

Explaining Scientific Collaboration: a General Functional Account

Thomas Boyer-Kassem* and Cyrille Imbert†

(Working paper – v3 – October 16, 2020)

Abstract

For two centuries, collaborative research has kept developing, which has been explained in various ways. We offer a novel functional explanation of this development, grounded in a sequential model of scientific research where the priority rule applies. Robust patterns about the differential successfulness of collaborative groups over their competitors are derived and it is argued that they feed the development of collaboration. This global mechanism may trigger an arms race and is compatible with some decrease of productivity of collaborative groups and some overcollaboration. The proposed explanation can integrate various factors usually associated with the rise of collaboration.

*Département de Philosophie, MAPP, Université de Poitiers, France.
thomas.boyer.kassem@univ-poitiers.fr

†CNRS, Archives Poincaré, France. cyrille.imbert@univ-lorraine.fr. The authors contributed equally to the paper.

1 Introduction

Scientific collaboration has kept developing for two centuries, especially since the 1950s, but inhomogeneously so across fields. Various epistemic or non-epistemic explanations of this collaborative trend have been proposed, such as professionalization, scientific specialization, desire to cross-fertilize research, changing patterns of fundings, better reliability of research, enhanced efficiency, or the need to access resources. As noted by Katz et al. (1997, 4), this list of possible explanatory factors, which are often conceived of as distinct and exclusive, is potentially endless.

By contrast, we present a general explanatory scheme that encompasses many such factors in a complementary way. This scheme is functional: it explains a feature by the function it serves, namely the beneficialness of one of its effects (in some technical sense to be presented). More specifically, we claim that the increase of scientists' collaborative behavior can be functionally explained by the increase in the individual successfulness of collaborative scientists over their competitors. That is, scientists collaborate because, *in suitable contexts*, they are more successful when collaborating (enough) than when not (enough).

Wray (2002) has already offered a functional explanation of collaboration. The novelty of our thesis lies in the social and epistemic mechanisms that ground the beneficialness of collaboration, and feed the functional explanation. We investigate these mechanisms through a model of scientific research introduced by Boyer-Kassem and Imbert (2015) that highlights the potentially significant beneficialness of collaboration in suitable contexts. Scientific inquiries are pictured as multi-step processes, which requires solving sequence of intermediate problems having communicable results. Even minor differences in the step-efficiency of collaborative groups are shown to make them more successful in competitive contexts. Robust patterns about the differential successfulness of collaborative groups over their competitors, in suitable contexts, are derived and it is argued that they feed the development of collaboration in science. Because of the priority rule, these mechanisms are compatible with a decrease of productivity of collaborative groups and some overcollaboration. All factors that increase the step-efficiency of groups can be integrated in this general explanatory scheme.

The paper is organized as follows. We review the literature about collaboration and discuss the application of functional explanations in the present case (sections 2-3). We introduce and investigate the model, and argue that it can feed a functional explanation of collaboration (sections 4-6). We further discuss modeling assumptions, analyze the scope of this explanation, and finally highlight limits of this inquiry (sections 7-9).

2 The development of collaboration in contemporary science

Collaboration is a complex multi-faceted phenomenon, and one may reasonably expect that no simple explanation covers all of its aspects. We review here various empirical findings about collaboration, and highlight the ones we target in the present paper.

For two centuries, co-authoring papers has become increasingly widespread in academia (Price, 1963, Beaver and Rosen, 1978). This increase of co-authoring is commonly taken as evidence of the increase of scientific collaboration, even if co-authoring and collaboration are not synonymous (e.g. co-authors with minor scientific contributions can be mentioned and genuine contributors omitted (Katz, 1997)). Importantly, the degree of collaboration varies significantly across fields. Wuchty et al. (2007) report that between the mid-50ies and 2000 the percentage of collaborative team grew from 50% to 80% in science and engineering, from 17.5% to 51% in the social science, but remained stuck to 10% in the humanities. Mean team size has also steadily grown in collaborative fields, e.g. from 1.9 to 3.5 in science.

These general trends supervene on different collaborative behaviors. For example, experimentalists collaborate more than theoreticians. Women have a higher percentage of female collaborators (Bozeman and Corley, 2004), a lower probability of repeating co-authors, and fewer co-authors over their career (Zeng et al. 2016). Researchers with large grants tend to have more collaborators (Bozeman and Corley, 2004). Overall, most people tend to collaborate in their immediate environment (Lee and Bozeman, 2005, 695). Individually, various collaborative profiles exist, e.g. “taskmasters”, “nationalists”, “mentors”, or “buddies” (Bozeman and Corley, 2004, 612; see also Li et al. 2013).

Policy-makers and institutions generally assume that collaboration is beneficial and should be encouraged. Indeed, Nobel prizes are known to collaborate more than average (Zuckerman, 1967), results by large teams are more cited (Wuchty et al., 2007), and the number of collaborators is strongly associated with that of publications. However, if one applies a fractional count, the increase of productivity is not that clear (Lee and Bozeman, 2005). Overall, how much collaboration is globally beneficial remains uncertain since the costs and benefits are hardly appraised (Katz 1997).

How can the development of collaboration be explained? A mere “rational” explanation of collaboration, based on its supposed epistemic or non epistemic beneficialness, cannot be taken for granted and is potentially unsuitable to account for the heterogeneity of these patterns. Simple explanations based on the beliefs of agents are in no better position. When asked why they collaborate, scientists provide many various reasons such as gaining knowledge,

accessing equipments, working necessity, or friendship relations (Melin, 2000). The development of big science and funding, or changes in norms of authorship, probably play a role, but wildly differ across fields (Wuchty et al., 2007, 1037–1038). Still, the pervasiveness of the increase of the percentage and size of collaborative means — to the exception of the humanities — strongly suggest that the conditions favoring team work are similar across subfields (ibidem).

In the remainder of the paper, we develop a general explanatory schema that can account for the general increase of the percentage of collaborative works and size of collaborative groups. It involves factors that are similar across fields but mostly absent from the humanities where collaboration is hardly observed.

