Mini-Symposium on Economic Approaches to Science:

Max Albert*

Methodology and Scientific Competition

ABSTRACT

Why is the average quality of research in open science so high? The answer seems obvious. Science is highly competitive, and publishing high quality research is the way to rise to the top. Thus, researchers face strong incentives to produce high quality work. However, this is only part of the answer. High quality in science, after all, is what researchers in the relevant field consider to be high quality. Why and how do competing researchers coordinate on common quality standards? I argue that, on the methodological level, science is a dynamic beauty contest.

1. THE PROBLEM

In this paper, I am concerned with a simple question: Why is the quality of scientific research so high on average?[1]

For those who accept the premise, the answer is probably obvious. Science as an institution is highly competitive. By science as an institution, I mean academic or open science, that is, the whole system of research-oriented universities, scientific journals, the peer review system, learned societies, and so forth. Open science is to be distinguished from proprietary science, that is, research under the protection of intellectual property rights like patents.[2]

In open science, the production of high-quality research is the way to rise to the top. Researchers compete for positions, research grants, journal space, and status. Thus, there is an incentive to aim at high quality but, of course, no sure-fire method to produce it, which explains why we also observe low quality.

However, pointing out that science is competitive is only part of the answer. High-quality contributions to science are rewarded. But there is no central agency deciding on the standards according to which contributions are to be judged.

Episteme 8.2 (2011): 165–183
DOI: 10.3366/epi.2011.0014

www.eupjournals.com/epi

So what is high quality? Where do quality standards come from? How are they enforced? What kind of standards can become established in science? Do they further or hinder the pursuit of truth or other aims – the aims outsiders have in mind when they admire the results of scientific research?

Quality standards in science emerge from the same competitive process as the scientific results that are judged according to these standards. The quality standards in a field are an important part of the field's methodology. Some methodological standards are specific to certain fields of research, others are common to many. By and large, competing researchers in a field accept common methodological standards. In most cases, these standards are demanding in the sense that high quality is not easy to produce[3] – otherwise, the quality standards could not be used to distribute scarce rewards.[4]

How can we explain the existence of these common standards? A conceivable explanation is that methodological standards are imported into science by scientific novices. After all, researchers believe in ideals and accept standards even before they enter science. Curiosity is often considered an important motivation for choosing a career in science, and it is hard to see how somebody could be curious without being interested in the truth about matters. Hence, a love of truth might be imported into science by those who enter science.

However, neither curiosity nor a love of truth is sufficient for judging the quality of a research paper. No statement – or, at the very least, no statement of scientific interest – can ever be proved or refuted, verified or falsified conclusively. Hence, even for those who just seek truth, it is not obvious how to judge the products of science, as everybody knows who has experience with the refereeing process (see, e.g., Seidl et al. 2008, Albert and Meckl 2008). It requires training to make these judgments. Methodologies are not easily explained to outsiders, and they are a matter for debate among insiders. Both facts suggest that methodological standards are endogenous to science and not just imported by scientific novices.

If methodological standards are not imported into science by scientific novices, they might be imposed from the outside, by final consumers of science. This is the traditional explanation for the quality of consumer goods. Consider the case of perfect competition, where among other conditions it is assumed that all parties are equally informed about the quality of all goods. Different qualities of the same good can then just be treated as different goods. In equilibrium, consumers select the qualities they want at given prices, and producers select inputs of a certain quality because these inputs allow them to satisfy consumers' wishes at minimal costs. This explanation extends to a production process with intermediate stages where the input quality selected upstream is determined by the output quality demanded downstream.

However, open science, in contrast to proprietary science, is not a market.[5] It uses the so-called voluntary contribution mechanism (VCM), not the price mechanism. Researchers publish their ideas and results, which can then be used free of charge by anybody who wishes to do so. Research outputs are non-rival goods anyway; publication turns them into public goods. Since researchers are typically not paid for their publications but receive a fixed income, this is a case of voluntary provision of public goods.

The VCM can be combined with different kinds of incentives. In science, status among one's peers is an important reward in itself and the key to most other rewards, like attractive positions or Nobel Prizes. The status of a researcher is determined by his impact on the field, that is, by the extent to which his ideas and results are used by other researchers.[6] Since all researchers face the same incentives when selecting the ideas and results on which to build their own work, everybody must try to produce inputs for the next step in research, in a potentially unending chain. From the internal perspective, all outputs are either discarded or used as intermediate inputs in further research.

There are, of course, users of research outside science: innovators, administrators, politicians, doctors, lawyers, and writers of popular science books. However, their preferences do not influence competition within science.[7] Indeed, this is the whole point of scientific competition. It is difficult, if not impossible, for outsiders to evaluate the quality of research, especially basic research. Scientific competition delegates the evaluation to those who have the required competence: the researchers themselves.[8]

Once scientific achievements have become established, even outsiders may be able to judge their value, especially if these achievements are closely related to technological innovations. The problem is to evaluate cutting-edge research, especially basic research. Usually, scientific developments are way ahead of outsiders' ability to understand them. It is this time lag that makes it impossible for outsiders to steer science in their preferred direction.[9]

Science is therefore characterized by a high degree of autonomy – by producer sovereignty, not consumer sovereignty.[10] Nevertheless, producer sovereignty in science works surprisingly well even for the outside users of research, at least on average. The reason is that competing researchers often coordinate on reasonable methodological standards and maintain them collectively. This brings us back to the central question: Why and how do they do it?

