "natural" steps toward better theories and findings. Karl Popper's view of science, for example, holds that a finding or theory can be defined as scientific to the extent that it is falsifiable (Popper, 1959). Therefore, at each given time, the body of scientific knowledge 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).

The history of science provides many examples of how accepted findings were challenged by subsequent research. 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 were made through progressive criticisms and falsifications of earlier results, for example obtained with less reliable empirical strategies (Holmberg, 2008). In yet 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). In a fascinating set of accounts, Livio (2014) describes "blunders" by some of the greatest scientists in history. 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. A further interesting case is given by the research on climate change and global warming. Although there is increasing acceptance that climate is changing and the nature of the change is anthropogenic, counterarguments and evidence of scholars who are more skeptic are contributing to improve the overall reliability of research in this area, with the obvious positive consequences that this bears (Sherwood, 2011). An 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 cited to point out limitations and qualifications of their claims about the origins, features and extent of climate change and its mechanisms.¹

Building upon previous research and potentially identifying its limitations thus appears as essential for a healthy working of the scientific community (Carpi and Eggers, 2011). But what are the incentives of scientists to perform research on existing, established topics to potentially exert control and criticism? Will these activities always improve upon or correct previous findings, or, conversely, shall we expect some degree of imprecision at any given time? How do these incentives interact with those to produce novel, high-quality findings? And what determines the incidence of imperfect science and the effort to improve upon it?

To answer these questions, we propose 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. Although quite simple, the model, described in detail below, delivers a number of results and comparative statics with several implications. They are summarized here:

1. 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

¹The papers (limited to research articles and letters) were parsed and phrases including a citation were isolated. Using natural language and sentiment analysis and based on an initial training set, we isolated citations that contained a negative element (contradicting previous evidence, criticizing methods, and so on). See Catalini et al. (2015) for the application of a similar methodology and for more details.

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.

2. 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. In contrast, reducing the value of a publication for the knowledge originator, as some scholars have suggested (for example by softening the "publish or perish" paradigm), does not have an impact on research quality, although it increases the fraction of low-quality papers that are identified.
3. The performance of verification activities by a high number of scientists may lead to the reduction of overall verification activities and expected quality of research if individual rewards from scrutiny are lower because they are shared among colleagues.
4. In scientific communities where interactions are frequent, scientists may "collude," i.e. avoid verifying each other's research and save on the investment required by expensive experimental procedures.
5. 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. There are, therefore, also limitations in promoting widespread investments in incremental research, replication activities, and so on.

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. In that paper it is assumed that the project's probability of success is exogenous: in case of an unsuccessful project, the scientist can nevertheless submit a paper, thus committing a fraud. Here the probability of a paper being of high quality is endogenous, because it depends on the scientist's effort. This different assumption significantly changes the nature of the game, as well as the results.

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; 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.

In the next section we present the basic model and discuss its assumptions. Section 3 reports the equilibria of the game and their consequences on the expected quality of research. A first set of extensions with modifications to the payoff structure is in Section 4, and modifications of the model's overall structure are studied in Section 5. In Section 6 we discuss the implications of the model for firms and policymakers interested in following the development of scientific discoveries; we also analyze some recent initiatives to propose more debates and evaluations of published studies, using the results of the model; finally, we propose empirical predictions emerging from our theoretical analysis, and strategies to test these predictions. Section 7 concludes, and all proofs are in the Appendix.

2 The model

2.1 The basic game

There are two 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. Through his choice of effort, S affects the quality of the knowledge that he produces, which can be high or low. A high-quality paper, if scrutinized by C , does not show errors or significant lack of robustness. Otherwise, the paper is of low quality. Absent C 's verification, high-quality and low-quality papers are indistinguishable both for S and C , and therefore provide the same benefit to the players.

More formally, S chooses between high effort (e_H) and low effort (e_L). If S chooses e_H , then the paper is of high quality with probability 1; if S chooses e_L , then the paper is of high quality with probability $p \in [0, 1)$. e_H and e_L denote both the feasible actions for S , and their associated costs, with $e_H \geq e_L \geq 0$. C chooses between verifying the quality of the results by S , which we denote as action v , or not verify (action nv). If C chooses v , then 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. Following 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. Both players are risk-neutral.

