WHAT'S WRONG WITH SCIENCE?

David Chavalarias

WHAT'S WRONG WITH SCIENCE ?
Modeling the collective discovery processes with the Nobel Game

David Chavalarias (CAMS/EHESS & ISC-PIF/CNRS)
Pre-print of DOI : 10.1007/s11192-016-2109-9, *Scientometrics* (2016)

Abstract
There is an increasing pressure on scholars to publish to further or sustain a career in academia. Governments and funding agencies are greedy of indicators based on scientific production to measure science output. But what exactly do we know about the relation between publication levels and advances in science ? How do social dynamics and norms interfere with the quality of the scientific production ? Are there different regimes of scientific dynamics ? The present study proposes some concepts to think about scientific dynamics, through the modeling of the relation between science policies and scholars' exploration-exploitation dilemmas. Passing, we analyze in detail the effects of the "publish or perish" policy, that turns out to have no significant effects in the developments of emerging scientific fields, while having detrimental impacts on the quality of the production of mature fields.

Keywords : collective discovery, distributed knowledge, social dynamics, science dynamics, publish or perish, reproducibility, science policy.

MSC Classification : 00A71, 91A06, 91A50, 91A10

JEL Classification : D79, C73

1 Background and Significance

1.1 Science dynamics in the era of big science

In a provocative paper, [Ioannidis(2005)] claims that most research findings are false for most research designs and for most fields. As bold as this statement might seem to be, more and more scientists are subscribing to this point of view. In the field of psychology, a collective of scholars [Open Science Collaboration(2015)] could reproduce only 37 percent of significant results from a sample of 100 papers published in three high-ranking psychology journals in 2008. From machine learning literature, specialists in the field [Pentland(2012)] find it plausible that up to 80 percent of the results could be wrong [1], whereas [Calude and Longo(2016)] have demonstrated that most correlations in Big Data analytics are spurious. In biomedical research, some scholars have estimated that the cost of irreproducibility in biomedical literature could raise up to $28 billion for only the United States [Freedman et al.(2015)Freedman, Cockburn, and Simcoe].

What's happening to Science ?

1. In particular because the researchers didn't understand that they were overfitting the data reported.

Since the work of Karl Popper [Popper(2002)], the relationship between science and truth has changed compared to the traditional Cartesian conception of truth. Most scientists have abandoned the illusion that their theories (described in publications) can be proven true, accepting that theories can only be proven false or be corroborated.

> *The way in which knowledge progresses, and especially our scientific knowledge, is by unjustified (and unjustifiable) anticipations, by guesses, by tentative solutions to our problems, by conjectures. These conjectures are controlled by criticism ; that is, by attempted refutations, which include severely critical tests. They may survive these tests ; but they can never be positively justified : they can neither be established as certainly true nor even as 'probable' (in the sense of the probability calculus). Criticism of our conjectures is of decisive importance : by bringing out our mistakes it makes us understand the difficulties of the problem which we are trying to solve.* [Popper(1962)]

We don't know. We can only conjecture. We should nevertheless ask ourselves what kind of dysfunction is transforming science, as an institution, for having reached the point where "true" results seems to be so scarce.

In this paper, we propose investigating this issue through modeling of scientific dynamics as a collective discovery process [2], which articulates the individual exploration-exploitation dilemma, with science policies and scientific norms. We will use Popper's ideas about conjecture and refutation to present our views, although we believe our approach does not require subscription to a frequentist epistemology.

1.2 It's all a matter of time

Scientific activity is both a private and public venture, one that is articulated by the act of publishing. On the private side, as Popper stated, "a scientist, let him be a theorist or an experimentalist, proposes some statements and tests them step-by-step" [Popper(1962)]. On the public side, a scholar describes his findings, carries out analyses and makes the results available to the public (journals, archives, blogs, conferences, etc.). From that moment, anybody can try to reproduce the proposed results, and possibly prove them false.

Independently of the field under study, this apparently simple process hides a high cognitive complexity. Since we can't prove our theories to be true, the reaching of truth cannot be taken as a landmark to guide our publishing activity.

Will somebody outstrip my work if I don't hurry to publish ? Is my data absolutely flawless ? Are my protocols bias free, my programs bug free, or in other words, are my results *sufficiently* reliable to be published ? The difficulty lies in the definition of "*sufficiently*", which has to take into account our perception of our own achievements, as well as the activity of the scientific community as a whole (what are the hot topics ? How will other scientists welcome my work ? etc.). Anyone familiar with research activity has already faced this kind of dilemma. We have to decide how far we can trust the work of others. Similarly, we have to decide at what point do we stop checking the results of the research of others ?

Underlying all these decisions is the exploration-exploitation dilemma, an ubiquitous dilemma occurring at all levels of behavior and time-scales of the decision making process, from deciding what to do next in the day to planning a career path [Cohen et al.(2007)Cohen, McClure, and Yu]. Exploration is the need to gather information about your environment to make better decisions in the

2. The original idea underlying this modeling framework was proposed by David Chavalarias (1998) in an unpublished study *La thèse de Popper est-elle réfutable ?*

future. Exploitation is the use of already-collected information to make decisions leading to some rewards or benefits.

When it comes to science, different people might have different preferences as to the degree of confidence they have before publishing. Nevertheless, if we operate under the premise that science policies have an effect on scholars' behavior, it should be expected that the ways publications and refutations are encouraged, either by social or by materialistic rewards, will have a great impact on the exploration-exploitation dilemma of scholars, and, *in fine*, on the dynamics of science.

These rewards exist in academia. Whereas social rewards might be subjective and rather difficult to estimate (for example people are heterogeneous with respect to the value they attribute to reputation, glory, and so on), different academic systems provide different levels of incentives for publishing, that are often common knowledge and easy to quantify. It is common for universities and other academic institutions to maintain some minimal requirements in terms of publication rate and journal quality for hiring and confirming full and associate professors and researchers. In some cases, publications are associated with teaching release opportunities for professors ; Christmas bonuses or future funding opportunities for researchers ; reduced scholarship fees for Ph.D. students. Outperforming the publication rate of colleagues appears to be so important in some cases that it has been compared to a life or death issue. The notorious mantra "*Publish or perish*" has even inspired software [Harzing(2007)]. The incentives for publishing are, however, very different from one academic system to another, and from one community to another, with the existence of local "publication cultures."

The case for refutation is similar. The incentives for checking another's results might vary from some recognition and publication payoff (if you can falsify famous results with a publication in a high-ranking journal) to nothing (some of the most important funding agencies, such as the EU ERC, have officially stated that the question of reproducibility is not in the scope of their funding schemes) or even penalties (depending on the stature of those you put into question and the degree of achievement of your own career, the animosity of colleagues could be very harmful). Attitudes seem to be changing, however, with a growing emphasis on reproducibility among the scientific community [3].

When it comes to the cost of being falsified, we see large discrepancies between the treatment received by scholars whose work has called into question. Depending on the degree of severity of the refutation and the degree of misconduct of the author of the publication being falsified, the sanctions range from mere disapproval to being ostracized. The sanctions also depends on applicable standards of the science. One example is the Olivier Voinnet affair. This French scientist, working at ETH Zurich, was an internationally renowned researcher, a member of the French Academy of Sciences, and 2009 winner of the *European Molecular Biology Organisation* gold medal. Investigation of his work discovered that since 1998, he published at least 20 papers that were proven to present false results and contain data manipulation [Larousserie and Morin(2015)]. The sanctions against him were very different between the CNRS (France), where he was tenured and has been disqualified for two years, and ETH Zurich (Switzerland), where he could continue his activity after agreeing to commit to good practices. Although quite rare, such an extreme case is not isolated.

Would this "breach of rules and good practices for the presentation of scientific data in 13 articles" [4] have fallen between the cracks if the many scientists who cited Voinnet's articles had carefully checked the results they relied on ? Worse, an anonymous referee of an article by Voinnet had expressed serious suspicions of fraud in 2003, but this notification had no effect on the publication of the article. This example illustrates how weak the refutation culture among the scientific community

3. See for exemple *Nature*'s Special on Reproducibility : http ://www.nature.com/news/reproducibility-1.17552
4. Sentence reported from the July 10 2015 CNRS national press release.

is. To no one's surprise, the race for publication was put forward to account for Voinnet's behavior.

1.3 Agent-based models of science dynamics

There is an increasing pressure on scholars to publish in order to further or sustain a career in academia. Governments and funding agencies are greedy for indicators based on scientific production to measure science output. But what do we know about the relation between publication levels and advances in science ? Do we understand how social dynamics and norms interfere with the quality of the scientific production ? Are there different regimes of scientific dynamics ?

Modeling scientific dynamics has become an active domain over the last several years [Scharnhorst et al.(2012)S Edmond, B. et al.(2011)Edmond, B., Gilbert, N., Ahrweiler, P., and Scharnhorst, A.], especially the question of the articulation between micro behaviors and collective dynamics, with the development of agent-based models (cf. the review of [Payette(2012)] and section 4). In most papers dealing with agents (the scientists) discovering an epistemic landscape, knowledge is modeled as a set of items for which value can be assessed exactly. The issue is then to find the optimal collective organization and individual behavior for spotting the best places of the epistemic landscape. The exploration-exploitation dilemma that scientists face when deciding to go forward with publishing and its relation with the science policy is thus not fully addressed in these works.