As we have seen, methodological standards are neither imported into science nor imposed from the outside. It might be argued, then, that they are maintained by implicit cooperation between researchers within scientific competition. A possible mechanism for cooperation is provided by the folk theorem(s) for infinitely repeated games (see, e.g., Fudenberg and Tirole 1991, 150–65). Since the status of a researcher is determined by the extent to which his ideas and results are used by other researchers, researchers could collectively enforce some standard by never making use of the work of a colleague who violated the standard.

Such a punishment strategy requires punishment for all those who fail to punish, which makes correct behavior difficult to monitor.[11] More importantly, however, punishment strategies just seem to be irrelevant in science. They would require researchers to ignore high-quality research outputs because somebody in the chain leading to this high-quality output used low-quality inputs. This is not the behavior we observe. Researchers are usually eager to build upon high-quality research. A good paper is not discarded just because the author misbehaved in some way. It is the quality of the ideas and results under consideration, and not the history leading to them, that determines whether researchers use them or not.

In this connection, it is interesting to note that the scientific community rarely punishes those who violate certain other well-established norms. For instance, plagiarism certainly violates the norms of proper scientific behavior. Nevertheless, it is usually not severely punished by the scientific community.[12] Of course, the folk theorem can easily accommodate any kind of behavior: in theory, there exist equilibria where one kind of norm violation would be punished and another would not. However, without further evidence in its favor, such an explanation is implausible. In my view, nothing points to the use of these punishment strategies in science. Norms of punishment are not part of the official methodology in any field, as far as I know, and they are not taught – quite in contrast to norms of proper citation and other kinds of appropriate behavior.

Subsequently, I explore a different mechanism that can explain coordination on common methodological standards. This mechanism is based on a production function for research with a simple but intuitively plausible property: input quality affects, at least stochastically, output quality. In other words, quality is hereditary in scientific production. If the costs of selecting the right kind of quality are not too high, hereditariness implies that the expectation that everybody prefers inputs of a certain quality is self-enforcing. Since everybody produces inputs for other researchers, it is in everybody's own interest to build upon research of the required quality, in order to maximize the chance to produce output of this quality. Under these assumptions, methodologies are self-enforcing or incentive compatible: if everybody adheres to the methodological rules, it is in everybody's interest to do so.

Hereditary quality features in most production processes. From food to clothing to housing, input quality determines output quality. For many quality characteristics in science, hereditariness is at least plausible: Using simple models as part of a new model increases the chances to produce a simple model. Combining consistent models increases the chances to come up with a new consistent model. Incorporating relations that have survived statistical tests increases the chances of finding a statistical model that will also stand up to empirical scrutiny.

A link between input and output quality provides a rationale for researchers to be selective in the choice of inputs. Researchers use the work of others if they think that it is good enough to build upon it (Hull 1988). Hence, if a paper gets a negative evaluation, it will not be used by other researchers, which means that the author fails in the quest for status. Thus, given hereditary quality, high-quality papers are used and low-quality papers are discarded in equilibrium just because everybody expects this kind of behavior.[13]

The paper proceeds as follows. Section 2 presents, in informal terms, a simple model showing how quality norms can become established in scientific

https://doi.org/10.3366/epi.2011.0014 Published online by Cambridge University Press

competition if these norms pick out hereditary properties. Section 3 is concerned with the interpretation of the model and indicates possible extensions. Section 4 concludes. The formal presentation and analysis of section 2's model is relegated to the appendix.

2. THE GAME OF SCIENCE

Let me sketch the mechanism of hereditary quality in the simplest terms.[14] Consider an infinite sequence of researchers, each of whom publishes one paper that may be used as input by the next researcher in the sequence. Alternatively, a researcher can discard his predecessor's paper and start from scratch.

We assume that researchers are all alike in their preferences and abilities. It may be the case that not everybody is able to become a researcher; but those who do enter research have identical abilities and preferences.

All other things being the same, each researcher is better off if his successor makes use of his paper. This reflects, in the simplest way possible, the idea that researchers gain status, and other rewards tied to status, if they have an impact on their field. However, research also has costs, which also enter the utility function and depend on what the researcher tries to achieve and whether he uses his predecessor's paper.

We now introduce different qualities. Let a paper fall into one of two categories, X or Y. For instance, the ideas in a paper may be either simple or complicated, or either consistent or inconsistent. Writing X-papers is the more demanding task. Researchers can tell an X-paper from a Y-paper; however, neither the quality of their own paper nor the quality of any other paper enters their utility function.

The costs of research have several components. There is a basic cost that is independent from the researcher's strategy. If a researcher discards the paper of his predecessor and starts from scratch, he bears an additional cost. And if a researcher tries to write an X-paper, which is more demanding, he also bears an additional cost.

The core assumption of the model is a production function, identical for all researchers, that introduces hereditariness of quality. If a researcher wants to write a Y-paper, success is certain. If he tries to write an X-paper, success is uncertain. The probability of producing an X-paper is high if a researcher makes use of an X-paper, low if he makes use of a Y-paper, and lies in between if he discards his predecessor's paper and starts from scratch.

We consider a sequential game where each researcher produces a paper that is used or not by the next researcher, who may or may not try to write an X-paper. Such a game has many equilibria. However, with the help of some plausible requirements, we can restrict considerations to two equilibria.

First, we consider only equilibria where all researchers use forward-looking strategies. A researcher who uses a forward-looking strategy may choose different actions depending on the quality of his predecessor's paper but, given this quality, will ignore all past events. The restriction to forward-looking strategies rules out the implausible punishment strategies discussed in the introduction.