3 Functional Explanations and Collaboration

We now present functional explanations and how to apply them here. Functional explanations explain the existence of a feature by one of its effects, typically its beneficialness. For example one can suggest that the State (resp. the practice of golfing) exists because it serves the interest of the ruling class (resp. enables upper classes to establish network). Obviously, pointing at a useful effect of a feature is not enough to explain it functionally — the usefulness of the nose to carry glasses cannot explain why humans have one. Nevertheless, if stringent conditions are met, functional explanations can be satisfactory, typically within biology and the social sciences.¹

How to characterize suitable functional explanations and get rid of spurious ones? Like Wray (2002), we follow Kincaid’s simple and widely accepted account. For Kincaid (1996, 110-111; 2006, 214-215), A is functionally explained by B , or A exists “in order to B ” if:

- (i) A causes B ;
- (ii) A persists because of (i), i.e. because it causes B ;
- (iii) A is causally prior to B .

As Kincaid notes (2006, 214), “The first claim is straightforwardly causal. The second can be construed so as well. [...] A ’s causing B causes A ’s continued existence.” Finally, the third condition says that the causal loop is initiated by A but not by B . For example, initiation rites bring about social solidarity, and therefore persist, but the converse is false.

To provide a functional explanation of scientific collaboration, we consider the credit corresponding for authors to the publication of papers, divided by

¹Even Elster (1983), who favors methodological individualism, considers functional explanations as acceptable in the social sciences.

the time spent to find these results. This quantity, which Boyer-Kassem and Imbert (2015, p. 672) call *individual successfulness*, depends both on the actual production of results by scientists and their ability to publish them first, and may vary in different contexts and for groups of various sizes (see below for details). It corresponds to the publishing productivity of scientists over their projects, even if the literature often uses this term to describe the sheer number of publications of authors.

Claim. The increase of scientists' collaborative behavior can be functionally explained by the increase of individual successfulness it brings; in other words, scientists collaborate "in order to" increase their individual successfulness.

Endorsing Kincaid's analysis, this claim amounts to:

- (i-c) in suitable contexts, the increase of scientists' collaborative behavior causes an increase of their individual successfulness, up to some point;
- (ii-c) scientists' collaborative behavior persists (and develops) because of (i-c), i.e. because it gives them a higher individual successfulness;
- (iii-c) the increase of scientists' collaborative behavior is causally prior to this increased individual successfulness.

As noted by Kincaid (1996, 115), the satisfaction of the third clause is often trivial, and our case is no exception (Wray 2002, 161). Many scientists have been successful without collaborating. Also, collaboration is recent while successfulness is not, so if successfulness was the initial cause of collaboration, one would need to explain why for so long it did not generate collaboration. So, we consider that (iii-c) is established. Claims (i-c) and (ii-c) will be argued for in sections 5 and 6. Their respective defenses are distinct issues. The central contribution of this paper is to provide strong evidence in favor of (i-c) (in section 5) thanks to a model of scientific collaboration, introduced in section 4.

4 The Sequential Competitive Model

In the model introduced by Boyer-Kassem and Imbert (2015), n competitors struggle over the completion of a research project composed of l sequential steps, in which only the last step is publishable. For instance, one of the steps corresponds to the design of an experiment, another one to its running, another one to the statistical analysis of the results, and finally the writing of the paper. What constitutes a step depends on the level of analysis. Here, l is fixed to 10

and one step corresponds to a task that would take between two weeks and a month.

Time is discrete, and at each time interval, agents have independent probabilities p of passing a step. If p is close to 1, passing steps is easy, if p is close to 0, it is hard. In forthcoming illustrations, p is set to 0.5^2 , which means that completing a project takes on average between 10 and 20 months for a single scientist. Agents can either stay on their own or gather into collaborative groups for the whole project. For instance, a community of 3 agents working on the same research problem can gather in a group of 3, or in a group of 2 and a loner, or just stay all on their own. These collaboration configurations are respectively noted (3), (2-1) and (1-1-1).

Importantly, the number n of competitors is not meant to represent the much larger number of researchers who belong to a scientific domain, such as high energy physics or social psychology, but stands for the subset of researchers who work on the very same problem and might collaborate. Further, an agent can be interpreted as an individual researcher, but also as a team. Thus, the model can also represent how collaboration develops at an institutional level. For example, $n = 5$ can stand for 5 individuals or 5 labs comprising 10 researchers, which may or may not decide to work together.

Is it reasonable to suppose that several individuals (or teams) compete over an identical problem? The number of scientific problems is potentially quite large. But in the framework of paradigms, research programs, or organized research communities, which scientific puzzles and questions are significant and which questions can plausibly be solved is common knowledge (Kuhn, 1962, Kitcher 1993). Overall, the assumptions that several individuals or teams often investigate identical problems remains plausible³. One might think that the model is only applicable to cases where several groups are in a credit race. Actually, the case where a collaborative group of n researchers works alone (“in a community of n ”) on a problem is also covered, and the model can be used to compare this case with *hypothetical* situations in which these n individuals would have worked in small groups or alone. So the scope of the model, and of our forthcoming explanation, is general and covers cases where collaborative groups do not have competitors.

Because of the priority rule, the first agent or group to reach the last step publishes it and wins all the scientific credit, while others get nothing. So all collaborative groups and loners compete to be the quickest. Within a collaborative group, the credit for a published paper (fixed to 1 for each race) is

²Evidence suggests that the main conclusions don't depend on this choice. See Boyer-Kassem and Imbert 2015, p. 674 and 684 for more details.

³Recent papers in the field of social epistemology of science make a similar assumption. For example, Zollman (2010), and Frey and Šešelja (2019) investigate a community of no more than a dozen of distinct researchers who assess experimentally the compared efficiency of two drugs.

assumed to be equally split among agents. Thus, the individual successfulness of an agent corresponds to her share of credit, divided by the time needed to complete the inquiry, and this quantity is assessed by averaging over millions of replicas.

An important choice bears on the value of the probability that a group of size k passes a research step, which is $p_g(k, p) = 1 - (1 - p)^k$. The initial justification by Boyer-Kassem and Imbert (2015) is that it represents a situation in which collaborative agents have independent probabilities p of passing a step, and share information at no cost: when an agent has passed a step, her collaborators pass it with her for free, and then all start trying to pass the next step. This interpretation is useful to grasp the model, but is not endorsed in this paper. Here, the mathematical quantity of $p_g(k, p)$ represents an advantage, though a relatively minor one, for collaborative groups. We discuss at more length this modeling hypothesis in section 7.