To complement these approaches, we present and explore a model that specifically investigates the mechanism of conjecture and refutation in a social context, and its consequences on scientific dynamics. We first describe a model of the organization of science that puts the social game of science at the forefront. We then investigate the dynamics of such a model with an agent-based approach and sketch the potential impacts of science policies on scientific dynamics.

Science dynamics can be of infinite complexity, and so can their models. We will, as far as possible, keep things simple. Remember, the two fundamental entities of our model are *agents* (the scientists) and *theories* (described in publications). Agents find their motivations in the rewards associated with publications. In the experimental versions of this model [Chavalarias et al.(2006)Chavalarias, Charron, C Chavalarias et al.(2014)Chavalarias, Leenart, and Panahi] [5], subjects receive their incentive from the perspective of earning monetary payoffs if they have accumulated the highest rewards at the end of the experiment. For that reason, this model has been called the *Nobel game*.

Game theory and agent-based modeling consider roughly two categories of agents, forward-looking eductive ones, and backward-looking adaptive ones. Sometimes agents are both, but modelers tend to choose sides. The first study of the Nobel game considered forward-looking agents in the pure tradition of game theory [Chavalarias, D.(1998)]. Under the representative agent hypothesis, in which agents maximize their payoffs over a lifetime, it was demonstrated analytically that the optimal strategy was in the form of evolving stopping times for publication and refutation processes, plus an evolving preference for publication vs. refutation activities. The hypothesis of infinitely rational agents, although widespread in the game theory, has been increasingly criticized as unrealistic [Aumann(1997), Rabin(2002)]. Furthermore, the representative agent hypothesis turned out to be an oversimplification that failed to account for heterogeneous strategies observed in experimental studies of the Nobel game [Chavalarias et al.(2006)Chavalarias, Charron, Gardelle, and Bourgine, Chavalarias et al.(2014)Chavalarias, Leenart, and Panahi].

To overcome the limitations of the first modeling attempt, [Chavalarias et al.(2006)Chavalarias, Charron, Garde and [Chavalarias and Gardelle(2008)] studied the Nobel game with agent-based simulations. Their preliminary results, based on a variant of the model presented hereafter (with different hypotheses

5. See also the online experiment : http ://nobelgame.org

H.1, H.4-6), account for heterogeneous populations and reported the existence of a trade-off between speed and quality of the discovery process.

The present study proposes some concepts to think about when it comes to science dynamics, investigates more deeply the dynamics of the Nobel game through sensitivity analyses, and confronts the results with empirical data.

Adopting Schelling's perspective on models, our main goal is "to illustrate the kind of analysis needed, some of the phenomena to be anticipated, and some of the questions worth asking" [Schelling(1978)], rather than to simulate the real dynamics of science with higher fidelity.

2 Modeling science dynamics with the Nobel Game

2.1 Hypothesis and model description

The basic idea for transmitting the core dynamics of evolution in science is simple : scholars publish theories and can falsify theories of others. Publication and refutation are rewarding, which serves as the base for their future decisions [6].

To give the reader some guidelines for understanding of model, we will illustrate our general description of collective discovery processes with the example of the discovery of genetic interactions in the yeast genome [7]. This discovery process has been studied empirically by [He and Zhang(2009)] with a clear definition of what a "theory" could be : "Two genes are said to interact genetically if the effect of one gene on a trait is masked or enhanced by the other." In this area of science, biologists have been working on the identification of all genetic interactions (GI) of the yeast's genes, and each discovery of such an interaction is recorded by a publication. These publications can stand for elementary theories that could be corroborated or falsified by other members of the community.

For clarity in what follows, we will number the main modeling hypotheses (e.g. H.1).

2.1.1 Publication activity

We will model the publication activity as a two-step process :

1. A scientist (or a team) establishes conjectures and tests them step-by-step until they are falsified or corroborated enough to be worth publishing (*e.g.* a sufficient number of experiments revealed a GI). This is the private part of the process of knowledge production.
2. Next comes the public part. When a theory is published, it becomes common knowledge throughout the community, and anybody can try to falsify it (*e.g.* replicate experiments). The theory is accepted until someone proves it false. We will note CK (for common knowledge) for the set of (temporarily) accepted theories (*e.g.* the set of all published GI).

2.1.2 Agent population

We will consider a population $\wp$ of N_a agents interacting on a network Γ (team, networks of collaborators, close colleagues, etc.). Agents can only observe the behavior and strategies of their neighbors in the network Γ. For any agent i, the set of its neighbors will be noted Γ_i.

6. We will consider only rewards and losses associated with the fact of being published or falsified. We won't consider other feature like citations, although they also play an important role in the dynamics of science.

7. This is only for illustration purposes, the model itself is not limited to this particular case of empirical science.

2.1.3 Modeling the knowledge space

We will consider Popper's definition of theories as statements dividing precisely *"the class of all elementary statements into two non-empty sub-classes : the one of all the statements with which it is in contradiction [...] and the one of all the statements with which it is not in contradiction"* [Popper(2002)].

Let us consider a set Θ of possible worlds and a set Φ of objects called theories. A possible world $\Omega \in \Theta$ is defined as a set of elementary statements ω, each describing issues of a particular experiment. These elementary statements can be used to refute the theory (for example, "Activation of $gene_1$ *always* enhances the activation of $gene_2$" can be falsified by "It has been observed that the activation of $gene_1$ does not enhance the activation of $gene_2$ when $gene_3$ is inhibited").

A theory is a statement about universal properties of the form "Every time $gene_1$ is activated, the activation of $gene_2$ is enhanced" or "Every time $gene_3$ is inhibited, the activation of $gene_1$ does not enhance the activation of $gene_2$". It is modeled as a function $\varphi : \{\omega | \omega \in \Theta\} \mapsto \{0, 1\}$, which describes whether a particular elementary statement contradicts or corroborates the theory. Given a theory $\varphi \in \Phi$, agents can design experiments to gain knowledge concerning the valid statement ω about their world, and they compare the output of these experiments to the predictions of φ. We can encode the results of such experiments with 1's and 0's depending on whether the theory is corroborated or refuted [8].

We will make the simplifying assumption that the set of all possible theories about a given class of phenomena (*e.g.* the GI interactions) is finite. Moreover, whereas former studies of scientific dynamic modeling (*e.g.* [Gilbert, N.(1997)] or [Edmond, B. et al.(2011)Edmond, B., Gilbert, N., Ahrweiler, P., and Sch have proposed to model the structure of theory spaces (the predictions of some theories are correlated - *e.g.* because one logically depends on an other), we will assume that theories are independent from each other.

H.1 : Theories are Bernoulli random variables. The simplest way to model the theory space is to consider a finite set Φ of random variables φ_i with the Bernoulli's law of parameter p_i (*i.e.* the output of each random draw is 1 with a probability p_i). The size of this set will be noted N_T (the total number of possible theories). An agent can perform an experiment concerning a theory φ_i by observing random draws of φ_i. φ_i will be corroborated with the probability p_i and falsified with the probability $1 - p_i$. Each random draw is time-consuming and will take one unit of time to be achieved.

In simulations thereafter, we will consider an *epistemic landscape* consisting of N_T theories that can only be of two types : "true theories" (type I : Bernoulli variables with parameter 1, every draw gives a positive results) and theories with some errors or inaccuracies (type II theories : Bernoulli variables with parameter $q < 1$, meaning that they can be falsified with a probability of $1 - q$ per time unit). We will note p to be the proportion of type II theories.

Along with H.1, we assume that all agents are evenly quick in the process of formulating and testing theories, which means that they all have the same skill for science, all provide the same level of effort in achieving their research and all have the same resources to conduct their investigations [9].

Theories can be in three possible states :

8. An alternative approach could be to consider predictions in probability rather than deterministic predictions.
9. Note that if the first two assumptions were true in real life, there would be no point evaluating scholars ; the last assumption is definitely false.

1. **Published.** The set of published theories that have never been falsified is common knowledge and named CK,
2. **Unpublished.** The set of such theories will be named $\overline{CK}$.
3. **Falsified.**

2.1.4 Research activities

Agents can undertake two kinds of activity, where the exploration-exploitation dilemma is formalized as *stopping times* (cf. fig. 1) :
— **Publication :** an agent draws a $\overline{CK}$ theory at random that he has not already examined, and decides to spend at most $\overline{\lambda}$ time units to test it. If after all these tests, the theory has never been falsified (no "0" has been drawn), he will publish this theory. Agents are assumed to be honest. If an agent obtains a negative result during his tests (a "0" outcome), he will stop immediately and begin a new research process. The record of this refutation will allow automatic falsification of this theory if it were to be published by an other agent in the future.
— **Refutation :** an agent draws a CK theory at random and decides to spend λ time-steps to find a negative result. As soon as he finds one, he publishes the negative result (*e.g.* "It has been observed that activation of $gene_2$ ***is not*** enhanced by the activation of $gene_1$ when $gene_3$ was activated") and begins a new research process. If he has not found any negative results after λ time units, he starts a new research process.

The choice between these two activities, publication and refutation, is modeled by a Bernoulli random variable ν that represents the preference of an agent for the publication activity. At the start of each new research process, if there are some published theories, an agent i chooses with a probability $1 - \nu_i$ to go for a process of type *refutation*. Otherwise, if there are some $\overline{CK}$ theories left, he begins a research process of the type *publication*.