Second, we consider only stationary equilibria, that is, equilibria where all researchers – with the exception of the first, who lacks a predecessor whose paper he could use – use the same strategy. A stationary equilibrium describes a situation where quality standards are established, nobody expects methodological change, and all researchers comply with the accepted standards. Non-stationary equilibria would describe situations where methodological change unfolds according to everybody's expectations. This seems implausible. Of course, methodological rules change over time. However, these changes seem to be unexpected. They should accordingly be modeled as an unexpected change from one equilibrium to another.

Two further requirements, subgame perfection and strictness, are rather technical and therefore explained in appendix B.

Thus, we focus on subgame perfect, stationary, and strict equilibria in forward-looking strategies, subsequently called "simple equilibria" for short.

There are just two simple equilibria: a non-discriminating equilibrium where all researchers (except the first) use their predecessor's paper and write Y-papers, and a discriminating equilibrium where only X-papers are used and every researcher tries to write an X-paper.

In both equilibria, researchers do, in their own best interest, what they expect their successors to do. In the non-discriminating equilibrium, each researcher expects his successor to use his paper no matter whether it is an X- or a Y-paper. Therefore, it makes no sense to do anything else but to produce a paper as cheaply as possible, which means using his predecessor's paper and writing a Y-paper. Obviously, the non-discriminating equilibrium always exists under our assumptions.

In the discriminating equilibrium, each researcher knows that his successor will only use X-papers and therefore tries to write such a paper. He will even discard his predecessor's paper if it is a Y-paper and start from scratch in order to maximize his own chances to produce an X-paper. However, a discriminating equilibrium may not exist if the costs of pursuing this strategy are too high in view of the expected payoffs.

The expected utility of a researcher in the non-discriminating equilibrium is, of course, maximal since he bears only the basic costs of research and his paper is always used. The expected utility of a researcher in the discriminating equilibrium, however, must be lower for two reasons: he bears additional costs, and he cannot be sure that his paper will be used.

3. METHODOLOGY IN EQUILIBRIUM

The model of scientific competition sketched above shows how endogenous quality standards can establish themselves in scientific competition without invoking game theory's folk theorem(s).

Let us call papers satisfying a demanding standard (the X-papers of the model) high-quality papers, in contrast to the other papers (the Y-papers of the model), called low-quality papers. In the discriminating equilibrium, researchers' strategies are described by the following methodological rule: Consider only high-quality papers as an input for your own research, and try to produce the same high quality. The quality standard referred to in this rule is demanding, but the rule is incentive compatible because quality is hereditary. The expectation that other researchers will stick to the rule becomes a self-fulfilling prophecy.

3.1 Equilibrium Analysis

It may seem odd to analyze scientific research under an equilibrium perspective.[15] After all, equilibrium requires rational expectations, which means that there are no real surprises. The development of research, in contrast, is nothing if not surprising.

However, there is no real contradiction here. Surprising scientific developments can easily occur in an environment where all researchers perfectly know the methodological standards that are used to evaluate scientific innovations. The model of section 2 considers methodology in equilibrium, not research in equilibrium. While there are methodological debates and innovations in methodological standards, it seems that we have on the methodological level more stability than on the level of scientific results. Thus, an equilibrium analysis may be appropriate as an approximation.

There is a further aspect involved. An important driving force in the acceptance of methodological standards is the fact that researchers are forward-looking: they think about the question of how to induce other researchers to build upon their work. At least as a first approximation, this aspect is best captured in a simple game-theoretic model with rational forward-looking agents. The obvious alternatives, simple evolutionary models of behavior, lack forward-looking agents.[16]

3.2 Non-Hereditary Quality

The model implies that not any kind of quality standard can be established in this way. Researchers must believe that quality is hereditary in the production process, that is, they must believe that using high-quality input increases their chances to produce high-quality output.

Once some quality standards have become established in this way, others may follow for different reasons. Procedural standards, like rules of proper experimentation, are a case in point. The procedural aspect of quality is not hereditary, and it can be achieved with certainty. However, following certain procedures is a means to an end. Before some procedures can become important, researchers must have accepted some quality standard according to which the respective end is important. The relevant ends are the subject matter of evaluative quality standards, which are based on certain "regulative ideas" (H. Albert 1978, 31) like truth, reliability, theoretical unity, consistency, mathematical correctness, simplicity, explanatory power, fruitfulness, and so on.

The model of this paper is only concerned with evaluative quality standards. These standards can become established if they pick out hereditary properties. Procedural standards become established because researchers believe that the recommended procedures are helpful in, or necessary for, achieving or demonstrating high quality in the light of evaluative standards. Thus, procedural standards ride piggyback on evaluative standards.

Further quality standards in science concern scientific writing and citation. Experienced researchers know how to write a paper and who should be cited for what. Beginners, or even experienced researchers who cross disciplinary boundaries, cannot easily comply with these standards. This explains why these formal standards exist and are important. Formal quality provides a reliable signal of experience, and experience, in turn, is highly correlated with the ability to produce high-quality ideas and results. In a situation where researchers differ in experience, many papers compete for the attention of researchers, and where digesting the scientific content of a paper costs valuable time, the optimal strategy for deciding what to read is to scan papers quickly and, in a first step, to discard all papers with the wrong references or a badly written abstract or introduction.

Quite plausibly, then, signaling becomes important once we extend the model and account for the fact that researchers must search for relevant papers of a sufficient quality. Again, the signaling explanation for standards of scientific writing and citation rests on the assumption that it is important for researchers to select high-quality inputs.