Because claims (i-c) and (ii-c) bear on individual successfulness (see section 3), we now review existing results about it. Boyer-Kassem and Imbert (2015) have computed the individual successfulness for all collaboration configurations up to a community of $n = 10$ agents. Their main finding is that minor differences in the probability of passing steps can be much amplified and that, even with not-so-favorable hypotheses, collaboration can be extremely beneficial for scientists. For example, in a community of 9 agents organized in the configuration (5-4), i.e. with a collaborative group of 5 and another one of 4, the group of 5 only has a 3% higher probability of passing a step compared to the group of 4, but the individual successfulness for an agent in the group of 5 is 25% higher. As another example, in a community of 3 agents in the configuration (2-1), whereas the difference in the probability of passing a step is 50% higher for the group of 2, the individual successfulness is 700% higher.

However, these results do not explain scientific collaboration by themselves. First, the beneficialness of a feature does not by itself explain its existence. Second, these results merely show that collaborating is beneficial for particular groups in particular collaborative configurations — for example, a group of 2 is very successful in configuration (2-1-1-1-1) but not in (7-2). Thus, the model merely provides results about what can be the case in possible configurations. But the explanandum is a general social feature of contemporary science, not some collaborative behavior in some particular case, so the explanans must also involve general statements about the link between collaboration and beneficialness. This is what we argue for in the next sections.

5 Does Collaboration Cause Successfulness?

In this section, we use the above model to provide strong evidence that the increase of collaboration causes an increase of individual productivity in

suitable contexts (i-c) in a sense that must be carefully specified. As will be argued, this is not incompatible with a global decrease of the productivity of the community. We first provide an analytical result that covers various cases and gives insight into the model, and then explore the model by a brute-force method.

A first general result is the following proposition (proof in the appendix):

Proposition. If m groups each comprising k agents merge, the individual successfulness of these agents increases.

That is, as soon as several collaborative groups of the same size exist, they improve the individual successfulness of their members by merging. A corollary is that single individuals always have interest in collaborating together, compared to staying on their own. A result of this type, though already general, merely provides a partial order over the successfulness of groups of size k across configurations, and cannot provide a general vindication for (i-c). However, its informative value should not be underestimated. It holds for any size and configuration of the rest of the community. For example, in configuration (3-2-2-1-1-1), the loners would get a larger individual successfulness if they merged into a group of 2 or of 3, and the members of the groups of 2 if they merged into a group of 4. The proposition also implies that individuals in a group of even size k have no interest in splitting into two groups of $k/2$. Also, the proof of the proposition gives an insight about why merging can be beneficial: new larger groups keep the gains of the subgroups from which they come, and in addition develop new gains either by saving time (if one of their subgroups would have finished first anyway, they are now quicker because they share progress) or by earning new gains (by finishing before their competitors in new cases because they now proceed quicker).

Another route to investigate the general link between collaboration and success is to analyze what happens *on average over all relevant configurations*, which we do in the remainder of this section. For all 138 possible configurations up to $n = 10$ competitors, we compute the average individual successfulness in all groups.⁴ For each configuration, the successfulness is computed by running many times the race between the different groups and assessing how frequently each group wins. These quantities are computed up to the point

⁴Technically, we compute the average individual successfulness of an agent in a collaborative group of k over all possible collaboration configurations meeting some constraint. For example, the average individual successfulness when in a group of 4 and in a community of 7 is computed as the average of individual successfulness in configurations (4-1-1-1), (4-2-1) and (4-3). Which weight should be attributed to each configuration? The number of combinatorial realizations of these configurations is not necessarily what matters, and the right weights ultimately depend on the empirical frequency of actual configurations. For lack of information about these frequencies, we privilege simplicity and give equal weights to all configurations.

where fluctuations become marginal, which takes between hundreds of thousands and millions of runs. Then, in order to establish the robustness of the causal claim, we study the influence of various parameters such as the total number of competitors, the collaborative group size, and the average level of collaboration in the community.

Successfulness and community size. Figure 1 shows the average individual successfulness as a function of community size n , for agents in collaborative groups of various sizes. A first observation is that the successfulness of loners quickly collapses with n and is much lower than that of other groups as soon as $n > 2$. This shows that loners are outraced, except in very small communities. Second, for all group sizes, individual successfulness decreases with the community size, which is expected since the number of competing groups and their sizes increase. Nevertheless, for each k , the successfulness of a group of size k starts high and does not decrease much, up to some community size s larger than k . In still larger communities, they are eventually outperformed by larger groups. Note that, for $n > 2$, the beginning of a curve of a group of size k is below a curve of some smaller group: collaborating maximally is not the best strategy.⁵ Third, when groups are larger, this initial plateau of successfulness is longer and flatter, and the subsequent decrease is less steep. Fourth, when n is much larger than k , the successfulness of an individual increases with the size of the group she is in. However, this increase is moderate and small groups still do reasonably well (this is somewhat unexpected given the general amplification mechanism, but see Figure 3 below for more refined analyses). For instance, in a community of 10 researchers, groups of 2 are suboptimal but remain somewhat viable since their average successfulness remains between 1/3 to 1/2 of that of groups of 3 or 4. Overall, the morale is that not collaborating is in general not a viable strategy. Collaborating moderately ($k = 2$ or 3) can be very rewarding when there are few competitors, for instance on ground-breaking questions that are only tackled by a handful of scientists. However, when communities of competitors become significantly larger, small groups tend to be outraced. Thus, not collaborating or collaborating moderately can be a risky strategy when uncertainty prevails about other competing groups. These findings suggest that normal science (when important problems are clearly identified within communities) is best developed by large teams, but achievements and success by individuals or small groups are still possible on ground-breaking issues⁶. Finally, being in large collaborative groups does not give very high successfulness but is a reasonably safe situation. Overall, these results support the claim that up to some point, collaborating

⁵See Boyer-Kassem and Imbert (2015, 679-80) for an analysis of over-collaboration in large groups.

⁶This is particularly in line with recent findings by Wu et al. (2019) that large teams develop science while small teams disrupt it.

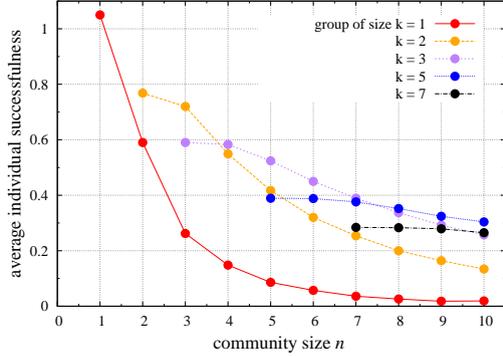


Figure 1: Variation of individual successfulness with community size.