The strategy of any agent i can thus be defined as a triplet $(\overline{\lambda_i}, \lambda_i, \nu_i)$ that determines the *stopping times* in each type of research process, plus the preference for publication.

Each process can end prematurely if someone else publishes before the agent is finished working on the theory, or if a negative result is published about the theory he is attempting to falsify.

H.2 : Agents' strategies. Agents perform publication and refutation activities separately. Their strategies are defined by *stopping times* $\overline{\lambda}$ and λ, and a preference for the publication activity ν.

Negative results found during publication processes are kept in mind for the potential refutation of forthcoming published theories (we will assume that this refutation is made at no cognitive cost, *i.e.* it is automatic).

2.1.5 Science policy and social reward

H.3 Agent's payoffs. Payoffs are associated with each publication event. They represent their utility from the point of view of the agents :
— P for the publication of a new theory,
— R for the refutation of an existing theory,
— L for being falsified.
In most cases, we can think of *P*, R and L as being respectively positive, positive, and negative (L is a loss). Several events can affect the total payoffs an agent receives during a given period. For example, he can publish a theory (winning *P*), but several of his theories can be falsified (losing L

several times) and one theory just published can be automatically falsified (winning R). The overall payoffs earned by an agent i at period t is noted g_i^t. A "science policy" will be defined by a triplet $\{P, R, L\}$ [10].

Agents will be ranked according to their aggregate score, π. In the present simulation study, we will set this aggregate score to be the sum of the agent's payoffs during his lifetime. Note that the definition of this score is very important and can have a great influence on scientific dynamic. There is a certain awareness in academia about this issue, in particular in scholar evaluation. Even if this were to take into account the sole publication rate, it is still not clear what should be the appropriate period for computing this index. Should it be a whole career ? The last three years ?

H.4 Social comparison Comparison between agents takes place through their cumulative payoffs $\pi(T)$ at time T. For an agent i, $\pi_i(T) = \sum_{t \leq T} g_i^t$.

2.1.6 Adaptive strategies

We will consider backward-looking agents, adapting their strategies through social comparison and imitation. Imitation is a key process in social learning and widely used in social modeling and simulation [Conte and Paolucci(2001), Chavalarias(2006)]. From the modeling perspective, it has the advantage of requiring weaker hypotheses than forward-looking agents models, both from a cognitive and computational point of view.

H.5 Agents' information. Agents know the higher π score in the population and their own rank on the π score scale. They know the π scores of their neighbors in Γ and, to a certain accuracy, their strategies (*i.e.* for an agent i, $(\overline{\lambda_j}, \lambda_j, \nu_j)_{j \in \Gamma_i}$).

The modeling of imitation rules features three components : when to imitate, how to choose the models and what is the copying process. We chose here some standard options, although the investigation of variants of these options could be interesting because they correspond to local cultures in academia.

— **H6.1 When to imitate.** Agents are more susceptible to change their strategy when they are ranked low down on the π score scale. Every time they end a research process, they engage in an imitation process with a probability proportional to their π rank. For example, for a population of 100 agents, the second-ranked agent will have a probability of $\frac{2}{100}$ to revise its strategy before starting a new research process, while the 98th ranked agents will have a probability of $\frac{98}{100}$.

— **H6.2 Who to imitate.** When updating their strategy, agents imitate one of their top scoring neighbors in Γ.

— **H6.2 How to imitate.** Agents will copy the best agent's strategy with some gaussian error (perception and implementation are noisy). The copying function for any X in $\{\overline{\lambda_j}, \lambda_j, \nu_j\}$ is $C(X) = X \times (1 + e \times \varepsilon)$ where ε is a random variable with standard normal distribution, and e a factor determining the precision of the copying process. We will consider $e = 0.05$ in all the following part of this paper. Ideally, the function C, and its parameters, should be chosen according to the research results in the field of experimental psychology.

10. Note that this is a double short-cut since, even if publication incentives were to be reduced to theses rewards, which there are not in the real world, P, R, L will have to include both endogenous rewards (reputation, recognition of peers, etc.) and exogenous reward from the academic establishment (the science policy).

To summarize the collective discovery process (cf. fig. 1), we created an evolutionary process in an evolving environment of a population of agents. At each time period :

— each agent performs a single test on the theory he is studying, either to publish or with the hope of falsifying it,

— for a publication process, if the stopping time $\overline{\lambda_i}$ of agent i is reached without any refutation event and the theory has not been published so far, the agent publishes the theory, which becomes common knowledge,

— for a refutation process, the agent tries to falsify a theory as long as the theory is in the CK set and its stopping time λ_i is not reached,

— at the end of every process, with a probability proportional to their rank in the whole population, agents perform social comparisons on the basis of their score $\pi_i = \sum_{n \leq t} g_i^{t-n}$ and eventually engage in an imitation process, copying the best agent in their neighborhood Γ_i, with some Gaussian error.

— agents ending a research process or having been interrupted by a concurrent publication or refutation of the theory they are working on, choose a new type of research process according to ν, and pick-up a new theory at random in the appropriate set.

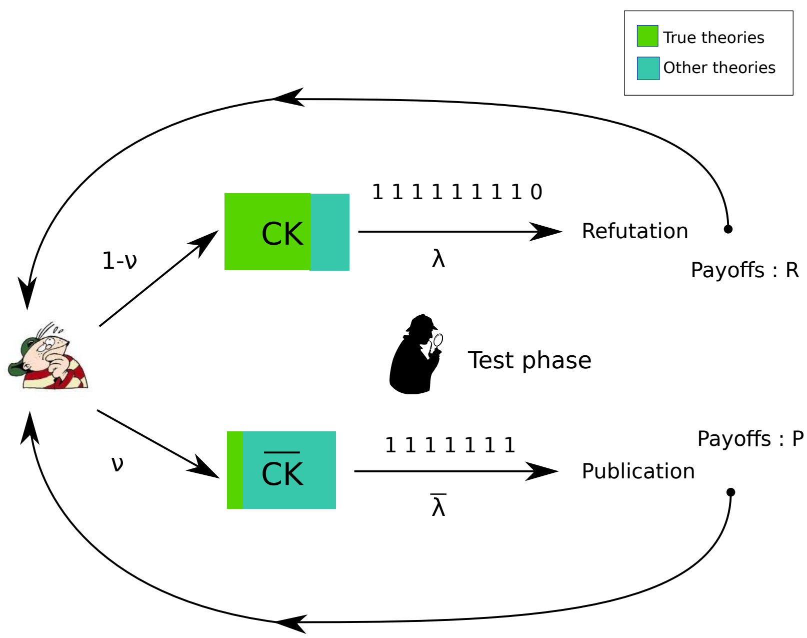

FIGURE 1 – Activity of a researcher : a scientist has the choice between working on an unpublished theory (from the $\overline{CK}$ set) and trying to falsify an existing theory (from the CK set). He takes the first option with probability ν each time he starts a new process. In the first case, he publishes his theory (and obtains P) if he acquires $\overline{\lambda}$ consecutive positive tests. In the second case, he publicly refutes this theory (and obtains R) if he succeeds in finding one negative result in less than λ time steps. Each of these processes could be interrupted by, respectively, a negative test or the publication of this theory by an other agent ; the publication of a refutation for this theory by an other agent.

2.2 Characterization of science dynamics

The Nobel game is a game in a non-stationary environment (not only the social environment changes, but also the epistemic landscape depletes), with some stochasticity (there is some noise on the agents' perception of other's strategies). To account for the rich variety of patterns observed in the sensitivity analysis, we will use the following notations and measures :

— $p^t_{published}$: the proportion of theories of type II in the CK set of published and accepted theories (falsified theories are not included),
— $n^t_{published}$: the total number of theories in the CK set of published and accepted theories,
— **Quality.** One way of assessing the quality of a set of theories is to estimate the probability that a prediction of a theory drawn at random from this set is revealed to be false. Different domains might have different levels of requirements in terms of the acceptable level of quality. The cost of being wrong differs if you are building nuclear plants or aiming to improve brewing. We define the quality of a set of published theories at time t as :

$$q(t) = -log10[(1-q) \times p^t_{published}]$$

A quality of 4 means that you have 10^{-4} chances to have a negative result when testing or applying a theory from this set at time t. The overall quality of a discovery process until the time T will be defined as the average quality of the published theories until *T*, *i.e.*

$$Q(T) = < q(t) >_{t \leq T}$$

For clarity of the plots, we will threshold to 10 the quality measures (which is equivalent to saying that theories with 10^{-10} or less chances of being false are indistinguishable from true theories).
— **Speed.** The speed of a discovery process will be defined as the average number of type I theories discovered per time unit per agent until t. The speed at time t will thus be defined as $S(t) = \frac{(1-p^t_{published}) \times n^t_{published}}{t \times N_a}$,
— **Achievement.** Achievement at time t is the proportion of type I theories that have been discovered so far. It is defined as $A(t) = \frac{n^t_{published} \times (1-p^t_{published})}{N_T \times (1-p)}$

In simulation studies, we can compute the earliest period where all type I theories have been published ($A(t) = 1$) and no type II theories remains in CK ($q(t) = \infty$). This period will be noted T_{end} and will be used as an upper bound to draw the plots.