Standards that emerge only for the purpose of signaling may of course be detrimental in the light of other standards. For instance, the use of advanced mathematics in an empirical science like economics is an excellent signaling device. Beginners and outsiders have great difficulties in mimicking experienced researchers in mathematical economics. Once mathematical sophistication has been established as a signal, however, it might crowd out other qualities that contribute more to empirical relevance. After all, investing into mathematical sophistication necessarily comes at the cost of investing into other skills. Moreover, the immigration of mathematicians and mathematically inclined researchers into a field can shift expectations and establish an equilibrium where everybody expects everybody else to use the evaluation standards of mathematics.

3.3 Good and Bad Equilibria

In scientific competition, production decisions are not governed by the evaluations of outside users of science. From the perspective of researchers, their output serves only as an input for further research; there is no final output. Outsiders may profit from research, or they may be hurt by it, but these effects are external to scientific competition. The value of research outputs in the eyes of the producers is determined by the demand of an infinite sequence of downstream producers who are all in the same position. Hence, the results of research have no "fundamental value" within science. Because quality is hereditary, every researcher tries to guess what the next researcher would like to use, and selects his own inputs accordingly. Since there are no contracts fixing a price for papers satisfying some quality norm, production proceeds on the basis of these guesses.

This game has the structure of a sequential guessing game or "beauty contest".[17] There is a certain similarity to stock markets, where a similar beauty contest also leads to the coordination on evaluation standards (Pratten 1993). However, the need to coordinate on evaluation standards is more obvious for stock markets, which use the price mechanism. Science, in contrast, uses the voluntary contribution mechanism. Every researcher produces a unique public good. Hence, there is no need to find a common price. Without the assumption of hereditary quality, which forges a link between input quality and output quality, there would be no beauty contest.

The equilibrium in scientific competition is not necessarily *ex ante* efficient from the researchers' point of view. The discriminating equilibrium is more competitive than the non-discriminating equilibrium: researchers discard some inputs and bear higher costs, which lowers expected utility. A similar effect exists in some other beauty contests, where coordination on an equilibrium reduces average payoffs. Thus, as in market competition, the competitive mechanism diminishes producers' rents. Whether this is good or bad for society as a whole depends, of course, on the quality standard itself. There is no guarantee that the quality standard established in science reflects the preferences of outsiders.

While the non-discriminating equilibrium of the model may describe the state in a small field, it cannot be taken seriously as a description of science as a whole. The utility a researcher derives from the use of his contribution by other researchers depends on the scarce rewards that are connected with use. Without discrimination, there would be no basis for distributing rewards on the basis of researchers' contributions, and some other distribution scheme that is independent from the quality of the contributions would have to emerge.[18]

3.4 Preferences for High Quality

The present model assumes that researchers are indifferent as to the quality of the papers they use and write. Even if the standards of science are unlikely to be imported into science by scientific novices, it might be that the preferences of novices favor some standards over others.

If one assumes that all researchers are only motivated by suitable ideals, it may become rational for them to adopt certain quality standards independently of the question whether the qualities in question are hereditary. However, it seems rather unlikely that researchers' motivations are that simple.

Let us, therefore, assume that many researchers have a preference for using and producing X-papers but also try to be successful in terms of the model above. This would make it more likely that a discriminating equilibrium establishes itself, for two reasons. First, introducing such a preference would be equivalent to reducing the costs of discarding low-quality papers and of trying to produce high-quality papers, thus making it more likely that the discriminating equilibrium exists. Second, a known preference of all or many researchers for high quality may be sufficient to establish the self-fulfilling expectation that the discriminating equilibrium will obtain. In other words, the discriminating equilibrium may become a focal point in Schelling's sense (1960).

Nevertheless, a non-discriminating equilibrium remains a possibility. Even if all researchers have a preference for demanding standards, they may find themselves locked into an equilibrium where nobody observes these standards because everybody expects everybody else to ignore them.

3.5 Methodological Change

The present model explains a methodological equilibrium. It does not offer an explanation of how such an equilibrium emerges. However, it is not inconsistent with methodological change, not even with methodological revolutions.

In terms of the model, change requires a shift in expectations. This can be brought about by methodological argument. Assume that the distinction between Y-papers and X-papers is not yet known. A researcher discovers the distinction, manages to write an X-paper and, in the introduction, explains the difference between X- and Y-papers and argues that everybody should use and write X-papers. This could change the expectations of his successors. If they believe the methodological argument to be effective, it is effective, and a new discriminating equilibrium emerges.

Of course, this story is much too simple. A new argument will typically change the expectations of some but not all researchers. There are many researchers working at the same time. Some of them may believe that the argument will be successful eventually; they may therefore take the risk and adopt the standards. This may convince others that, at least in the long run, their chances to have an impact on their field are greater if they join the movement.

Thus, the bandwagon may slowly begin to move as more and more researchers jump on it. Jumping early is risky because too few researchers may join and the movement may fail. On the other hand, if the movement is successful, early adopters on average will have a greater impact on the field than late adopters because many late adopters build upon the work of a few early adopters.

4. CONCLUSION

The present paper makes two main points. First, a methodology must be incentive compatible in order to establish itself in scientific competition. Second, the most plausible mechanism ensuring incentive compatibility is based on properties that are hereditary in the production of research.

Which hereditary properties, then, are picked out by actual scientific methodologies? I have already mentioned some, like simplicity and consistency. The most important such property, however, is confirmation by empirical tests, or corroboration, to use a term introduced by Karl Popper in order to emphasize the requirement that empirical tests should be as severe as possible.

It is well known that hypotheses cannot be tested in isolation. The derivation of a testable prediction will, implicitly or explicitly, always make use of further hypotheses, so-called auxiliaries. This leads to the famous Duhem problem: it is always possible to blame the auxiliaries if a hypothesis fails a test.