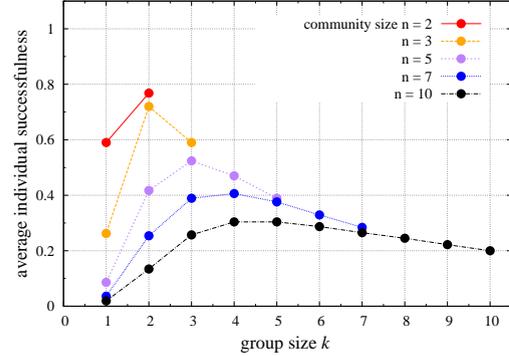


Figure 2: Variation of individual successfulness with the size of groups.

more causes more successfulness, and more robustly so.

Successfulness and group size. Figure 2 shows for various community sizes how individual successfulness changes with the size of collaborative group. For $n > 2$, the successfulness curve has a one-peaked (discrete) form, which means that there is an optimal collaborative group size, different from the community size, over which adding new collaborators decreases successfulness. Note that a single-peaked preference function is usually *assumed* in the literature about coalitions, or in social choice more generally. Here it emerges from a micro-model, without an *ad hoc* mechanism for overcollaboration. This shows that the model remains reasonable even if the costs of collaboration are idealized away. In particular, in a given context, groups have no incentive to grow indefinitely larger.

Note also that the decrease after the peak is slower than the increase before it: over-collaborating is *globally* less harmful and risky than under-collaborating. For stable large networks that collaborate a lot, adding or removing some collaborators does not substantially change their (high) successfulness. So, the successfulness of these groups is robust. Also, the position of the maximum grows with the community size, which suggests that in larger, more international communities, optimal collaborative groups are larger.

Successfulness in more or less collaborative communities. Figure 3 displays how the successfulness of members of collaborative groups of size k varies with the mean group size in the community, i.e. with the degree of collaboration in the group environment⁷. Consider how the various curves stand relatively

⁷For each collaboration configuration, we compute the mean group size, for instance 1.75 for collaboration configuration (3-2-1-1). Then, for each group size k , we compute the average successfulness over configurations having a degree of collaboration within intervals $[1, 1.5]$ (represented at coordinate “1.25” on the x -axis), $[1.25, 1.75]$ (represented at coordinate “1.5”), $[1.5, 2]$... $[3.5, 4]$.

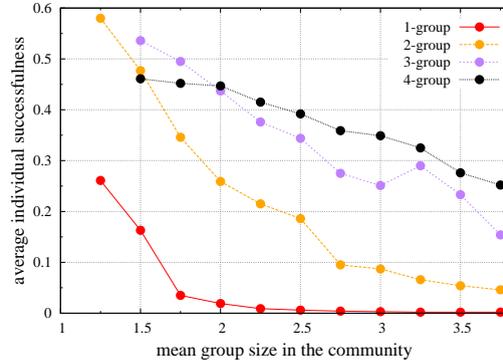


Figure 3: Variation of successfulness with the degree of collaboration in communities.

to one another. Whatever the mean group size of the community, loners are always outstripped by collaborative groups. Individuals in groups of size 2 get a much lower successfulness (compared to larger groups) as soon as the mean group size is larger than 1.5. Individuals in groups of size 3 have lower successfulness than those in groups of 4 as soon as the mean group size is larger than 2 or so. Overall, this shows that successfulness depends less on the absolute size of the group than on how the relative size of the group compares to the mean group size of competitors. The main point is that scientists who collaborate more than average are very successful; those who collaborate as their peers do reasonably well; those who collaborate less than average are outraced by a large margin. This general result is not unexpected given previous results, but the graph highlights that the success for intensively collaborating scientists, and underachievement for under-collaborators, can be very large. This shows that collaborating clearly causes successfulness, compared to not collaborating at all. Moreover, collaborating more than average causes more successfulness too. Resources do not simply accrue to successful scientists but to more successful ones, and the reward structure favors the degree of collaboration going upward, creating an unintentional arms race of sorts.

The above patterns have been obtained for competing communities up to $n = 10$, but remember from section 4 that n stands for the number of scientists who compete on the very same problem (not the whole scientific community) and that agents can further be interpreted as teams or institutions. Hence, our results cover a large number of real cases.

Let us sum up the findings of this section. Within a competing community, when collaboration is possible, it entails successfulness. This relation is robust under changes in the size of communities or in the exact size of groups. Also, those who collaborate more than average are more successful; collaborating too

much is slightly problematic, while under-collaborating is strongly so, which makes collaborating a lot a safe working habit, especially in the absence of information about competitors. All this makes claim (i-c) adequately supported: in suitable contexts, when scientists collaborate, this causes an increase of their individual successfulness, up to some point.

These results are compatible with the empirical finding that when adopting a fractional count (dividing the number of publications by the number of authors, like here), the number of collaborators does not clearly significantly boosts publishing productivity (Lee and Bozeman, 2005), even if it increases the number of journal papers of collaborating authors. Indeed, we do not claim *simpliciter* that the productivity of scientists is positively correlated with the size of their group. Our claim is relative to a given community: productivity can be higher for groups of size 3 surrounded by groups of average size 1.5, than for a group of size 4 surrounded by groups of average size 2.5 (see Fig. 3). In brief, if they remain in an ecological niche in which collaborative is low, individuals and small groups can do better than larger groups.

6 Collaborative Behaviors Develop Because of the Differential Successfulness of Collaborative Scientists

We now argue that, because of the differential successfulness of collaborative scientists, collaborative habits persist and develop in scientific communities (ii-c). This feedback loop can be caused by a variety of well-attested social mechanisms across scientific contexts, and we review in this section various evidence about these mechanisms.

Transmission. Knowing when, with whom, and how to collaborate is not straightforward. Like other know-how skills, it is usually developed by exercising it with people who already possess it. People who already collaborate endorse the role of cultural transmission for colleagues (Thagard 2006), and above all for students, as collaborating with them is an efficient way to train them (Thagard 1997, 248–50). Then, the cultural transmission of collaborative practice does not require any particular effort on top of that⁸. The very circumstances that make collaboration possible and beneficial also make its transmission easier: when a research project can be divided into well-defined tasks, the solutions of which can be publicly assessed and shared, it is easier to enroll other people and thereby transmit collaborative skills to them (*ibidem*). Thus, collaborative habits and skills can be passed over and need not be reinvented by newcomers.