The effort choice of S is not observed by C . Therefore, this is a game of imperfect information akin to a simultaneous-move game, with Nash equilibrium as a solution concept. The payoff matrix in normal form is presented in Table 1.

| | | | |
|-----|-------|--|----------------|
| | | 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 .

2.2 A discussion of the model's assumption

Before we solve the basic game and explore its implications, it is worth discussing the key assumptions 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.² 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. This assumption (which we relax in one extension in Section 4) implies that S does not take into account the expected quality of the paper, which he knows given the effort exerted. Thus, any intrinsic reward from high quality that S may receive is not considered here, although the same effect may be captured by a lower value of the cost e_H . In our interpretation, the value of B_S can be seen as primarily influenced by the prestige of the journal where the research is published, by the institutional context (such as the "publish or perish" culture) or by personal characteristics of S , such as his career stage.

Third, notice that B_S does not depend on effort. Therefore, higher effort does not lead to "better" scientific results, e.g. results that are more general or more relevant in some dimensions and that could lead to more cited publications, or appearing in more prestigious journals (Ellison, 2002). In our model, higher quality is associated to a characteristic of research, i.e. its reliability, which becomes evident only if the paper is scrutinized. High effort by S , therefore, can be interpreted as "internal replication" (Hamermesh, 2007) and for that 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 (Wible, 1998) or scientific (Hamermesh, 2007) replication, where a different experiment or analysis is conducted, but in a way that might inform about the solidity of the original

²To 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.

result.³ 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.

Finally, the parameter β measures the relative magnitude of verification cost for C with respect to "internal" verification by S . 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, which can be invalidated by a single counterexample or by identifying a logical error in the proof. However, we assume that β is not too low, i.e. $\beta \geq \frac{B_C}{B_S}$.⁴ β can also be affected by 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. Note also that assuming that the quality of the original paper is known with certainty after C 's verification excludes from the analysis the fact that verification activities themselves are subject to uncertainty about their reliability.

3 Results

3.1 The equilibria of the game

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.

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*

³The recently started initiative "Behavioral Economics Replication Project" (<http://sciencepredictionmarkets.com/>) is an example of this type of replication exercises.

⁴The fact itself of modeling the verification costs of C with Δe (up to a multiplicative factor) might come across as a strong restriction and simplification. However, note first 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. Second, as explained here, it makes sense to establish a simple comparison between the "internal verification costs" by S and the verification by an external peer.

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 graphically represents the equilibrium regions in the $(\Delta e; B_C)$ space.

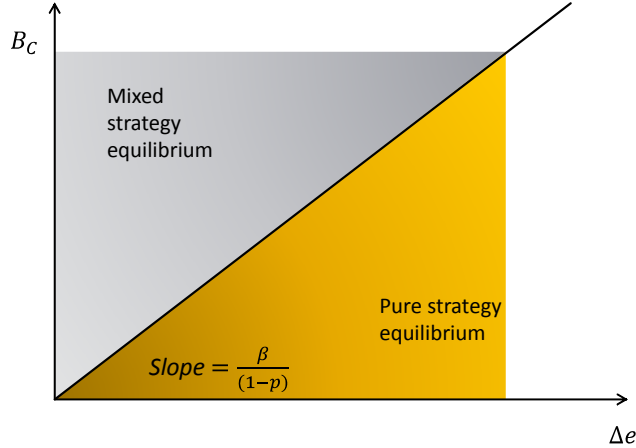


Figure 1: The equilibria of the basic game in the $(\Delta e, B_C)$ space.

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

A first implication of Proposition 1 is that a situation in which low-quality papers have zero probability to be produced is not an equilibrium. In other words, errors, flaws, limitations and other forms of unreliable or incomplete results are a distinctive feature of the scientific endeavor as captured by our model. If the incentives for verification are low ($\Delta e \geq \frac{(1-p)B_C}{\beta}$), then low-quality papers are not identified (verification does not occur), and constitute a fraction $(1-p)$ of all papers. Conversely, if the incentives for verification are high enough ($\Delta e < \frac{(1-p)B_C}{\beta}$), then verification activities are performed with positive probability: in turn, this leads S to exert high effort with positive probability.