Neither speed nor achievement is directly observable from empirical data, particularly because the quality of published results cannot easily be assessed. As emphasized in the introduction, it is impossible in real life to know when T_{end} is reached for a given scientific question. Nevertheless, we can compare the dynamics of the total number of publications observed empirically with simulation data (see section 5).

There are several ways to reach T_{end}. The science community could go there by achieving different values for the quality and speed of the discovery process, their dynamics being influenced by scientific policies. One of the challenges of science dynamic modeling is to explain or predict the influence of scientific policies on these social dynamics. Moreover, scientific policies often take into account some objectives in terms of speed and quality. If you want to build or improve nuclear plants, you can't afford to have anything but the highest quality results all along the way to the end of the discovery process ; in contrast, if you are improving brewing, you might prefer to go faster and tolerate lower quality at every step of the discovery process.

Due to the complexity and stochasticity of the dynamics for a given science policy $\{P, R, L\}$, the collective dynamics can still exhibit a wide variety of patterns (different speed and quality), as it does in simulations.

3 Agent-based exploration of the Nobel game

We present an exploration of the Nobel game dynamics using agent-based simulations, for which pseudo-code is described in appendix A2.5. [11].

Empirical studies about networks of scientific collaborations revealed that they have a *small world* structure [Newman(2001)], with short average path lengths between two nodes of the network and a high clustering coefficient. In these simulations, we generated Γ with the simplest model of *small world* network, described by [Watts and Strogatz(1998)]. In these networks, agents are arranged in a ring and connected to their n closest neighbors, with some of these connections being rewired once and for all with a probability of ρ. The mean number of collaborators per author for the scientific fields studied by [Newman(2001)] range from 3.59 to 18.1, except in high-energy physics where it reaches 173. In [Watts and Strogatz(1998)], $\rho = 0.1$ is about the threshold where small world networks have the shortest characteristic path length but still have a high clustering coefficient of up to 80 percent of the value of the associated regular lattice. For these reasons we considered $n = 6$ and $\rho = 0.1$ as being reasonable values to generate, for each simulation run, a random *small world* network Γ.

The *epistemic landscape* consists of N_T theories with a proportion $(1 - p)$ of type I theories and p of type II theories with a parameter $q < 1$. This bimodal distribution is the most parsimonious configuration that makes it possible to explore the influence of the "difficulty" of the field on science dynamics : domains with a high proportion of type II theories with a high q parameter (it's hard to find a type I theory and it's difficult to distinguish between the two types) are more difficult to study than those with a high proportion of type I theories and low q parameters for type II theories. In addition, the number N_a of agents working in the field and the richness of the field (number N_T of theories) adds a supplementary difficulty by increasing the competition level. The higher is the ratio $\frac{N_a}{N_T}$, the stronger is the competition.

In this agent-based exploration of Nobel game dynamics, we have studied the influence of the difficulty of the field (influence of p, q) and of the science policy $\{P, R, L\}$ on science dynamics. The set of parameters used for these simulations is described in table 1.

3.1 Single experiments analysis

To familiarize the basic concepts associated with Nobel games, we first present two case-studies corresponding to two simulation runs with very different values for *P*. All the parameters of these simulations are given in table S1 (except those indicated as fixed in table 1).

As can be seen in supplementary figures S1 to S7, the two simulations have few common features :

— **Strategies $(\overline{\lambda_i}, \lambda_i, \nu_i)$ are heterogeneous in the population.** (fig. S7). This kind of heterogeneity has a been observed experimentally in [Chavalarias et al.(2006)Chavalarias, Charron, Gardelle, and Bo Chavalarias et al.(2014)Chavalarias, Leenart, and Panahi].

— **Average stopping times keep evolving throughout the simulation with large variations over short time periods.** These oscillations in average stopping times and average publication propensity are typical of mimetic dynamics and can be understood to reflect the social game of science, as described in figure 2. Agents are constantly adapting to the evolution of the environment and strategies of others who they imitate. Mimetic dynamics are known to produce strong positive feedback and consequently abrupt changes.

11. These simulation have been implemented with MatLab. Pseudo code for the algorithm is given in appendixes.

$Parameters$	Description	Values	Remarks
N_a	Number of agents	100 ; 1,000	Can be measured empirically
N_T	Number of theories	3,000 ; 30,000	Other values have been investigated
P	Payoffs for a publication	1 < P < 100	P=1 ; 5 ; 10 ; 20 ; 50 ; 100
R	Payoffs for a refutation	1 < R < 100	R=1 ; 5 ; 10 ; 20 ; 50 ; 100
L	Loss for being refuted	1 < L < 100	L=1 ; 5 ; 10 ; 20 ; 50 ; 100
p	Proportion of Type II theories	0.9 < p < 0.999	p=0.9 ; 0.95 ; 0.99 ; 0.999
q	Bernoulli parameter for type II theories	0.9 < q < 0.99	q= 0.9 ; 0.95 ; 0.98 ; 0.99
Γ	Parameters for the small world network	$n = 3$; $\rho = 0.1$	Same for all simulations. Values chosen according to [Newman(2001)] and [Watts and Strogatz(1998)].
e	Error rate of agents when copying neighbors strategies	$e = 0.05$	Same for all simulations. Could be measured experimentally.

TABLE 1 – Summary of parameters used for the reported sensitivity analyses.

TABLE 2 – Summary of the parameters used for case-studies 1 and 2.

Description	Values
Number of agents	$N_a = 1,000$
Number of theories	$N_T = 30,000$
Payoffs for a publication	$P = 1$ (case1) ; $P = 100$ (case2)
Payoffs for a refutation	$R = 10$
Loss for being refuted	$L = 10$
Proportion of Type II theories	$p = 0.9$
Bernoulli parameter for type II theories	$q = 0.98$

— **The average instantaneous score $< g_i^t >_{i\in\wp}$ peaks at the beginning of the discovery process** and then decreases significantly. This means that with regards to the payoffs associated to publications, scientific fields are much more attractive in their earliest years than when they are mature. This is in line with the "hot topic" effect that we all perceive when a new field emerges. Although the two discovery processes have very different speeds, the early stages have about the same duration (cf. insets of figures S6).

— **The distribution of the number of papers per authors is highly skewed** (with a Gini index of between 0.3 and 0.45, cf. fig. S3). It is like a Zipf law (cf. fig. S2), although the range covered (between 0 and 20 publications per authors) is not large enough to significantly fit a power law. These highly skewed distributions are a well-known empirical fact, and this was one of the first patterns modelers tried to reproduce, such as in the case of [Gilbert, N.(1997)]. It is noteworthy that the Nobel game, as well as previous research, demonstrates that differences in publication levels could have nothing to do with scholars' skills or abilities. In our case, it only reflects the contingencies of the publication and refutation process. Thus, one might ask what is the meaning and usefulness of scientometrics indexes based on these publication

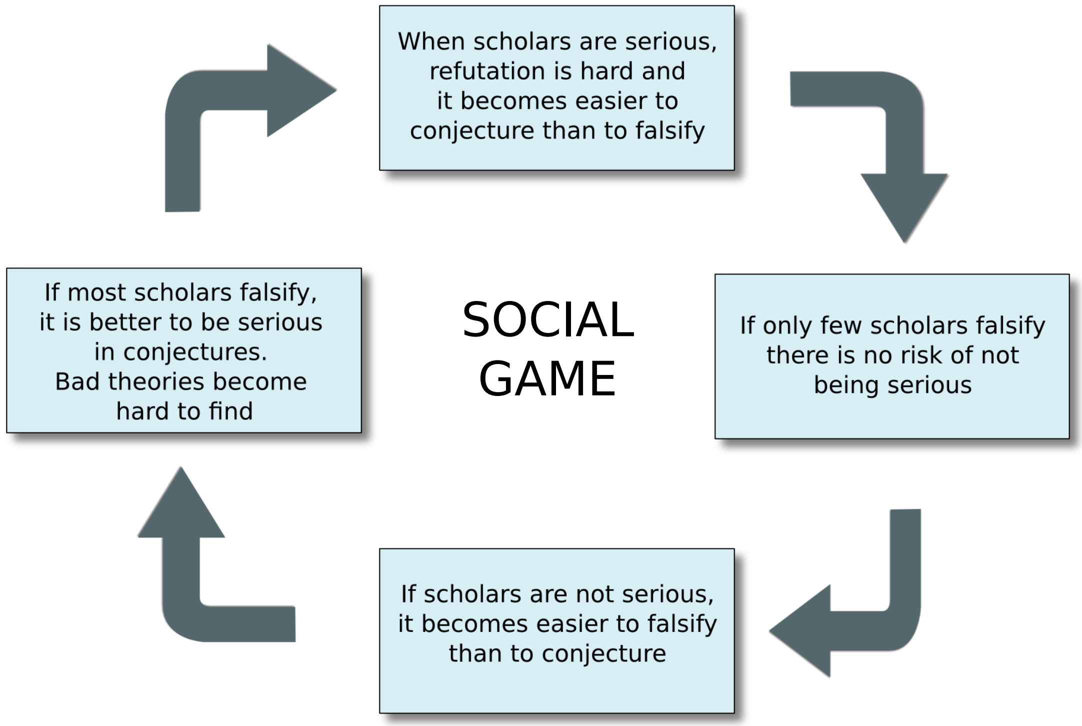

FIGURE 2 – The social game of scientific discovery

levels, such as the $h - factor$.