⁸The fact that students reproduce the practices of their supervisors is a key ingredient of models analyzing selection mechanisms in science, see e.g. Smaldino et al., 2016 or O'Connor, 2019.

Transmission opportunities. We now give evidence that collaborative scientists, because they are more successful, are more often in a position to transmit their collaborative habits, hence the persistence and development of collaboration.

Within pure science, because the extension of certified knowledge is the official goal of science, successful scientists can be expected to stand better chances to get good positions and grants, develop research programs, and pass over their collaborative habits. And although scientific institutions are imperfect, it seems that this is actually often the case. Within applied science, in which collaboration is also widespread (Wuchty 2007), research projects are usually directed at finding profitable and patentable applications. Thus, fund providers are strongly interested in hiring and providing resource to successful scientists. Note that the link between average success and pragmatic rewards is needed on average only, and it may be that *some* epistemically successful scientists get little resource while *some* unsuccessful scientists get a lot.

In practice, non-epistemic factors may even tend to over-credit successful scientists, and in particular collaborative ones. First, individual success has been assessed in the model with a conservative estimate, and an agent's publication within a group of size k is often assessed more than just $1/k$ a single-authored publication. For instance, a large French research institution in medicine officially weighs the citations of a paper with “a factor 1 for first or last author, 0.5 for second or next to last, and 0.25 for all others” (Inserm 2005). Second, a publication within a group of 10 is generally more visible than one single-authored publication, as more people can promote it. Third, sociology of science suggests that scientific credit tends to accrue to a subset of scientists who are perceived as extremely successful — the Matthew effect (Merton, 1968). Precisely, the model shows that collaborating more than average can be extremely rewarding. Then, to the extent that access to resources increases with scientific credit, successful collaborative scientists can be expected to benefit more from this effect.

Other types of mechanisms may contribute to this process, like conscious ones.⁹ Once winners of the scientific race publish co-authored articles, it becomes easy for others to see that successful scientists are highly collaborative. For instance, if agents in a group of 3 are 4 times more successful than a single agent, this means that their group publishes 12 more articles than this agent. Accordingly, the belief that collaborating is beneficial can even be acquired by non-collaborating scientists just by looking at journals. Furthermore, resources may accrue to scientific institutions that host individually successful scientists, and indirectly to these scientists. Agents in the model can be reinterpreted as teams or collective entities which decide to share results or to combine

⁹Kincaid mentions that “complex combinations of intentional action, unintended consequences of intentional action, and differential survival of social practices might likewise make these conditions [(i)–(iii) in our section 3] true” (Kincaid 1996, 112).

their expertise to produce collective articles. Then, these institutions and their members are more successful, attract resource, and keep developing and transmitting their working habits.

Overall, there is a wealth of evidence that scientists and scientific institutions which are successful (in terms of published results) have more opportunities to develop, to school research students, and to place them in the academic world¹⁰. Hence, the causal connection between the success of collaborative scientists and the persistence and development of collaborative practices (ii-c) is highly plausible.

7 Discussion: modeling assumptions

The model does not purport to represent faithfully all situations in which collaboration occurs. As we shall argue, this choice does not threaten our general argument.

First, the winner-takes-all scheme is not always strictly implemented, as identical or close findings may be published in less renowned venues. For this reason, we discuss at length in section 8 the case of the humanities. However, note that scientists usually get far less credit for such publications, and major journals explicitly require novel material. This legitimates investigating what takes place in the idealized situation in which the priority rule is strictly applied.

Second, the modeling choice concerning the step-efficiency of groups should be taken with some grain of salt. Actual collaboration takes place in various ways. All collaborative mechanisms are not equally efficient, rewarding, or beneficial to all types of members of a group (see e.g. Abbasi et al., 2011, Li et al., 2013, Parish et al., 2017). So it would be preposterous to assume that one single step-efficiency function can represent adequately all cases. But this is not what we do here. As we study how the effects of the priority rule can feed a functional explanation of collaboration, it is argumentatively better to assume that the epistemic benefits of collaboration are low. Accordingly, the step-efficiency of collaborating groups in the model is chosen on a pessimistic basis so that if collaboration is efficient in this case, it can be expected to be even more efficient in more realistic ones. In this way, our results highlight the successfulness and potential development of groups in cases their differential advantage is larger than (or equal to) the low key estimate within the model. The fact that the step-efficiency function is not favorable can be seen in two ways.

Firstly, remember it can be interpreted as the step-efficiency of a somewhat inefficient group of researchers that work independently and sequentially on the

¹⁰This feature is also used in formal models that investigate potential selection mechanisms in science, see again Smaldino et al. 2016 and O'Connor, 2019.

same subtasks, share intermediate results, and do not interact otherwise. But experimental research in social psychology has shown that collaborative groups are often better at solving problems than individuals, even if various pitfalls can spoil the epistemic activity of these groups. For highly intellectual tasks, which have demonstrably correct solutions, a group performs better than the best of its individuals (see Laughlin 2011 for a review). Thus, groups are faster at solving a succession of intellectual tasks. By contrast, in the model, at each step, a group does just as well as its most successful member, not better. Further, evidence shows that group performance increases with the number of high ability members within the group (*ibidem*, chapter 4). In collaborative groups, specific tasks can be ascribed to experts, i.e. high ability members. Accordingly, the probability to pass steps in such groups is higher than that chosen in the model, in which all agents have the same probability to pass steps. Finally, the model assumes that research is sequential whereas, in actual groups, it can sometimes be made parallel. All these differences point in the same direction: real collaborative groups tend to complete research projects faster than individuals and the modeling choice is reasonably pessimistic (see also Boyer-Kassem and Imbert 2015, 672). So, if collaboration is beneficial with these hypotheses, it should be even more so with more realistic ones.

Secondly, the mathematical features of the step-efficiency function are rather unfavorable and the beneficialness of groups is not bluntly assumed in the model. Indeed, $p_g(k, p)$ has a decreasing growth rate below 1, is bounded, and converges quickly. Consider these features in turn. First, because the growth rate is below 1, k heads are not k times more efficient than 1 (as e.g. Thagard 2006, p. 194). Second, because the game is a race against others, what primarily matters are the relative, not absolute, values of p_g for groups of different sizes. Since the marginal increase of $p_g(k, p)$ vanishes as the size k of the group increases, after a while, having more collaborators makes a tiny difference for larger groups. Third, the boundedness of $p_g(k, p)$ means that the speed of larger teams cannot go beyond one stage per time step, whereas in real cases adding more people can help passing thresholds and complete the task more quickly. Two persons can push a car, ten can carry it¹¹. Finally, because the convergence is quick, the advantage of adding more collaborators shrinks quickly. In any case, our choice of $p(k)$ is aimed at explaining the development of collaboration in general, not at making quantitative predictions about the general efficiency of collaborative groups or specific types of them having different age, gender, country-type, collaborative strategies, etc.