From Proposition 1, we also determine the likelihoods of two events that will be the subjects of our comparative exercises below: i) the probability that a paper is of high quality, which is a measure of actual reliability of scientific knowledge (independently of what is observed); and ii) the probability that papers are of low quality and are identified as such.

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

$$\Pr(\text{high quality}) = \begin{cases} 1 - \frac{\beta\Delta e}{B_C} & \text{if } \Delta e \leq \frac{(1-p)B_C}{\beta} \\ p & \text{if } \Delta e > \frac{(1-p)B_C}{\beta} \end{cases} . \quad (1)$$

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

$$\Pr(\text{low quality and identified}) = \begin{cases} \frac{\beta(\Delta e)^2}{(1-p)B_S B_C} & \text{if } \Delta e \leq \frac{(1-p)B_C}{\beta} \\ 0 & \text{if } \Delta e > \frac{(1-p)B_C}{\beta} \end{cases} . \quad (2)$$

Figures 2 and 3 below report $\Pr(\text{high quality})$ and $\Pr(\text{low quality and identified})$ as a function of Δe , for different values of B_C and for $\beta = 1$.

The probability that a paper is of high quality is non-increasing in the verification cost Δe (Figure 2). If Δe is large (relative to C 's expected gain from verification), then no verification occurs, and the fraction of high-quality papers only depends on the exogenous probability p . If Δe is low (with respect to C 's expected gain from verification), then the lower Δe , the larger the fraction of high-quality papers because exerting higher effort is less costly for S . Note, however, that verification activities by C , although being less costly, are less frequent because the probability to find a low-quality paper is smaller.

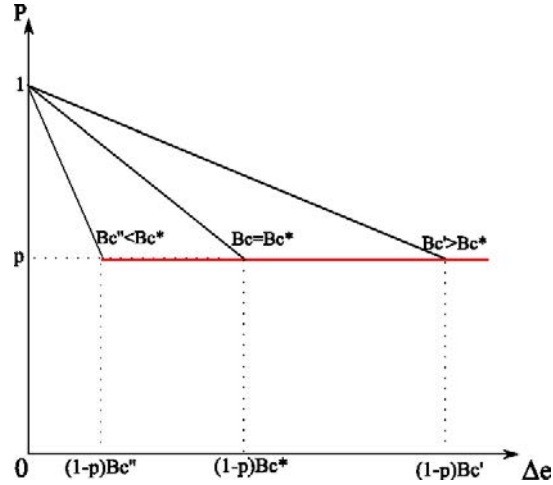


Figure 2: Probability of a high-quality paper (P) for different values of Δe and B_C (e.g. for three hypothetical levels B_C'' , B_C^* and B_C'). P is equal to p for $\Delta e > (1-p)B_C$.

In the mixed-strategy equilibrium region, moreover, larger benefits B_C from identifying a low-quality paper (or lower costs via a reduction in β) increase the fraction of high-quality papers; because verification is more rewarding (or less costly), S increases his effort in order to reduce C 's incentives to verify. However, the probability that the paper is of high quality does not depend on B_S , i.e. the value of a publication. To understand the intuition for this, consider that for a given intensity of C 's control (i.e. for given value of r), the marginal effect of B_S on the S 's payoff is 1 when he exerts high effort, and $rp + (1-r) < 1$ when he exerts low effort, because in this case S must take into account that the value of publication is lost

if the 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 of this, 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.

We next look at the probability that a paper is of low quality and is identified as such (Figure 3).

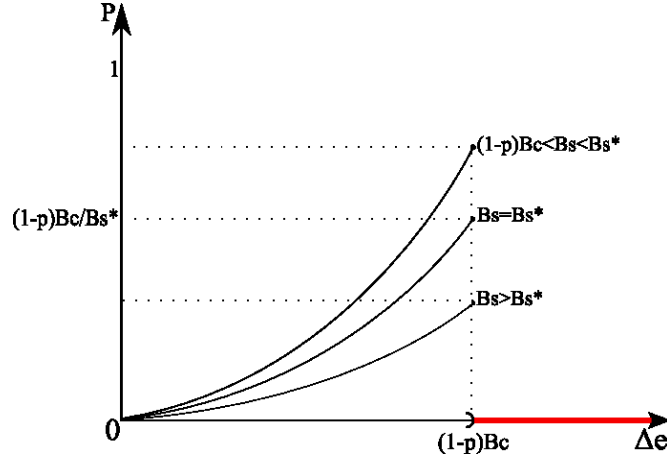


Figure 3: Probability P that a low-quality paper is identified, for different values of Δe and B_S . P drops to zero for $\Delta e > (1 - p)B_C$.