However, the dynamics of these two discovery processes are very different, as can be seen in figures S1 to S7. When incentives for publication are low, the discovery process is slow with a high quality at all times. When incentives are high, by contrast, the discovery process is as much as five times faster on the achievement scale, but the quality drops considerably, as summarized in table 3. These differences in speed and quality are due both to differences in agents' propensities for publication ν_i and differences in stopping times. When incentives for publication are increasing, agents tend to neglect refutation processes. Meanwhile, their stopping times are decreasing on average, both for publication and refutation processes (cf. supplementary fig. S4). These two phenomena lead to a decreased overall quality of publications. This can be clearly seen in supplementary figure S5. Whereas, when publication incentives are low, type II theories are kept at quite a low level all along the discovery process in the CK set of published theories, they are in a significant proportion when incentives are high, even for high values of achievement.

TABLE 3 – Differences in speed and quality between case-studies 1 and 2 at $A = 0.5$. All values of parameters are the same for both cases, except P.

Simulation	Number of periods	Quality	Speed
Case 1 ; $P = 1$	$4,734$	4.2	1.0^{-4}
Case 2 ; $P = 100$	$1,499$	2.63	3.34^{-4}

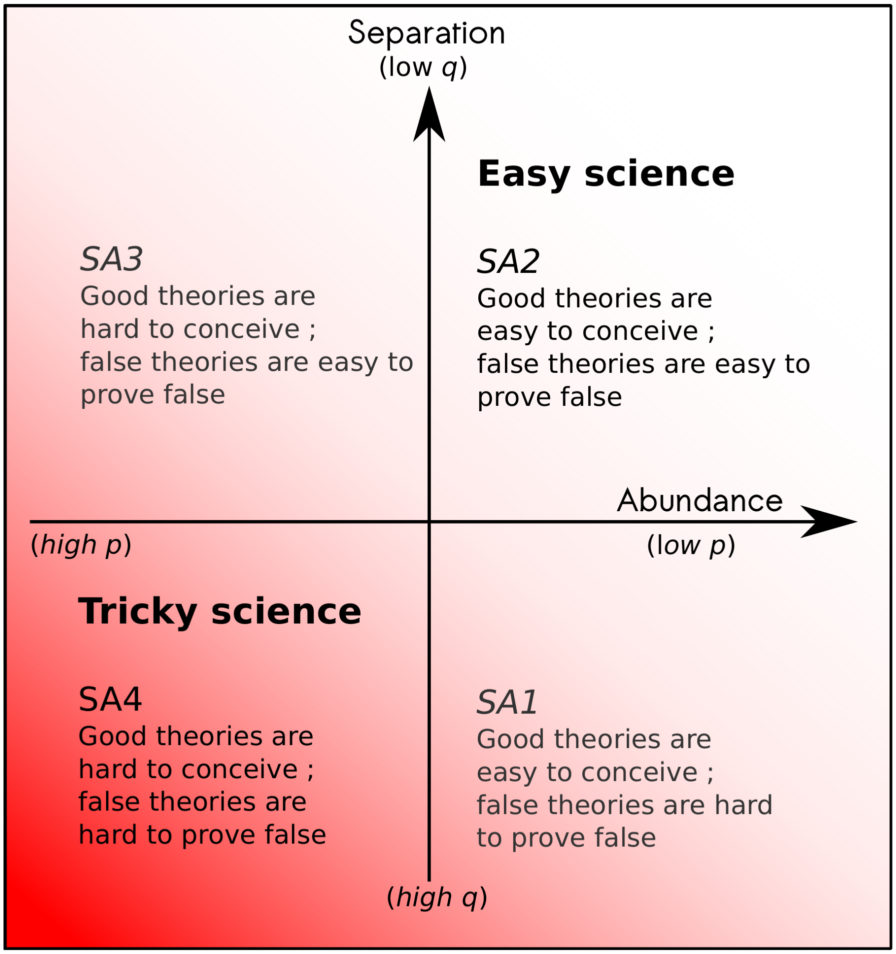

FIGURE 3 – Characterization of different types of science. Sensitivity analyses $SA1$ to $SA4$ have been designed to explore different areas of the abundance/separation space.

3.2 Scientific policies and the speed-quality dilemma

The two previous case-studies, as well as our daily scientific life experience, teach us that at the individual level there is a trade-off between the speed and the quality of the research we are conducting. Contemporary science policies are orientated toward productivity, as we have discussed in section 1.1. How do science policies influence the social game of science, the individual exploration-exploitation dilemma, and *in fine* the quality of the scientific production ? These questions can be investigated by exploring the influence of the incentives for publication P in the Nobel game as well as the reward and loss (R and L) for refutation.

We should also expect that the impacts of scientific policies depend on the nature of the scientific field to be explored. If we think in terms of our simple models of knowledge spaces, when p is low, type I theories are abundant and agents have a high chance of finding good theories before they proceed to the test phase. When q increases, it becomes more and more difficult to distinguish between type I and type II theories in the test phase ; we can say that different types of theories are hardly separable. *Abundance* and *separation*, along with the size of the field (total number of theories) are parameters that are likely to influence the knowledge discovery process.

For example, if the gene-gene interactions network of a given species was sparse, but the protocol for testing a given relation was quite standard and easy to perform with modern equipment, the issue of finding all the gene-gene interactions for that species would define a field with high *separation* but with a low *abundance* (a lot of interactions do not exist due to the sparsity of the GI network). In contrast, if theories are dealing with objects that are at the limit of what our technology can measure (such as with particle physics), it should be expected that *separation* is weak.

To give a sense of how *abundance* and *separation* influence the sensitivity of the collective discovery process on the incentive level for publication, we have conducted four sensitivity analyses on P (cf. figure 3) for high and low values of p (resp. 0.9 and 0.99) and q (resp. 0.95 and 0.98).

These analyses present some common features that we illustrate with the sensibility analysis $SA1$ defined in table 4. The parameters are the same as for case-studies 1 and 2, except that we took 10 times fewer theories and agents, to limit the computational load. In what follows, this sensitivity analysis will be referred to as $SA1$ (*sensitivity analysis 1*).

TABLE 4 – Summary of the parameters used for $SA1$ sensitivity analysis.

Description	Values
Number of agents	$N_a = 100$
Number of theories	$N_T = 3000$
Payoffs for a publication	$P = 1,\ 5,\ 10,\ 20,\ 50,\ 100$
Payoffs for a refutation	$R = 10$
Loss for being refuted	$L = 10$
Proportion of Type II theories	$p = 0.9$
Bernoulli parameter for type II theories	$q = 0.98$
Number of independent simulations per points	100

$SA1$ confirms the trends identified in case-studies 1 and 2, and gives insights into several important phenomena.

First, the same scientific policy $\{P, R, L\}$ can result in very different outputs (cf. fig. 4). Although the increase in the incentive for publication tends to decrease the quality while increasing the speed, the relation between speed and quality is very constrained, the variance in the speed to quality space for a given policy and for a given level of achievement is quite high. This is due to the strong path dependency of social dynamics driven by imitation, where local contingencies can generate very different evolutionary paths. This can be observed in figure S8. This property of Nobel game dynamics, with this choice of social learning by imitation, has major consequences for the interpretation of empirical data. Scientific dynamics and this version of Nobel game dynamics could have similar properties. If this were so, then predictions of scientific dynamics knowing part of the publication history could only be done in probability with a high variance, even with the best estimation of all parameters.

Second, there is a clear trade-off between the speed and the quality of the scientific discovery process, with fast discovery processes typically showing lower quality (cf. fig. 4).

Third, as demonstrated in supplementary figures S11 and S12, when the incentives for refutation are unchanged, raising the incentive for publication drives, on average, the collective dynamics of science toward higher speed and lower quality. The strength of this trade-off depends on achievement. For $SA1$, in the early stages of the discovery process, the *publish or perish* policy has almost no influence on the speed or the quality, especially when the process peaks at its maximal speed. But quite quickly (around $A = 0.1$), the impact of the speed-quality trade-off becomes significant, and

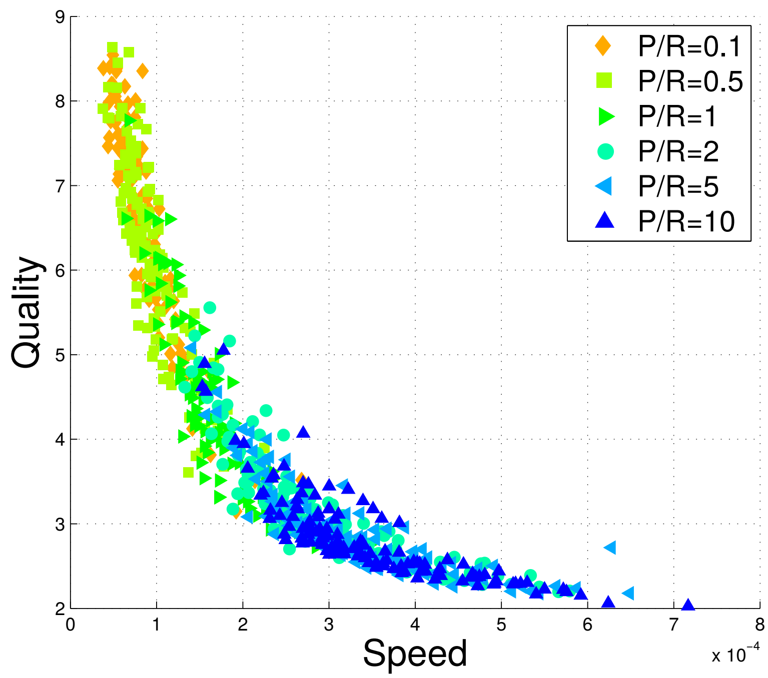

FIGURE 4 – **Speed-Quality at $A = 0.5$ for $SA1$ simulations.** Each point corresponds to the output of a simulation at $A = 0.5$. For each value of publication payoff P, 100 simulations are plotted. Only the incentives for publication P vary in the scientific policy, R and L being constant at 10. We observed the clear trade-off between speed and quality of the scientific discovery process, influenced by P. However, the same scientific policy $\{P, R, L\}$ can result in very different outputs.

any increase in speed or in publication incentives is done at a huge cost in terms of quality of the scientific production. This means that in early stages of the collective discovery process, the *publish or perish* policy has almost no influence on science output, and thereafter it significantly decreases the quality of the published theories with only moderate benefits in terms of speed.