Let us finally discuss collaboration costs, for instance those of communicating intermediate results within the team. By not including them in the model, we by no means claim that collaboration costs are never obstacles to collaboration

¹¹The speed of groups could be made higher than 1 by using an appropriate speed probability distribution, but this would be more favorable to groups.

nor are always negligible. This would be unreasonable, since potential collaborators may not share a common scientific culture, be geographically far apart, or may lack communication facilities, as it used to be the case. In other cases, though, neglecting collaboration costs in the model is a fairly reasonable assumption, typically for a team of scientists who belong to close institutions, and whose research problem can be divided into well-defined subproblems, the solutions of which can easily be communicated. So, the model, though ideal, does fairly well apply to an important set of real cases. In such cases, the overall balance of favorable or unfavorable modeling assumptions in the model should be clear: the benefits of collaboration are strongly underestimated, which can make up for some collaboration costs. Moreover, as the benefits of collaborating are eventually large in the model, the conclusion that collaborating is beneficial (in suitable situations) is safe. We further discuss the issue of collaboration costs in the next section.

8 Scope of the explanation

What is really the scope of the explanation we have defended? We first discuss the case of the humanities, in which collaboration hardly develops. Next, we turn to collaborative costs, and finally, we discuss to what extent our explanation can encompass collaborative mechanisms discussed in the literature.

Objection: what about the humanities? As mentioned above, collaboration is widespread in the sciences, less developed in the social sciences and almost absent in the humanities. Yet, research in the latter fields arguably also involves some successive steps (e.g. reading the literature, developing ideas, drafting a paper), and researchers in these fields also strive for publication and credit. Do not these facts falsify the proposed explanation?

No, because several other important hypotheses of the model are not met in this case. First, the conditions for a general, straightforward and uncontroversial application of the priority rule are hardly present. In the humanities, well-posed problems can be hard to single out, their solutions difficult to assess uncontroversially, and the contributions of co-authors uneasy to delineate. Subfields where the priority rule can be applied (to some extent), like formal or analytical philosophy, are something like exceptions. Similarly, the importance of interpretative methods and the coexistence of incompatible traditions may prevent consensus on the nature of significant problems and what counts as a solution. For instance in philosophy, the question of freedom of human beings may be approached from a naturalistic philosophy of science viewpoint or from a phenomenological tradition, and philosophers from one tradition may not accept as satisfactory the “solutions” advanced by researchers from the other

one.

Second, the model assumes that research problems should be dividable into subtasks, and the solutions of these subtasks should be communicable. But in the humanities, where interpretative practices play a prominent role, it is not that easy to communicate results of subtasks, if that notion makes sense at all. In history of philosophy, almost each specialist of Hegel has her own interpretation of the master's writings for instance. So collaborating Hegelian scholars would just not agree on the validity of intermediate results! Similarly, Thagard (1997, 249) notes that the humanities do not obviously lend themselves to the division of labor.

Third, the successfulness of scientists should be well identified by institutions for (ii-c) to hold. But in the humanities, scholars generally do not share paradigms, methods or norms about what is scientifically sound and significant, and cultural or linguistic barriers can restrain the existence of unified communities.

Communication and collaboration costs. Commonly valued publication venues are needed for the priority rule to apply, and these venues should be easily accessible to the scientists. Arguably, these conditions have progressively been fulfilled by the development of scientific journals and by more efficient means of transportation and communication over the 19th and 20th centuries, which corresponds to the period in which scientific collaboration started and kept developing.

As mentioned above, the model does not include collaboration costs, in particular communication ones. We remarked that the benefits of collaboration are large in the model, thus enabling some reasonably low collaboration costs. But what if these costs are not low? Some empirical input is here useful. The burden of collaboration costs can be studied historically by looking at episodes such as technology shocks, which have lowered transmission costs and reshaped the network structure of scientists who can have informational exchanges. An example is the development of air transportation and cheaper flights (Catalini et al., 2016). Another example, analyzed by Agrawal and Goldfarb (2008), is the introduction of the network Bitnet among American universities in the 1980s and 1990s, which increased collaboration within engineering communities. Bitnet disproportionately benefited middle-tier universities, especially those co-located with top-tier institutions, and had a democratization and equalizing effect (Ding et al., 2010). Conversely, top-tier universities benefited less from such shocks, because their internal collaborative potential was already exploited. These examples suggest that the parallel between the decrease in collaboration costs and the increase of collaboration is not accidental. Too high collaborative costs remove much of the interest of collaborating.

Back to our model, the fact that no collaboration costs are present should be seen as representing an ideal situation. And the closer the actual situation is to this ideal one, the more the conclusions of the model apply, the more

collaboration is beneficial according to the model, and the more our functional explanation is supported. This is by the way consistent with the historical material just presented. So, to come back to the initial worry, our explanation is not limited to the cases where transmission costs are negligible — it is actually relevant for the whole historical spectrum of collaboration patterns. In the context of other inquiries, collaboration costs, spatial constraints or factors unfavorable to collaborative activities and their development could easily be integrated to the model, but we leave this for future work.

A general, encompassing explanation. As pointed above, scholars studying collaborations have proposed a wealth of “explanations of collaboration”. All explanations are not compatible, but the coexistence of a plurality of explanations is also normal. Typically, different types of explanations can be considered depending on the scale at which an explanation is sought, and micro and macro explanatory schemes can be compatible. Also, different causal processes can contribute to some target phenomenon (e.g. global warming). Further, within one and the same causal process, different factors can be correlated with the advent of a target explanandum. Finally, a factor can be embedded in different ways within distinct explanatory schemes (e.g. functional, causal, rational, unificatory, etc.).