First, note that this probability is non monotone and discontinuous as a function of Δe ; it is positive and increasing in Δe when Δe is "low" and zero when Δe is "high," because we enter the pure-strategy equilibrium region. Combining Figure 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 (false, flawed, 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.

However, in the mixed-strategy equilibrium region, lower verification costs Δe imply both higher expected quality and lower rate of low-quality discovery; this occurs because both low-quality papers and verification activities are less frequent. The same effect occurs for an increase in B_C , but in this case this is due to the increased effort by S . The probability of verification by C , indeed, does not depend on B_C ; an increase in the reward from detecting a low-quality paper operates as a threat that "disciplines" S 's behavior, leaving the intensity of

the verification activity unaffected. As a consequence, observing a reduction of the frequency of low-quality papers that are discovered (but not their disappearance) is associated to higher expected quality of research.

Note, finally, that 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). The intuition is that, when publications are more valuable, the opportunity cost of low effort is also higher; as a consequence, C may save on verification activities while leaving S indifferent between high and low effort.

4 Extensions – 1: rewarding confirmatory, low-quality, and confirmed research

The model can be extended in several directions to capture additional aspects or policies that characterize the scientific community and the production of research. In this section we discuss modifications to the payoff structure, of which the basic model is a special case.

4.1 A value for merely confirmatory results

Suppose that C obtains a positive reward even when verifying a paper that is of high quality. In this case, the normal form of the game is expressed by the payoff matrix in Table 2, where B_C^H (B_C^L) corresponds to the value for C of discovering a high (low)-quality publication. We assume $B_C^H \leq B_C^L$; although positive, the benefit from confirming a high-quality result is not greater than spotting a lower-quality study.

| | | | |
|-----|-------|---|----------------|
| | | C | |
| | | v | 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 2: 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 .

The payoff structure above implies that, in this case, it is possible to sustain an equilibrium where no low-quality papers are produced. The proposition below formalizes the solution.

Proposition 3 *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 prefers to verify even if the research by S is of high quality with probability 1. Thus, if confirmatory results are positively valued by the scientific community, it is possible that low-quality papers are not produced. Moreover, 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.

4.2 A value for low-quality research

In the second extension, we assume that S obtains a positive benefit also from low-quality research. Denote the benefit from high-quality (low-quality) research as B_S^H (B_S^L), with $B_S^L \leq B_S^H$. The payoff matrix is in Table 3.

| | | | |
|-----|-------|--|------------------|
| | | C | |
| | | v | 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 3: 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 .

This extension allows us to consider two aspects that were not captured in the basic version of the model. First, it is not uncommon that, although a scientific result is not fully confirmed by subsequent research, it may nevertheless maintain some validity; the scientist 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). Second, 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 takes place, so that the non-confirmatory result cannot be taken as an undisputable proof of the unreliability of the original research.

The proposition that follows presents the equilibria.

Proposition 4 *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 (being $p < 1 - \frac{\beta \Delta e}{B_C}$ since $\Delta e < \frac{(1-p)B_C}{\beta}$), whereas the fraction of low-quality research that is identified as such is higher (because $\frac{\Delta e^2}{(1-p)[B_S^H - B_S^L]B_C} > \frac{(\Delta e)^2}{(1-p)B_S B_C}$ 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.

4.3 A value for confirmed results

We now allow S to obtain a higher reward when his research is of high quality and is verified (and thus confirmed in its validity). 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 $B_S^v \geq B_S^{nv}$.⁵ Table 4 reports the payoff matrix for this version of the game and Proposition 5 provides the solution.

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

Table 4: 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 .