The standard solution to the Duhem problem is to require that the auxiliaries are themselves well-corroborated or, in other words, that they are part of the background knowledge. Thus, in order to establish a newly invented hypothesis as corroborated, a researcher must combine it with well-corroborated hypotheses in order to derive an empirical prediction, which then has to be tested.

The new hypothesis may survive or fail the test. However, the requirement to use other corroborated hypotheses as an input when one wants to have a chance to produce a new corroborated hypothesis means that corroboration is hereditary in the sense of section 2's model.

The model of this paper, then, shows how the division of labor works in the case of the pursuit of knowledge. Researchers have an incentive to build upon the empirically successful work of other researchers by proposing new hypotheses that fit with the best older ideas in generating new testable predictions. This mechanism does not even require that those involved in scientific competition pursue knowledge for its own sake. However, if scientific competition works as it should, the pursuit of knowledge is not in conflict with the pursuit of material interests.

APPENDIX: MODELING THE GAME OF SCIENCE

A. ASSUMPTIONS

Consider an infinite sequence of researchers $t = 0, 1, \ldots, \infty$, each publishing one paper, which may either be an X-paper or a Y-paper. All researchers have identical utility functions $u(v, c) = v - c$, where v is a researcher's payoff from publishing his paper (already net of the basic costs of research) and c are additional costs that depend on the researcher's decisions.

The behavior of researcher 0 is irrelevant for subsequent considerations. We therefore ignore him.

Each researcher $t > 0$ faces two decisions: whether to use the paper of his predecessor $t - 1$ (decision variable P), and whether to try to produce an X-paper (decision variable Q, for "quality").

Researchers who try to produce an X-paper (decision $Q = X$) succeed only with a certain probability w. If they fail, or if they do not even try (decision $Q = Y$), their output is a Y-paper. Trying to produce an X-paper increases the costs of research by $c_X > 0$.

Researchers $t = 1, 2, \ldots$ who discard their predecessor's paper (decision $P = D$) bear extra research costs $c_D > 0$ and have a success probability $w = q > 0$ of producing an X-paper if they try.

For researchers $t = 1, 2, \ldots$ who use their predecessor's paper as an input for their own research (decision $P = U$), the probability w of success for the production of X-papers depends on the quality of this input. If they use an X-paper, $w = p$; if they use a Y-paper, $w = r$. We assume $1 > p > q > r \geq 0$, which means that quality is hereditary.

If a researcher's own paper is discarded by his successor, he receives a payoff v_0; if his paper is used, he gets $v_1 > v_0$.

Researchers $t = 1, 2, \ldots$ have four possible actions, resulting from the four possible combinations of decisions P and Q, namely, UX, DX, UY, DY, where, for instance, UX means "use your predecessor's paper and try to write an X-paper".

If they try to write an X-paper, their success probability w is determined by the input they use. If they use a Y-paper, their success probability is r. If they start from scratch, their success probability is q. If they use an X-paper, their success probability is p. We assume $0 < r < q < p < 1$.

Consider a researcher $t > 0$. Let x be the probability that his paper is used by his successor if it is an X-paper, and let y be the probability that his paper is used if it is a Y-paper. The expected payoff of a researcher who produces an X-paper with probability w is $\pi(w) - c$, where c are the researcher's costs and $\pi(w) := [wx + (1 - w)y]v_1 + [w(1 - x) + (1 - w)(1 - y)]v_0$ with $\pi(w) - \pi(w') = (w - w')(x - y)(v_1 - v_0)$. We then find the following expected payoffs $EU(PQ|Z)$ for the relevant decisions, where $Z = X, Y$ denotes the category of the predecessor's paper and PQ stands for the possible actions as before:

$$
\begin{aligned}
EU(UX|X) &= \pi(p) - c_X \\
EU(UX|Y) &= \pi(r) - c_X \\
EU(UY|X) &= \pi(0) \\
EU(UY|Y) &= \pi(0) \\
EU(DX|X) &= \pi(q) - c_D - c_X \\
EU(DX|Y) &= \pi(q) - c_D - c_X \\
EU(DY|X) &= \pi(0) - c_D \\
EU(DY|Y) &= \pi(0) - c_D.
\end{aligned}
\tag{1}
$$

We assume perfect information. All researchers know all the other researchers' costs and payoffs, and researcher $t > 0$ knows what all researchers $s < t$ did and

https://doi.org/10.3366/epi.2011.0014 Published online by Cambridge University Press

whether they produced an X-paper or a Y-paper. The game, then, is a sequential game.

B. SIMPLE EQUILIBRIA

Simple equilibria, as defined in the text, are equilibria in forward-looking strategies that are subgame perfect, stationary, and strict.

A researcher's pure strategy is a complete contingent plan selecting one of his four possible actions for each possible information set (that is, combination of known past events). As explained in the text, we consider only forward-looking strategies, which, given the predecessor's paper, ignore all other information about the past. A researcher $t > 0$ has four possible actions. His forward-looking pure strategies select one action for the case that his predecessor produced an X-paper, and one action for the case that his predecessor produced a Y-paper, yielding sixteen different forward-looking pure strategies.

The mixed strategies of a researcher result from randomization among the pure strategies. Below, it is argued that we can focus on pure strategies.

An equilibrium (or Nash equilibrium) is a list of strategies, one for each researcher, with the property that every researcher's strategy maximizes his expected utility given the strategy choices of the others. As explained in the text, we require equilibria to be stationary equilibria, that is, we require that all researchers $t = 1, 2, \ldots$ play the same strategy.