To see why our account is compatible with several others from the literature, note that, from a formal point of view, the model shows that differences in the step-efficiency p_i are greatly amplified by the sequential race, *whatever the origin* of these differences. In other words, any factor, whether epistemic or not, that causes an increase in p_i for group i (e.g. if a collaborator is an expert concerning specific steps, if increased resources improve the probability of passing steps, etc.) can provide a decisive differential advantage to this group. This grounds the generality of the model: the information sharing hypothesis of the initial model (Boyer-Kassem and Imbert, 2015) no longer applies and only the step-efficiency differences matter. So, readers can interpret the model in terms of their favorite accounts of what boosts the step efficiency of groups (e.g. pointed expertise or access to resource) and get the amplification-of-differences result with the general robust patterns. Note also that the differences in collaborative patterns across fields may be rooted in differences concerning the factors that boost efficiency in these fields. For example, if collaborating helps harvesting money (see Wray, 2002) and if money makes a step-efficiency difference in science but not in the humanities (or if money is a booster but is rare in the humanities), collaboration will develop in the former but not in the latter.

This is what enables the present explanation to encompass several existing accounts. Consider first the explanation proposed by Muldoon (2017): because acquiring new skills comes with large costs for scientists, hard problems which require a wide range of skills are optimally solved by groups who gather

scientists with varied specialties.¹² This can be framed in the model we consider here: a hard problem involves steps which can only be solved by specific agents, so a larger group of scientists with different pointed fields of expertise, has a higher probability of passing steps — or even allow the group to just pass a step, in a more dramatic way. Then, the results concerning how differences in probabilities are amplified by the sequential race make the job. So Muldoon’s explanation of scientific collaboration is compatible with, and encompassed by, our own.

Consider next the explanation by Thagard (1997, 251), that “peer-similar collaborations can improve reliability by virtues of members of a team noticing mistakes that would get past them working alone”. That is, a larger group has higher chances to correctly pass research steps. Since our model assumes that a step can only be passed correctly or not passed at all, Thagard’s point is just that larger groups have larger probabilities to pass steps, and it is another possible interpretation of the origin of the larger speed of larger groups. Again, his account is compatible with the model we consider, and his explanation is encompassed by ours.

Wray (2002) argues that scientific collaboration’s function is to enable scientific communities to realize their epistemic goals, in particular by accessing the substantial resources that are generally required to conduct research more efficiently. That is, one can imagine that, with more resources, collaborative groups may have better instruments, more administrative and teaching assistants, which overall helps researchers to pass steps more quickly. So, Wray’s explanation, which may apply in fields where big money can make a difference, can be made to fit in the present explanation, with larger groups having higher probabilities to pass steps.

Overall, the explanation of collaboration is probably a multi-factorial issue, and the various authors cited may have identified some part of the truth. The advantage of our functional explanation is that it is able to encompass various such factors because it relies on a formal model in which the origin of the differences in probability to pass steps can be interpreted in various ways. This makes the present explanation general and unifying.

9 Limits, morales, and future work

Some limits of the present model are worth highlighting. The model assumes a general step-efficiency function which only depends on the number of agents, and makes no specific assumptions about the way collaborative groups actually

¹²Appealing to specialization to explain collaboration has a longer history than Muldoon’s recent paper (see e.g. Beaver and Rosen (1978, p. 69)). Muldoon’s specificity is to compare the various costs, and to discuss the relation between recent models of the division of cognitive labor in science and collaboration.

work. This is a strength as far as generality and scope are concerned, but the downside is that it does not provide fine-grained information that could explain collaborative patterns or yield normative implications about scientists with particular features such as academic age, gender, place in collaborative network, access to technical resource or funding. Arguably, this would require developing fine-grained micro-models that would represent the various inner workings of collaborative groups and their epistemic effects. In any case, our results suggest that participation in collaborative groups is a crucial step in the pathway of individuals to scientific success. Thus, (proactive) inclusion of (apprentice) scientists in collaborative groups could be a leverage to counter existing social biases.

Let us now discuss whether the above results about the development of collaboration should be welcomed. It is widely believed that collaboration should be encouraged because it bolsters scientific productivity. A more nuanced picture emerges from our results: a global *increase of collaboration* remains compatible with a global *decrease of productivity* of collaborative groups. Indeed, when the conditions of the models are met, collaborative groups have a differential advantage, and even moderately efficient groups may thereby develop. However, this advantage vanishes when the degree of collaboration increases in the field. It can be kept only if groups grow even bigger, which feeds a potentially detrimental arms race. The culprit seems to be the priority rule. As it rewards only the first competitor to reach a scientific goal, it may disproportionately favor quicker groups, notwithstanding their productivity, which may actually be low. Thereby, the priority rule can encourage overcollaboration, if not research waste. This phenomenon can be even more amplified if funding is tied to apparent indicators such as the sheer number of publications. This suggests that the historical development of collaboration may be, for some part, artificial and counterproductive, as a side effect of the priority rule. However, collaboration also brings distinct advantages above and beyond productivity, such as the quicker completion of scientific inquiries (e.g. when looking for a medical treatment), cross-fertilization and the development of novel ideas, transmission of skills, or the completion of inquiries that individuals cannot do alone. Overall, our results highlight the need to investigate further the pros and cons of scientific collaboration and come as another “reminder that scientists and policy-makers need to have something more than a knee-jerk reaction to the presumed benefits of collaboration” (Bozeman and Corley, 2004).

Regarding the priority rule, philosophers of science have highlighted that it could have unexpected benefits, for example by leading to a better division of cognitive labor (Strevens, 2003) or providing an incentive system that makes the best of scientist’s non epistemic motivations (Zollman, 2018). The present results highlight other important and potentially detrimental effects of the priority rule (see also Muldoon and Weisberg, 2011), which suggest that more

systematic investigations are needed.

Finally, the above findings show that the sequentiality of inquiries can be a core ingredient of the dynamics of scientific communities. Former research in philosophy of science, like the pioneering works by Kitcher (1990) or Strevens (2003), tend to represent research problems as piecemeal units. The present results extend existing works by Boyer (2014) or Heesen (2017) that highlight that sequentiality can be crucial.

10 Conclusion

We have provided a general, functional explanation of the development of collaboration. Its core ingredients are the existence of common scientific goals, the competition triggered by the priority rule, and the possibility to divide research projects into subtasks and to share intermediate results. In this framework, collaborating more than other competitors provides a differential advantage. Then, to the extent that the successfulness of researchers gives them more opportunities to transmit their research habits within scientific subfields, the existence and the development of collaborative practices in communities is favored and a kind of unintentional arms race may take place. This functional scheme is encompassing since any factor that increases the step-efficiency of groups (including features that are usually related to the development of scientific collaboration) can provide such a decisive differential advantage. Finally, this global mechanism is compatible with a global decrease of the productivity of scientific groups. In other words, it provides more evidence that the beneficialness of scientific collaboration cannot be blindly assumed and requires more scrutiny in the future.