Proposition 5 *The game has a unique Nash equilibrium. i) If $\Delta e > \frac{(1-p)B_C}{\beta}$, then the unique, pure-strategy Nash equilibrium is $(e_L; nv)$; 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 $r^* = \frac{\Delta e}{(1-p)B_S^v}$.*

⁵In line with the basic model, we assume $\frac{B_C}{\beta} \leq B_S^v$.

Propositions 1 and 5 coincide for $B_S^v = B_S$. By distinguishing between B_S^v and B_S^{nv} , Proposition 5 further shows that the probability that C scrutinizes depends on the reward for confirmed high quality. 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 .

5 Extensions – 2: The structure of the game

The analysis so far has concerned a static, two-player game. This was a natural starting point that, although very simple, already pointed out a number of non-obvious trade-offs and results concerning, in particular, the expected quality of research that is produced. The extensions that we now propose modify some major features of the game’s structure in order to model additional, relevant aspects of the operating of the scientific community. First, we let more than one scientist perform verification activities. Second, we analyze a repeated version of the game where scientists take in turn the role of originator and scrutinizer of new results. Finally, we endogenize the decision of S to conduct research in the first place.

5.1 More than one verifying player

Research findings are usually built upon and verified by more than one other scientist. The model considered in this section modifies the basic set-up by considering the simplest extension in this direction, i.e. the presence of two researchers, C_1 and C_2 , 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. 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 .⁶ The equilibrium 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,

⁶For $\frac{(1-p)B_C \left[\frac{1}{2} + \frac{1}{2} \left(1 - \sqrt{1 - \frac{\beta \Delta e}{(1-p)B_S}} \right) \right]}{\beta} < \Delta e \leq \frac{(1-p)B_C}{\beta}$, 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.

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.

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^* .

5.2 Repeated interactions

A second realistic modification of the basic model is to allow each player to be in the position of the originator of new research as well as of the scrutinizer, and for this to happen multiple times. Consider two researchers, 1 and 2, interacting repeatedly for $t = 1, 2, \dots, \infty$. 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 described in Section 4.1, which allows for equilibria in pure strategies where verification occurs. 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 $\delta \geq \frac{-\Delta e + \sqrt{(\Delta e)^2 + 4(1-p)(B_C^L - B_C^H)[pB_C^H + (1-p)B_C^L - \beta\Delta e]}}{2(1-p)(B_C^L - B_C^H)}$. (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. It is interesting to observe how Propositions 6 and 7 make opposite predictions as of the impact of community size. Also, we should expect findings to be less reliable in scientific communities which are isolated and base recognition on "local" journals.

5.3 Modeling a scientist's decision to conduct research in the first place

So far we assumed that scientist S would always exert effort. In fact, a scientist can always decide to not start a project at all. 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. The game that includes the entry decision is represented in extensive form in Figure 4.

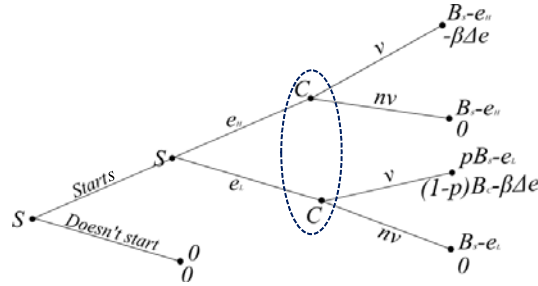


Figure 4: The basic game with the entry decision: extensive form representation. The top payoff at each end node refers to S , and the bottom payoff to C . The dotted circle indicates nodes that are part of the same information set.

The game is solved by backward induction. The subgame following the decision of S to start the project is equivalent to the basic game introduced in Section 2.1, thus the expected payoff of S depends on which type of equilibrium prevails. 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 her possible strategies). Therefore, S will 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 provide a graphical representation of the decision to start a project in the $(B_S, \Delta e)$ space in Figure 5.