This speed-quality trade-off was confirmed by $SA2$ to $SA4$ with slight variations in patterns for low values of achievement ($A \leq 0.1$, cf. S3 to S5).

As summarized in figure S10, it should be expected that increased pressure on the scholar to publish is immediately translated into a higher speed of discovery, but also a much lower quality. The exploration-exploitation dilemma that operates at the individual level has its counterpart dilemma at the collective level : a speed-quality dilemma.

Many will say that this conclusion is no surprise. Yet we all continue to play a game of science that is getting faster and faster.

The Nobel game can teach us much more than that. Are there better practices to regulate science dynamics than just to increase the incentives for publishing ? How should we reward refutation and deal with refuted scholars ? Is there a difference in the policy to be applied depending on the age of the field, its difficulty, the size of the community ? There is no space here to give a detailed answer to these questions, but from the many studies we have conducted so far with this model, through an analytic approach [Chavalarias, D.(1998)], multi-agents approach [Chavalarias et al.(2006)Chavalarias, Charron, Gard Chavalarias and Gardelle(2008)] or experimental approach [Chavalarias et al.(2014)Chavalarias, Leenart, and Pana we can claim that this model undoubtedly has an important heuristic power to help us answer these

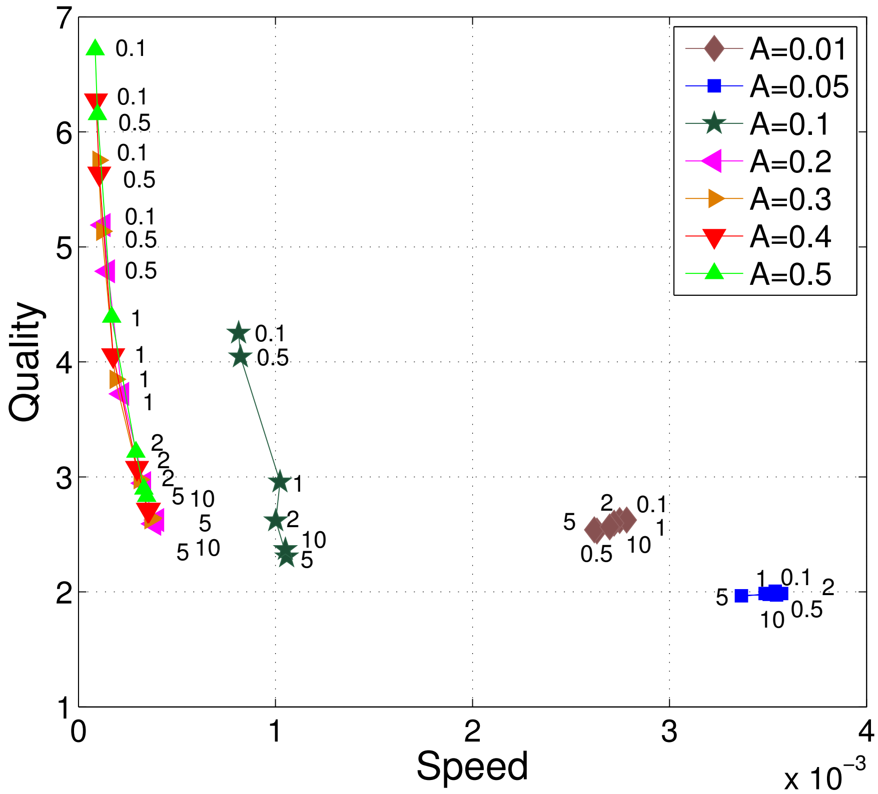

FIGURE 5 – **Average speed-quality diagram for $SA1$.** Influence of the incentives for publication (P) on the speed and the quality of the discovery process for different degrees of achievement A. The ratios P/R are indicated next to the markers. The population size is $N_a = 100$. Each data point is an average of 100 simulations. Only the incentives for publication P vary in the science policy, R and L being constant at 10. Except for emerging fields ($A < 0.1$), there is a clear trade-off between the speed and the quality of the discovery process, with a rapid decrease in quality as soon as the incentive to publish increases. Confidence intervals for speed and quality averages are given in supplementary figure S11.

non trivial questions.

4 Comparison with previous work

The results of the analysis of the Nobel game dynamics are in line with previous findings in science dynamics modeling.

[Gilbert, N.(1997)], [Sun and Naveh(2009)] and [Borner et al.(2004)Borner, Maru, and Goldstone], among others, have modeled the dynamics of co-authorship and citation networks. In particular, these models could reproduce the highly skewed distribution of papers per authors. They demonstrated that this kind of distribution could emerge, in the absence of cognitive differences between scholars, and for a wide range of cognitive settings [Sun and Naveh(2009)].

[Zollman(2007)] considered agents exchanging views about states of the world, while making theories about it. His model analyzes the convergence of the beliefs within the population, and he demonstrated, with agent-based simulations, that the topology of interactions between agents influences the speed and quality of the discovery process. This influence also takes the form of a trade-off between speed and quality : an increase in network connectivity increases the speed of the collective discovery process, but it decreases its quality. This results could guide future analyses about the influence of topology in the Nobel game.

[Edmonds(2008)] developed a model in which scientists are represented as theorem provers, generating new theorems by inference from existing premises. This model is an attempt to understand an explicit epistemic landscape in which some locations are more difficult to discover than others, although few results exist about the behaviors of this model. Moreover, as in [Weisberg and Muldoon(2009)], it is assumed that "there *is* some independent 'correct' knowledge to be discovered and that it is checkable" [12]. Thus the exploration-exploitation dilemma that scholars face in everyday life could not be addressed within this class of models.

The most significant difference between the Nobel game and previous models of science is the emphasis on the relation between scientific policies and the social dynamics of science. Whereas some other aspects and associated results can be intersected with previous work, this aspect is, as far as our knowledge is concerned, specific to the Nobel game.

5 Comparison between Nobel Game dynamics and empirical data

We can question the relevance between the Nobel game dynamics and real science dynamics. We already emphasized that many other factors contribute to the dynamics of science (such as scholars mobility, scholars' turn-over, funding policies, fads, etc.). They could all interfere with conjecture and refutation dynamics. Consequently, it should not be expected that a Nobel game, in such a minimal form moreover (with a simple knowledge landscape and only one field of expertise), could account for *in vivo* science dynamics. Besides, if Nobel game dynamics were indeed a key component of science dynamics, it should be expected to be a non-deterministic correspondence between the policy or environmental settings and the evolutionary path of science would make the inference of hidden parameters non-trivial.

We can, however, check that Nobel game dynamics are compatible with science dynamics when observed empirically. There are some cases where the Nobel game dynamics have some chances to be predominant over the other factors. The case of ambulance chasing in particle physics is one of them. Ambulance chasing refers to situations where some recent data disagree with the Standard Model of particle physics, and researchers come up with an interpretation in terms of new physics [Allanach(2014)].

As one can imagine, ambulance chasing are, in fact, races for being the first to propose the "good" interpretation for weird results. Since particle physics scholars usually post their papers in a pre-print archives as soon as they think it is "sufficiently" reliable to be published, we have a daily record of the publications and can compare the evolution of the number of publications in some ambulance chasing cases with some Nobel games dynamics.

[Backović(2016)] proposed a theory of ambulance chasing at the macro level, reconstructing the evolution of the number of publications. His model assumes that publications follow a Poisson process with two parameters that reflect the evolution of the available number of topics and the interest of scholars in the field. This approach is to the Nobel game what macro-economy is to

micro-economy – very different but potentially complementary. The data collected by Backović to test his model corresponded to nine recent instances of ambulance chasing from inSPIRE and arXiv repositories. He obtained the cumulative number $N_T(t)$ of published papers on a topic as a

12. [Edmonds(2008)] , p66.

function of time, by extracting the lists of citations to the result that initiated the ambulance chasing instance. The author recognizes some limitations of the method, but they are assumed “to result in systematic errors of $O(10)$ papers in total (per data set) and will hence typically be smaller than the statistical error”.