We consider only strict equilibria, that is, equilibria where unilateral deviation of a researcher from his equilibrium strategy leads to an expected loss for this researcher. Strict equilibria are much more plausible than non-strict equilibria, where players are indifferent between playing their equilibrium strategy and deviating from it. Strictness also means that we can focus on pure-strategy equilibria because other equilibria are never strict.

Under the assumption of perfect information, equilibria in forward-looking strategies are also subgame perfect, meaning that they are not based on incredible threats.

C. FINDING THE SIMPLE EQUILIBRIA

The sixteen forward-looking pure strategies of researchers $t = 1, 2, \ldots$ are denoted by a combination of four letters $P_X Q_X P_Y Q_Y$. The pair $P_Z Q_Z$, where $P_Z \in U, D$ and $Q_Z \in X, Y$, denotes the researcher's action in the face of a Z-paper, $Z = X, Y$. Thus, $UXDY$ means "use an X-paper and try to produce one; discard a Y-paper and produce a Y-paper".

Consider the expected payoffs (1). In the case of pure strategy play, the probabilities x and y are, of course, 0 or 1. Obviously, $EU(UY|Z) - EU(DY|Z) = c_D > 0$, that is, it is always better to use the predecessor's paper if one writes a Y-paper. Moreover, $EU(UX|X) - EU(DX|X) = c_D > 0$, that is, if a researcher finds an X-paper and tries to write an X-paper, he is always better off

https://doi.org/10.3366/epi.2011.0014 Published online by Cambridge University Press

using the paper. By excluding the strictly dominated options, we get the following reduced set of relevant options and expected payoffs:

$$
\begin{aligned}
EU(UX|X) &= \pi(p) - c_X \\
EU(UX|Y) &= \pi(r) - c_X \\
EU(UY|X) &= \pi(0) \\
EU(UY|Y) &= \pi(0) \\
EU(DX|Y) &= \pi(q) - c_D - c_X.
\end{aligned}
\tag{2}
$$

This means that X-papers will always be used ($x = 1$). Since we focus on stationary pure-strategy equilibria, there remain two possibilities: an equilibrium where all papers are always used ($x = y = 1$), and equilibria where X-papers are always used and Y-papers are always discarded ($x = 1$, $y = 0$).

If $x = y = 1$, we have $\pi(w) = v_1$, which implies $EU(UY|X) > EU(UX|X)$ and $EU(UY|Y) > EU(UX|Y) > EU(DX|Y)$. Thus, utility maximization requires that a researcher uses any paper and writes only Y-papers. A researcher facing $x = y = 1$, then, will play the strategy *UYUY*, which implies that his predecessor faces $x = y = 1$ and will also play *UYUY*. Therefore, all researchers playing *UYUY* is a simple equilibrium.

If $x = 1$ and $y = 0$, we have $\pi(w) = wv_1 + (1 - w)v_0$ with $\pi(w) - \pi(w') = (w - w')(v_1 - v_0)$. In order to ensure that researchers expecting $x = 1$ and $y = 0$ will themselves strictly prefer to use X-papers and to discard Y-papers, further conditions must be fulfilled.

Discarding Y-papers is strictly preferred if and only if $EU(DX|Y) > EU(UX|Y)$ and $EU(DX|Y) > EU(UY|Y)$, which leads to the conditions $\pi(q) - c_D - c_X > \pi(r) - c_X$ and $\pi(q) - c_D - c_X > \pi(0)$. These conditions can be transformed into

$$
\begin{aligned}
s_1 &< 1, \quad \text{where} \quad s_1 := \frac{c_D}{(q - r)(v_1 - v_0)} \\
s_2 &< 1, \quad \text{where} \quad s_2 := \frac{c_D + c_X}{q(v_1 - v_0)}.
\end{aligned}
\tag{3}
$$

Of course, given that it is optimal, when finding a Y-paper, to discard it and to try to write an X-paper, it cannot be optimal to write a Y-paper when finding an X-paper. Not surprisingly, then, $s_2 < 1$ implies $EU(UX|X) > EU(UY|X)$.

Thus, under all parameter constellations, there exists a simple equilibrium where all researchers $t = 1, 2, \ldots$ play *UYUY*. If the existence conditions (3) are fulfilled, there exists a further simple equilibrium where all researchers $t = 1, 2, \ldots$ play *UXDX*. In the latter equilibrium, all researchers use only X-papers and try to write X-papers because they all foresee that Y-papers will always be discarded. We call this the discriminating equilibrium. In the other equilibrium, everybody is prepared to use both kinds of papers and nobody tries to write an X-paper; therefore, only Y-papers are used and written. We call this the non-discriminating equilibrium.

D. EXPECTED UTILITY IN EQUILIBRIUM

In the non-discriminating equilibrium, all papers are used and only Y-papers are written. Therefore, the expected utility of each researcher $t > 0$ is $EU_N = v_1$.

The discriminating equilibrium can be analyzed as a regular Markov chain with two states: in state 1, an X-paper has been produced; in state 2, a Y-paper has been produced. The transition matrix between states is $T = \begin{pmatrix} p & 1-p \\ q & 1-q \end{pmatrix}$.

The stationary distribution, to which this chain converges independently from its starting point, is determined by the equation $(\chi, 1 - \chi)T = (\chi, 1 - \chi)$. The probability χ is the long-run probability of state 1. Solving for χ, we find $\chi = \frac{q}{q+1-p}$, where $\chi \in (q, 1)$ since $1 > p > q > 0$.

Thus, independently from the starting point of the process, the probability that a researcher can use an X-paper converges, in the discriminating equilibrium, to χ as t goes to infinity. The long-run expected utility of a researcher in the discriminating equilibrium is, therefore, $EU_D = \chi v_1 + (1 - \chi)v_0 - c_X - (1 - \chi)c_D$.