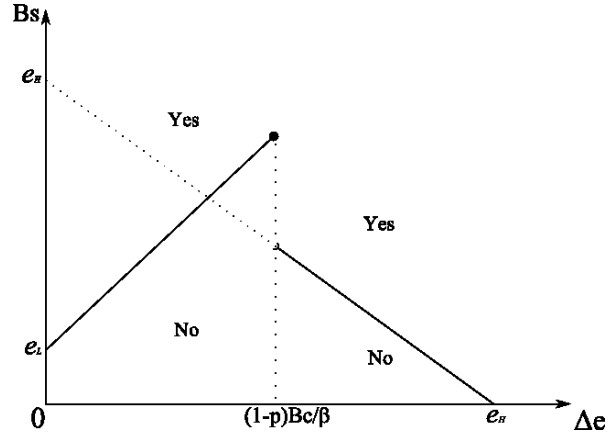


Figure 5: The decision to start a project. S starts a project in the regions indicated with "Yes", and does not start a project in regions denoted with "No".

Note from the graph that, 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; 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 could instead backfire if the decision of S to start a project is considered.⁷

⁷Mueller-Langer and Andreoli Versbach (2014) propose a model where mandatory (and immediate) data

6 Assessing and improving the quality of scientific production: implications and predictions

"Jack, if you think you have a good idea, publish it! Don't be afraid to make a mistakes. 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).

In recent years, numerous discussions about the functioning of the scientific community have concerned how to enhance the quality and reliability of research. Our model provides insights on some of the proposals that have been advanced as well as about some of the current trends in the scientific community. The model also informs companies and investors exploring scientific advances for business opportunities.

Some authors have identified a main driver of potential scientific unreliability in the lack of proper incentives towards incremental research and replication. For example, the 2013 Medicine Nobel Laureate Randy Schekman announced that he would not send his papers to some of the major journals, in particular because they excessively select "novel," "newsworthy" findings at the expense of rigor and depth of inquiry, which requires instead additional incremental work (Schekman, 2013). Similar considerations were expressed in the past by other prominent scientists (see for example Ioannidis, 2014). Several initiatives have been undertaken to enhance further reviews and research on existing studies. An increasing number of journals (as well as public funding agencies) now have data-sharing policies. The platform PubPeer (<https://pubpeer.com/>), allows scientists to review, comment and potentially propose corrections to published papers. Post-publication comments are also encouraged in the journals of the Public Library of Science, especially *PLOS ONE* (<http://www.plosone.org/>), whereas other journals, such as the *New England Journal of Medicine*, include a "Journal Watch" section on the website to stimulate the collection and discussion of interesting published findings (<http://www.jwatch.org/>). Also, the "Behavioral Economics Replication Project" (<http://sciencepredictionmarkets.com/>) is an initiative to promote replication of published lab experimental studies in economics, associated to a "prediction markets" 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. Increasing the expected rewards from verification activities (via an increase in B_C) or reducing their costs (via a reduction in β) increase the expected quality of scientific knowledge. This does not happen because verification is indeed more frequent, but because knowledge originators exert higher disclosure policies might inhibit researchers to undertake research in the first place.

effort in response. The analysis in Section 4 further qualifies this claim, by showing that a positive utility from confirmatory results (which, to a large extent, does not characterize the current incentive structure of the scientific community) might lead to the reduction of low-quality papers, while keeping verification activities in place. In terms of discovery of low-quality research, our model suggests that we should expect an increase, if starting from a situation when verification activities were not performed; or a decrease otherwise. A limit of these policies, however, is that making incremental or replication research less costly might discourage undertaking certain research in the first place. 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. So for example, mandatory data disclosure policies may backfire. We also showed that a higher number of potential verifiers might crowd out incentives to perform these activities at all. Again, this points out a countervailing (and somewhat counterintuitive) effect of verification activities that will need to be considered in devising optimal policies. Recent debates about psychology scholars on how too much attention to "dissecting" existing studies might come at the expense of more innovative exploration appear to highlight the tradeoffs derived in the analysis above (Bartlett, 2014).