Our model features several parameters for describing the micro-dynamics of science. Trying to find the best fit for given empirical data would require specific optimization methods, such as genetic algorithms. Moreover, particle physics scholars often publish in large groups or consortia [13], which means that inferring the relevant unit for the private research process is not straightforward. Overcoming these two issues goes beyond the objectives of this paper. Our goal is to gain a first insight about the possible fit between Nobel game dynamics and real science dynamics. Thus, we have confined our analyses to determine the best fits in terms of publication dynamics among the simulation runs of our sensitivity analyses for each of the nine cases of ambulance chasing ($SA1$ to $SA4$). One important parameter for science dynamics is the characteristic time scale of experiments, which can be modeled by the average duration of an experiment. It should be expected that this characteristic time-scale varies across fields of science. To find this parameter s, we stretched the time axis of the simulations so that there was at least one common point between empirical and simulated data in addition to the origin $(0, 0)$. The number of published theories being a monotonous function, we chose s such that the number of publications in the Nobel game equals the number of arxiv.org pre-prints at the most recent data point of the empirical data. s gives the equivalent (in days), of the duration of an experiment in the Nobel game, $1\ experiment = s * Days$, whatever the notion of experiment could stand for. Note that s cannot be properly inferred without a precise estimation of the size of the community, and consequently, an interpretation of the s values would be risky at this stage.

This scaling being done, for each case of ambulance chasing we measured the gap between the empirical data and each of the Nobel game simulations of $SA1$ to $SA4$. This gap is determined in terms of the average relative deviation of empirical data to simulated data. A gap of 4% means that on average, the number n_s^t of publications in the simulation at period t was in the interval $[0.96\ n_e^t\ ; 1.04\ n_e^t]$, where n_e^t is the number of papers published in arxiv.org measured in the number of days from the time of publication of the initial result. The best fit for each empirical case study is plotted in figure 6. As can be seen, the Nobel game dynamics fits these empirical data quite well, with an average gap below 5% for most studies ; which should be within the order of magnitude of the error in data collection. This is despite the low sophistication of the fit procedure ; no search for the best fit, only among simulated data in $SA1$ to $SA4$. Although there is a risk of over-fitting the data in the absence of additional data in the community to constrain the model, we can claim that at least empirical data do not contradict Nobel game dynamics.

An interesting question, one left for future investigations, would be to know whether a fit of this empirical data, with state-of-the-art techniques, could make it possible to estimate the characteristics of a scientific domain in a way consistent with experts’ intuition, both for observable parameters (such as the number of scholars, the duration of an experiment, the Gini index of publication distribution, the payoff policies, etc.), and also for non-observable parameters (*e.g.* degree of achievement of the discovery process, size, *abundance* and *separation* of the field, etc.). Another question would be to give a prediction, in probability, of the future development of a field.

13. cf. S8 and supplementary references for the analysis of these *ambulance chasing* event

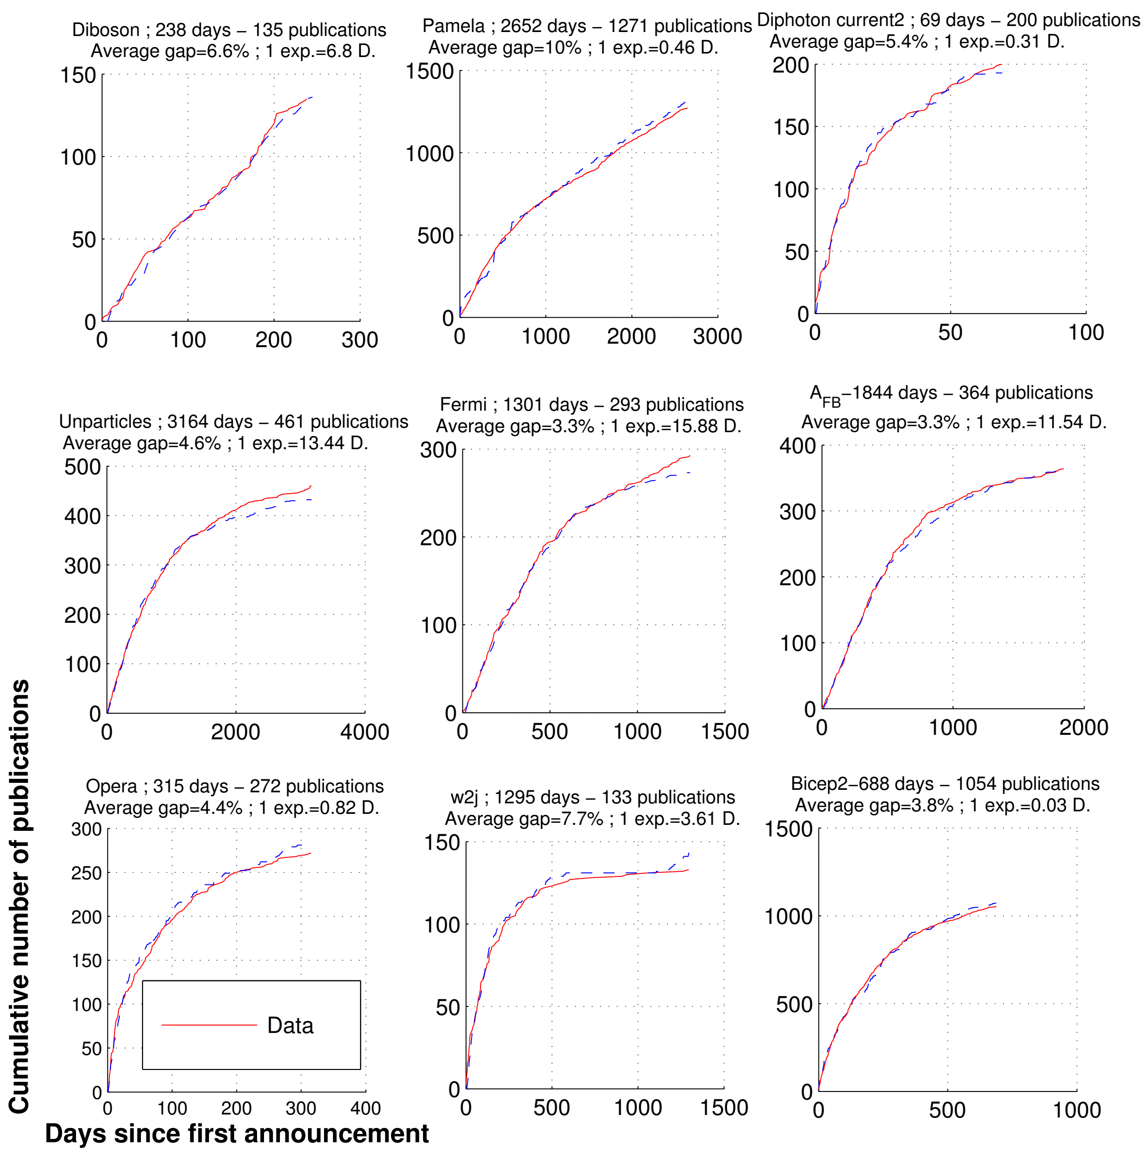

FIGURE 6 – ***Ambulance chasing* and Nobel game dynamics.** Comparison between publication dynamics of few particle physics fields with some Nobel games publication dynamics. Data (number of publications per day) have were collected by [Backović(2016)] from Arxiv.org. They concern 9 cases of *ambulance chasing* in particle physics. Simulated data correspond to the best fit in $SA1$ to $SA2$. More details about these cases are given in S8.

6 Conclusions

The *Nobel game* is a generic model for thinking about the effects of scientific policies on science dynamics. The analysis of its dynamics in the framework of agent-based simulations reveals the coupling between micro-dynamics and macro-dynamics in science through the scholars' exploration-exploitation dilemma, which translates at the level of the collective discovery process into a speed vs. quality dilemma.

In a context where more and more scientists insist on the importance of reproducibility in science, the Nobel game reveals an inescapable trade-off between the speed and the quality of the discovery process, and it highlights the importance of the social value given to refutation for sustaining science quality.

We analyzed in detail the effects of the "*publish or perish*" policy, which turns out to have detrimental impacts on the quality of the production of mature fields (*i.e.* for $A > 0.1$). On emerging fields, its influence depends on the characteristics of the field (*e.g.* separation, abundance, size) and varies from no influence ($SA3$) to weak influence in terms of speed or quality.

This observation, along with the one that the Nobel game reproduces the highly skewed distribution of the number of publications per scholar, without any assumption of scholar heterogeneity in terms of skill, also poses the question of what is really measured when science and scholar outputs are measured in terms of numbers of publications and citations.

Many other questions can now be addressed by computational studies. For instance the incidence of the network topology underlying scholars' interactions, the difficulty of the field, the effects of the community size and timing of scientific policies according to the maturity of the field can be studied.

For clarity, we tried to keep the model as simple as possible, and this means it has the limitations of the assumed simplifications. Many variations and adaptations are nevertheless possible, and through this article we hope to have triggered the curiosity of our readers for this kind of model.

The interest of the Nobel game, as a model for collective discovery processes, goes far beyond academia. Indeed, the activity of the scientific community can be taken as a prototypical example for a wide scope of distributed work situations, from knowledge elaboration to artifacts building. Well-known examples are collective development of software (ex. Linux [Raymond(2001)]) and content elaboration on the Internet (blogs, Wikipedia), newspaper publications, knowledge management, etc. Each of these knowledge spheres might have its own rules, timescale and specific dynamics, not to mention the varying nature of the publication, testing and refutation processes.