The expected welfare loss for researchers through the quality standard, then, is $EU_N - EU_D = c_X + (1 - \chi)(v_1 - v_0 + c_D)$. This is easily interpreted: c_X is the cost of trying to produce an X-paper; $v_1 - v_0 + c_D$ is the cost that results from discarding a Y-paper, namely, $v_1 - v_0$ for the discarded paper's author and c_D for his successor, who has to start from scratch.

REFERENCES

Albert, Hans. 1978. *Traktat über rationale Praxis*. Tübingen: Mohr Siebeck.

Albert, Hans. 2006. "Die ökonomische Tradition und die Verfassung der Wissenschaft." *Perspektiven der Wirtschaftspolitik* 7: 113–31.

Albert, Hans. 2010. "The Economic Tradition and the Constitution of Science." *Public Choice* 144: 401–11.

Albert, Max. 2006. "Product Quality in Scientific Competition." *Papers on Strategic Interaction* 6–2006. Jena: Max Planck Institute of Economics.

Albert, Max. 2008a. "Introduction." In M. Albert, D. Schmidtchen, and S. Voigt (eds.), *Scientific Competition. Conferences on New Political Economy* 25, pp. 1–9. Tübingen: Mohr Siebeck.

Albert, Max. 2008b. "Product Quality in a Simple OLG Model of Scientific Competition." *MAGKS Joint Discussion Paper Series in Economics* 04–2008. Giessen: Justus Liebig University Giessen. http://www.uni-marburg.de/fb02/makro/forschung/magkspapers/04–2008_albert.pdf

Albert, Max and Jürgen Meckl. 2008. "What Should We Expect from Peer Review?" In M. Albert, D. Schmidtchen, and S. Voigt (eds.), *Scientific Competition. Conferences on New Political Economy* 25, pp. 257–61. Tübingen: Mohr Siebeck.

https://doi.org/10.3366/epi.2011.0014 Published online by Cambridge University Press

Albert, Max, Dieter Schmidtchen, and Stephan Voigt (eds). 2008. *Scientific Competition. Conferences on New Political Economy* 25. Tübingen: Mohr Siebeck.

Broad, William and Nicholas Wade. 1982. *Betrayers of Truth: Fraud and Deceit in the Halls of Science*. New York: Simon and Schuster.

Camerer, Colin F. 2003. *Behavioral Game Theory*. New York: Russell Sage Foundation; Princeton: Princeton University Press.

Congleton, Roger D. 1989. "Efficient Status Seeking: Externalities, and the Evolution of Status Games." *Journal of Economic Behavior and Organization* 11: 175–90.

Dasgupta, Partha and Paul A. David. 1994. "Toward a New Economics of Science." *Research Policy* 23: 487–521.

David, Paul A. 1998. "Common Agency Contracting and the Emergence of 'Open Science' Institutions." *American Economic Review* 88: 15–21.

David, Paul A. 2004. "Understanding the Emergence of 'Open Science' Institutions: Functionalist Economics in Historical Context." *Industrial and Corporate Change* 13: 571–89.

Diamond, Arthur M., Jr. 2008. "Economics of Science." In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, 2nd ed. London: Palgrave Macmillan.

Fudenberg, Daniel and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.

Gambetta, Diego. 2009. *Codes of the Underworld: How Criminals Communicate*. Princeton: Princeton University Press.

Güth, Werner and Menahem E. Yaari. 1992. "Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach." In U. Witt (ed.), *Explaining Process and Change*, pp. 23–34. Ann Arbor, MI: University of Michigan Press.

Hackl, Franz, Martin Halla, and Gerald J. Pruckner. 2007. "Volunteering and Income – The Fallacy of the Good Samaritan?" *Kyklos* 60: 77–104. Longer version: *Working Paper* 415. Linz: Department of Economics, Johannes Kepler Universität, 2005.

House, Christopher L. and Emre Ozdenoren. 2008. "Durable Goods and Conformity." *RAND Journal of Economics* 39: 452–68.

Hull, David L. 1988. *Science as a Process*. Chicago: University of Chicago Press.

Marmot, M. 2004. *The Status Syndrome: How Social Standing Affects Our Health and Longevity*. New York: Holt and Company.

Mayer, Thomas. 1993. *Truth versus Precision in Economics*. Aldershot: Edward Elgar.

Merton, Robert K. 1973. *The Sociology of Science*. Chicago: University of Chicago Press.

Pratten, Cliff. 1993. *The Stock Market*. Cambridge: Cambridge University Press.

Schelling, Thomas. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Seidl, Christian, Ulrich Schmidt, and Peter Grösche. 2008. "A Beauty Contest of Referee Processes of Economics Journals." In M. Albert, D. Schmidtchen, and S. Voigt (eds.), *Scientific Competition. Conferences on New Political Economy* 25, pp. 235–55. Tübingen: Mohr Siebeck.

Sokal, Alan D. 1998. "What the *Social Text* Affair Does and Does Not Prove." In N. Koerttge (ed.), *A House Built on Sand: Exposing Postmodernist Myths about Science*, pp. 9–22. Oxford: Oxford University Press.

Stephan, Paula E. 1996. "The Economics of Science." *Journal of Economic Literature* 34: 1199–235.

Stephan, Paula E. Forthcoming. "The Economics of Science." In B. H. Hall and N. Rosenberg (eds.), *Handbook of Economics of Technical Change*.

Vanberg, Viktor J. 2010. "The 'Science-as-a-Market' Analogy: A Constitutional Economics Perspective." *Constitutional Political Economy* 21: 28–49.