Another frequently held belief, sometimes considered equivalent to the one just discussed, is that the quality of scientific research may be negatively affected by too high-powered incentives to publish; proposals have therefore been advanced to soften the "publish-or-perish" paradigm (Abelson, 1985; Giles, 2007; Schekman, 2013). Our model shows, however, that acting directly on the incentives to publish may be different from increasing the incentives for incremental or confirmatory research, for example. In particular, the expected quality of research would be unaffected by softening publication incentives alone (a reduction in B_S). However, the fraction of low-quality papers that could be recognized as such would increase when publication incentives are weaker because of a more intense control activity by C . A similar effect (i.e. an increase in verification activities by C due to lower publication incentives) is obtained when also low-quality research provides a positive benefit. More scrutiny will be socially valuable as well, since it will reduce the uncertainty concerning the quality of scientific research. Also, a reduction in the publish-or-perish attitude can be interpreted as a reduction of the relative value of (supposedly) path-breaking research with respect to more incremental research, causing a simultaneous decrease in B_S and increase in B_C . This, as demonstrated above, would simultaneously increase research quality and the identification of low quality research.

As for companies and investors scouting the scientific landscape in search for opportunities, the main message from our analysis 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 where incremental research is observed may not necessarily be a mature or declining area; it 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.

These considerations 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.

7 Conclusions

Our model conveys a number of relevant 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 the incentives provide further work upon and, in the process, possibly question existing and established results. We also show an interesting asymmetry of effects between lowering incentives to produce new research, and increasing incentives to do additional work and verification on existing findings. Finally, we point out some countervailing effects of encouraging verification activities.

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.

References

- [1] Abelson, P. (1990). Mechanisms for evaluating scientific information and the role of peer review. *Journal of the American Society for Information Science*, 41(3), 216-222.
- [2] Aghion, P., Dewatripont, M., & Stein, J. C. (2008). Academic freedom, private-sector focus, and the process of innovation. *The RAND Journal of Economics*, 39(3), 617-635.
- [3] Azoulay, P., Furman, J. L., Krieger, J. L., & Murray, F. E. (2015a). Retractions. *Review of Economics and Statistics*, forthcoming.
- [4] Azoulay, P., Bonatti, A. & Krieger, J. L. (2015b): The career effects of scandal: Evidence from scientific retractions. Working paper.
- [5] 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.
- [6] 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.
- [7] Bartlett, T. (2014). Replication crisis in psychology research turns ugly and odd. *The Chronicle of Higher Education*, June 23.
- [8] 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.
- [9] Broad, W., & Wade, N. (1982). *Betrayers of the Truth. Fraud and Deceit in the Hall of Science*. Simon and Schuster, New York.
- [10] Catalini, C., Lacetera, N. & Oettl, A. (2015). Beyond citation counts: Exploring negative citations in science. Working paper.
- [11] Collins, H. (1992). *Changing Order: Replication and Induction in Scientific Practice*. Chicago: The University of Chicago Press.
- [12] Dasgupta, P., & David, P. A. (1994). Toward a new economics of science. *Research Policy*, 23(5), 487-521.
- [13] Ellison, G. (2002). Evolving standards for academic publishing: a q-r theory. *Journal of Political Economy*, 110 (5),994-1034.

- [14] Giles, J. (2007). Breeding cheats. *Nature*, 445(18): 242–243
- [15] Hamermesh, D. S. (2007). Replication in economics. *Canadian Journal of Economics*, 40(3), 715–733.
- [16] Hanson, R. (1990). Could gambling save science? *Proceedings of the Eighth International Conference on Risk and Gambling*, London, July.
- [17] Hartemink, A. E. (2000). Publish or Perish (3)–Fraud and ethics. *Bull Internat Union Soil Sci*, 97, 36–45.
- [18] Häussler, C., Jiang, L., Thursby, M. & Thursby, J. (2014). Specific and general information sharing among competing Academic researchers. *Research Policy*, 43(3), 465–75.
- [19] Henry, E. (2009). Strategic disclosure of research results: the cost of proving your honesty. *The Economic Journal*, 119, 1036–1064.
- [20] Henry, E., & Ottaviani, M. (2013). Research and the approval process. Working paper.
- [21] Holmberg, S. D. (2008). *Scientific Errors and Controversies in the HIV/AIDS Epidemic*. Greenwood Publishing Group.
- [22] Howson, C., & Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach*. Open Court Publishing Co.
- [23] 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.
- [24] Ioannidis, J. P. A. (2014). How to make more published research true. *PLOS Medicine*, 11(10), e1001747
- [25] Jensen, R. & Thursby, M. (2001). Proofs and prototypes for sale: The licensing of university inventions. *American Economic Review*, 91(1), 240–259.
- [26] Lacetera, N., & Zirulia, L. (2011). The economics of scientific misconduct. *Journal of Law, Economics, and Organization*, 27(3), 568–603.
- [27] 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.
- [28] Lakhani, K. R., & Von Hippel, E. (2003). How open source software works: “Free” user-to-user assistance. *Research Policy*, 32(6), 923–943.
- [29] 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.
- [30] Lu, S. F., Jin, G. Z., Uzzi, B., & Jones, B. (2013). The retraction penalty: Evidence from the Web of Science. *Scientific Reports*, 3, 3146.
- [31] 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.
- [32] Merck (2015). Experts in your field, close to your location: <http://www.merck.com/licensing/contact-us/home.html> (last accessed 04/10/2015).
- [33] Mialon, H. (2010). What’s in a title? Predictors of the productivity of publications. working paper.