For example, there is a well-known journalistic adage from Pierre Lazaref, according to which one information plus one denial gives two pieces of information. When self-refutation is rewarding, shall we expect strange dynamics ? Exotic dynamics also exist in academia, such as in sociology where criticizing someone makes the critic's theories better known, and is a form of reward. Thus, some members of this community are reluctant to criticize people they don't agree with, to avoid giving them credit.

Beside presenting in detail the behavior of this model, future work will focus on quantifying all these effects by large-scale experimental and empirical approaches. This should lead to corroboration or refutation of our model, and a finer tuning of its parameters and hypotheses.

acknowledgements

This material is based upon work supported by the Complex Systems Institute of Paris Île-de-France (ISC-PIF) and the *Science en poche* project funded by the city of Paris *Emergence(s)* funding

scheme. The author is very grateful to Mihailo Backović for sharing his data about *ambulance chasing* cases.

Références

[Allanach(2014)] Ben Allanach. Ambulance-chasing Large Hadron Collider collisions | Ben Allanach | Life & Physics. https://www.theguardian.com/science/life-and-physics/2014/sep/17/ambulance-chasing-large-hadron-collider-collisions, 2014.

[Aumann(1997)] Robert J. Aumann. Rationality and Bounded Rationality. In Sergiu Hart and Andreu Mas-Colell, editors, *Cooperation : Game-Theoretic Approaches*, number 155 in NATO ASI Series, pages 219–231. Springer Berlin Heidelberg, 1997. ISBN 978-3-642-64413-9 978-3-642-60454-6. doi : 10.1007/978-3-642-60454-6_15. 00278.

[Backović(2016)] Mihailo Backović. A Theory of Ambulance Chasing. *arXiv :1603.01204 [hep-ph, physics :physics]*, March 2016.

[Borner et al.(2004)Borner, Maru, and Goldstone] K. Borner, J. T. Maru, and R. L. Goldstone. The simultaneous evolution of author and paper networks. *Proceedings of the National Academy of Sciences*, 101(Supplement 1) :5266–5273, April 2004. ISSN 0027-8424, 1091-6490. doi : 10.1073/pnas.0307625100.

[Calude and Longo(2016)] Cristian S. Calude and Giuseppe Longo. The Deluge of Spurious Correlations in Big Data. *Foundations of Science*, pages 1–18, March 2016. ISSN 1233-1821, 1572-8471. doi : 10.1007/s10699-016-9489-4.

[Chavalarias(2006)] David Chavalarias. Metamimetic Games : Modeling Metadynamics in Social Cognition. *Journal of Artificial Societies and Social Simulation*, 9(2) :5, 2006. ISSN 1460-7425.

[Chavalarias and Gardelle(2008)] David Chavalarias and Vincent Gardelle. Social simulation of collective a discovery process : the Nobel Game. In *5th conference of the European Social Simulation Association*, Brescia, Italy, 2008.

[Chavalarias et al.(2006)Chavalarias, Charron, Gardelle, and Bourgine] David Chavalarias, Sylvain Charron, Vincent Gardelle, and Paul Bourgine. NOBEL, Le jeu de la découverte scientifique, Une approche analytique, experimentale et computationnelle. In *conférence Modélisation, Optimisation et Simulation des Systèmes, Défis et Opportunités (MOSIM'06)*, Rabbat, 2006.

[Chavalarias et al.(2014)Chavalarias, Leenart, and Panahi] David Chavalarias, Jean-Baptiste Leenart, and Maziyar Panahi. Publish or Perish - Is that all ? What is the social game behind science ? In *European Conférence on Complex Systems*, Lucca, Italy, 2014.

[Chavalarias, D.(1998)] Chavalarias, D. La thèse de Popper est-elle réfutable. Master thesis, Ecole Polytechnique, Paris, 1998.

[Cohen et al.(2007)Cohen, McClure, and Yu] Jonathan D. Cohen, Samuel M. McClure, and Angela J. Yu. Should I stay or should I go ? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B : Biological Sciences*, 362(1481) :933–942, May 2007. ISSN 0962-8436, 1471-2970. doi : 10.1098/rstb.2007.2098. 00426.

[Conte and Paolucci(2001)] Rosaria Conte and Mario Paolucci. Intelligent Social Learning. http://jasss.soc.surrey.ac.uk/4/1/3.html, 2001.

[Edmond, B. et al.(2011)Edmond, B., Gilbert, N., Ahrweiler, P., and Scharnhorst, A.] Edmond, B., Gilbert, N., Ahrweiler, P., and Scharnhorst, A. Simulating the Social Processes of Science. *JASSS*, 2011.

[Edmonds(2008)] Bruce Edmonds. Artificial science : A simulation to study the Social Processes of Science. In Bruce Edmonds, Cesareo Hernandez, and Klaus G. Troitzsch, editors, *Social Simulation - Technologies, advances, and new discoveries*, pages 61–67. Hershey - New York, 2008.

[Freedman et al.(2015)Freedman, Cockburn, and Simcoe] Leonard P. Freedman, Iain M. Cockburn, and Timothy S. Simcoe. The Economics of Reproducibility in Preclinical Research. *PLOS Biology*, 13(6) :e1002165, June 2015. ISSN 1545-7885. doi : 10.1371/journal.pbio.1002165. 00053.

[Gilbert, N.(1997)] Gilbert, N. Gilbert : A Simulation of the Structure of Academic Science. *Sociological Research Online*, 1997.

[Harzing(2007)] A.W. Harzing. Publish or Perish, 2007.

[He and Zhang(2009)] Xionglei He and Jianzhi Zhang. On the Growth of Scientific Knowledge : Yeast Biology as a Case Study. *PLoS Computational Biology*, 5(3) :e1000320, March 2009. ISSN 1553-7358. doi : 10.1371/journal.pcbi.1000320.

[Ioannidis(2005)] John P. A. Ioannidis. Why Most Published Research Findings Are False. *PLoS Medicine*, 2(8) :e124, 2005. ISSN 1549-1676. doi : 10.1371/journal.pmed.0020124.

[Larousserie and Morin(2015)] David Larousserie and Hervé Morin. Olivier Voinnet, star de la biologie végétale, sanctionné par le CNRS. *Le Monde.fr*, July 2015. ISSN 1950-6244. 00000.

[Newman(2001)] M. E. J. Newman. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2) :404–409, January 2001. ISSN 0027-8424, 1091-6490. doi : 10.1073/pnas.98.2.404.

[Open Science Collaboration(2015)] Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*, 349(6251) :aac4716–aac4716, 2015. ISSN 0036-8075, 1095-9203. doi : 10.1126/science.aac4716.

[Payette(2012)] Nicolas Payette. Agent-Based Models of Science. In Andrea Scharnhorst, Katy Börner, and Peter van den Besselaar, editors, *Models of Science Dynamics*, pages 127–157. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-23067-7 978-3-642-23068-4.

[Pentland(2012)] Alex Pentland. Big Data's Biggest Obstacles. https://hbr.org/2012/10/big-datas-biggest-obstacles, 2012.

[Popper(1962)] Karl R. Popper. *Conjectures and refutations : the growth of scientific knowledge*. Basic Books, 1962. ISBN 978-0-415-28593-3 978-0-415-28594-0. 00030.

[Popper(2002)] Karl R Popper. *The logic of scientific discovery*. Routledge, London ; New York, 2002. ISBN 978-0-203-99462-7 978-0-415-27843-0 978-0-415-27844-7.

[Rabin(2002)] Matthew Rabin. A perspective on psychology and economics. *European Economic Review*, 46(4-5) :657–685, May 2002. ISSN 00142921. doi : 10.1016/S0014-2921(01)00207-0. 00586.

[Raymond(2001)] Eric S. Raymond. *The Cathedral & the Bazaar : Musings on Linux and Open Source by an Accidental Revolutionary*. O'Reilly Media, Beijing ; Cambridge, Mass, 1 edition edition, January 2001. ISBN 978-0-596-00108-7.

[Scharnhorst et al.(2012)Scharnhorst, Börner, and van den Besselaar] Andrea Scharnhorst, Katy Börner, and Peter van den Besselaar, editors. *Models of Science Dynamics*. Understanding Complex Systems. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-23067-7 978-3-642-23068-4.

[Schelling(1978)] Thomas C. Schelling. *Micromotives and macrobehavior*. Fels lectures on public policy analysis. Norton, New York, 1st ed edition, 1978. ISBN 978-0-393-05701-0 978-0-393-09009-3.

[Sun and Naveh(2009)] Ron Sun and Isaac Naveh. Cognitive simulation of academic science. pages 3011–3017. IEEE, June 2009. ISBN 978-1-4244-3548-7. doi : 10.1109/IJCNN.2009.5178638.

[Watts and Strogatz(1998)] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684) :440–442, June 1998. ISSN 0028-0836. doi : 10.1038/30918. 28244.

[Weisberg and Muldoon(2009)] Michael Weisberg and Ryan Muldoon. Epistemic Landscapes and the Division of Cognitive Labor. *Philosophy of Science*, 76(2) :225–252, April 2009. ISSN 0031-8248, 1539-767X. doi : 10.1086/644786.

[Zollman(2007)] Kevin J. S. Zollman. The Communication Structure of Epistemic Communities. *Philosophy of Science*, 74(5) :574–587, December 2007. ISSN 0031-8248, 1539-767X. doi : 10.1086/525605.