Ziman, John. 2000. *Real Science: What It Is, and What It Means*. Cambridge: Cambridge University Press.

NOTES

* Paper based on a talk at the conference "Collective Knowledge and Epistemic Trust", Greifswald, 6–8 May 2010. For suggestions and comments, I am indebted to Volker Gadenne, Thomas Grundmann, and an anonymous referee.

1 The term "science" is used in the European sense here and therefore includes the social sciences and parts of the humanities, even though the observation about the high average quality of scientific research may not apply with the same force in all these fields.

2 See (Stephan 1996, forthcoming) and (Diamond 2008) on the economic analysis of science and, specifically, on the characteristics of open science. See (Albert 2008a) for a compact summary and statement of the position underlying the present paper.

3 Of course, even demanding standards may be low standards from a different perspective. Parts of the humanities are characterized by very low standards if judged from the perspective of the natural sciences; see, for instance, (Sokal 1998). However, here, as in any other competitive environment, it cannot be easy for the average participant to rise to the top. Sokal does not argue that internal quality standards do not exist in the fields under attack or that quality perceptions do not govern the distribution of rewards. Thus, even if his criticism is well-founded (which I believe to be the case), standards might be quite demanding and, consequently, quality may be high from the perspective of an insider who accepts the quality standards.

4 It is certainly possible that the distribution of rewards proceeds on some other basis. Gambetta (2009, 42–5) argues that the distribution of academic positions in Italy is dominated by those, and favors those, who produce low-quality output. Interestingly, quality standards as such seem to be the same as elsewhere. Since well-paid jobs are probably as hard to get as everywhere else, and producing no output or only low-quality output is easy, selection must be based on other criteria, like a candidate's ability and willingness to be of service to the power brokers. In Gambetta's analysis, producing low-quality output enhances trustworthiness because one becomes unable to opt out of the system (by, for example, going abroad).

5 See (Albert 2008a). Cf. (Vanberg 2010) for a comprehensive discussion of the science-as-a-market analogy from a constitutional-choice perspective.

6 On incentives for voluntary contributions, see (Hackl et al. 2005, 2007). On status as a reward in itself, see (Marmot 2004). See (Merton 1973) on status as an incentive in science and (Hull 1988, 283 and elsewhere) on use as the basis of status. See (Congleton 1989) for a model of status seeking and competing VCMs. See (H. Albert 2006, 2010) for the relation between economics and philosophy of science.

7 Researchers might, of course, wish to produce outputs that are of immediate use outside science. However, if their outputs are not also inputs for further research,

they are at a dead end within open science. They might turn to proprietary science and the market instead (and some science policies aim at facilitating this step; cf. Ziman 2000). Those who have this option therefore face mixed incentives; they are at the border between open and proprietary science. However, this does not imply that market incentives spill over into open science, because those who are one step removed from proprietary science receive no share of the profit if their research results are used to produce marketable research. Those one step removed from marketable research must still hope that they have an impact on academic research. Thus, even if quality standards in open science tended to favor results that are useful outside science, we would have to explain how these quality standards can become established within science.

8 Cf. the theoretical and historical explanation of the self-regulating character of open science offered by (Dasgupta and David 1994) and (David 1998, 2004), which centers on the inability of the patrons of science to evaluate its results. This explanation must be supplemented, however, by a model of scientific competition explaining how self-regulation can actually deliver high quality. This explanation problem becomes even more severe once it is recognized that scientific quality standards are endogenous to science.

9 In the long-run, of course, the patrons of science, be they private or public, can allocate resources to fields that, in the past, delivered results that were valuable to them (cf. also Vanberg 2010, 42–3). Given the above-mentioned time lag and the self-interest of politicians, it seems highly doubtful to me that a political process of resource allocation in science will deliver better results from a social point of view than decisions made by self-interested researchers. However, in order to say more on this problem, we would need models of different resource allocation mechanisms in science.

10 The term producer sovereignty in this context is due to Mayer (1993, 10), who considers it to be a problematic feature, at least in economics (see also Vanberg 2010, 45).

11 In some games, there are simple punishment strategies that require less monitoring: everybody punishes everybody else once a single norm violation occurs. However, such strategies cannot work in the present model because the threat not to use one's predecessor's paper if some norm violation occurred in the past is not credible: given that his work is not to be used anyway, a researcher would not hesitate to use the work of others because this lowers his own costs of research and more severe punishment is not possible anyway.

12 Hull offers an explanation for this fact that is in line with the explanation of methodological standards offered below: plagiarism, in contrast to counterfeit data, does not affect the quality of the research building upon the offending publication. See the critical discussion of (Broad and Wade 1982) by Hull (1988).

13 A similar process can lead to the adoption of common standards in the production of durable consumer goods, especially houses; see (House and Ozdenoren 2008).

14 For full-fledged models along the lines indicated here, see (Albert 2006, 2008b).

15 Cf. (Vanberg 2010, 32) and the literature discussed therein.

16 Of course, more complicated evolutionary models may incorporate forward-looking agents whose preferences and expectations are subject to evolutionary forces ("indirect evolutionary approach", cf. Güth and Yaari 1992).

17 On simultaneous beauty contests, together with experimental results, see, e.g., (Camerer 2003, 209–18).

18 Cf. n. 4 above on Gambetta's (2009) analysis of the Italian academic system.

Max Albert is professor of institutional and behavioral economics at the Justus Liebig University in Giessen, Germany. His areas of research include the philosophy and economics of science. He is also co-editor of the free online journal *Rationality, Markets and Morals*.