- [34] Mirowski, P., & Sklivas, S. (1991). Why econometricians don't replicate (although they do reproduce). *Review of Political Economy*, 3(2), 146-163.
- [35] Moser, P., Voena, A. & Waldinger, F. (2014). German Jewish émigrés and US invention. *American Economic Review*, 104(10), 3222-55.
- [36] 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.
- [37] Muma, J. R. (1993). The need for replication. *Journal of Speech, Language and Hearing Research*, 36(5), 927.
- [38] Pfizer (2015). Innovative partnering models:
http://www.pfizer.com/partnering/areas_of_interest/innovative_partnering_models (last accessed 04/10/2015).
- [39] Popper, K. R. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.
- [40] Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5 pt. 2), S71-S102.
- [41] Rousseau, D. L., & Porto, S. P. S. (1970). Polywater: Polymer or artifact?. *Science*, 167(3926), 1715-1719.
- [42] Ryan, P.K. (2013). *How Venture Capital Works*. Rosen Publishing.
- [43] Sauermann, H. & Stephan, P. (2013). Conflicting logics? A multidimensional view of industrial and academic science. *Organization Science*, 24(3), 889-909.
- [44] 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>
- [45] Sherwood, S. (2011). Science controversies past and present. *Physics Today*, 64(10), 39-44.
- [46] Stephan, P. (2013), *How Economics Shapes Science*. Harvard University Press, Cambridge, MA, USA.
- [47] Stern, S. (2004). Do scientists pay to be scientists?. *Management Science*, 50(6), 835-853.
- [48] Taubes, G. (1993). *Bad Science: The Short Life and Weird Times of Cold Fusion*. London: Random House.
- [49] Wible, J. R. (2004). *The Economics of Science: Methodology and Epistemology as if Economics Really Mattered*. Routledge.

A Proofs⁸

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. $\Delta e > \frac{(1-p)B_C}{\beta}$. Note, in contrast, that the strategy pair (e_L, v) is never an equilibrium; if S chooses e_L , then C prefers v if $(1-p)B_C - \beta\Delta e > 0$, i.e. $\Delta e < \frac{(1-p)B_C}{\beta}$. However, in order for S to play e_L in response, the condition is that $B_S - e_H < pB_S - e_L$ or $\Delta e > (1-p)B_S$. Because by assumption $\frac{B_C}{\beta} \leq B_S$, 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:

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

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:

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

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, $\Pr(\text{high quality}) = q^* + (1-q^*)p$. 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 $\Pr(\text{low quality and verified}) = (1-q^*)(1-p)r^*$.

Proof of Proposition 3 First, we determine the existence of pure-strategy equilibria. If S chooses e_H , C will play v as long as $B_C^H - \beta\Delta e > 0$, i.e. $\Delta e < \frac{B_C^H}{\beta}$, and nv otherwise. 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

⁸The 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.

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 4 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 5 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 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:

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

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 :

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

by deviating, 2 obtains instead:

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

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

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

which simplifies to:

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

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

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

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

It is immediate to check that $\underline{\delta}$ is negative (recall that $B_C^L \geq B_C^H$). As for $\bar{\delta}$, 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

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

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.