Appendix: Proof of the proposition

Proposition (formal version). Consider a community of n agents, including m groups of k agents in a particular collaborative configuration. Let $\mathbf{E}(P_k)$ be the expected individual productivity of an agent in a group of size k in this configuration, and $\mathbf{E}(P_{mk})$ the expected individual productivity of an agent if these m groups of k agents merge. Then $\mathbf{E}(P_k) < \mathbf{E}(P_{mk})$.

Proof. The general idea of the proof is that when collaboration occurs, agents win when they would have won without collaborating and, in addition, can save time. Though this is not needed for the proof, it is also noted that, by saving time, they can also win in more races.

By definition, $\mathbf{E}(P) = \Sigma(p(r)\rho_r)/\Sigma(p(r)\tau_r) = \mathbf{E}(\rho)/\mathbf{E}(\tau)$, where ρ_r , τ_r respectively denote the reward an agent gets and the time she spends in a particular race r , and $p(r)$ the probability of this race.

Consider first the simple case when two individual agents, α and β , compete with no other competitors: $n = 2$, $k = 1$, $m = 2$. To model the passing of step s by agent α in race r , consider the infinite sequence of independent random binary variables $V_{s,i}^{\alpha,r}$, with $p(V_{s,i}^{\alpha,r} = 1) = p$ and $p(V_{s,i}^{\alpha,r} = 0) = 1 - p$. By definition, the time $\tau_s^{\alpha,r}$ for agent α to pass step s alone is $\text{Min}(i|V_{s,i}^{\alpha,r} = 1)^{13}$. Thus, the time $\tau^{\alpha,r}$ α would need in race r to complete the inquiry if β was not competing is $\sum_s \tau_s^{\alpha,r}$. When actually competing against agent β in race r , agent α wins and gets a finite reward $R^{\alpha,r}$ if $\sum_s \tau_s^{\alpha,r} \leq \sum_s \tau_s^{\beta,r}$ and gets 0 otherwise. For a given race, as soon as an agent wins the race, all agents stop the inquiry. Thus, for any race r , the time τ_r that α and β actually spend on inquiry r is equal to $\text{Min}(\tau^{\alpha,r}, \tau^{\beta,r})$. Thus, both agents have the same expected spent time $\mathbf{E}(\tau^{\{\alpha,\beta\}})$. Further, by symmetry, $\mathbf{E}(\rho^{\alpha,r}) = \mathbf{E}(\rho^{\beta,r}) = R/2$, where R denotes the reward in each run. Overall, their expected individual productivity is identical: $\mathbf{E}(P^{\{\alpha;\beta\}}) = R/(2\mathbf{E}(\tau^{\{\alpha;\beta\}}))$.

Suppose now that α and β collaborate (still without other competitors). The group has probability 1 to win the race and the reward is shared between α and β . Thus, their individual expected reward $\mathbf{E}(\rho^{\alpha+\beta})$ is still $R/2$. However, they are now faster: when one agent passes a step, the other one passes it too at the same time, so the time $\tau_s^{\alpha+\beta,r}$ spent to pass step s in race r is $\min(\tau_s^{\alpha,r}, \tau_s^{\beta,r})$. Thus, it takes all agents $\sum_s \min(\tau_s^{\alpha,r}, \tau_s^{\beta,r})$ to complete a race r . Since $\sum_s \min(\tau_s^{\alpha,r}, \tau_s^{\beta,r}) \leq \min(\sum_s \tau_s^{\alpha,r}, \sum_s \tau_s^{\beta,r})$, this establishes that $\mathbf{E}(\tau^{\{\alpha+\beta\}}) \leq \mathbf{E}(\tau^{\{\alpha;\beta\}})$, and $\mathbf{E}(P^{\{\alpha;\beta\}}) \leq \mathbf{E}(P^{\{\alpha+\beta\}})$ (inequality (1)).

To establish that the inequality is strict, note that for a particular race $\sum_s \min(\tau_s^{\alpha,r}, \tau_s^{\beta,r}) = \min(\sum_s \tau_s^{\alpha,r}, \sum_s \tau_s^{\beta,r})$ exactly if for all s , $\tau_s^{\alpha,r} \leq \tau_s^{\beta,r}$ or if for all s , $\tau_s^{\alpha,r} \geq \tau_s^{\beta,r}$ (condition A). For a particular s , $p(\tau_s^{\alpha,r} \leq \tau_s^{\beta,r}) = \gamma$ with γ finite and strictly inferior to 1 because of the independence of the random variables describing the attempts by individuals to pass a step. The two parts of condition A overlap only if for all s , $\tau_s^{\alpha,r} = \tau_s^{\beta,r}$, which has probability ϵ strictly inferior to 1. Overall, the probability that condition A is not met and that inequality (1) is strict is $1 - 2*\gamma^k + \epsilon$ (by the standard rules for calculating the probability of compound events, here the passing of the various steps), and, therefore, it is finite. Further, since for all configurations in which condition A is not met, the gain in time is finite (superior or equal to one time step), eventually on average, the total gain of time is also finite. Thus, on average, the gain in productivity is finite, which establishes the strict inequality in (1).

Mutatis mutandis, the same reasoning holds i) for a number of m individual agents competing against each other; ii) for m groups of size k (with no other competitors), for example by considering that a group of k is equivalent to an agent having probability $1 - (1 - p)^k$ to pass a step; iii) when there are

¹³For instance, in this representation all sequences that start with $(0, 0, 0, 1)$ correspond to the event in which the agents takes $3+1$ steps to pass the step, and this event has probability $p \cdot (1 - p)^3$.

other competitors besides the m groups: a collaborative group of size $g.h$ wins whenever one of its g subgroups would have won, and in addition $\mathbf{E}(\tau)$ strictly decreases. QED.

Addendum. Note that in the context of the existence of other competitors, because collaborative groups works quicker than their subgroups, they win the race in a larger finite fraction of cases in which other competitors would have won otherwise, thus, $\mathbf{E}(\rho^{\{\alpha;\beta\}}) < \mathbf{E}(\rho^{\{\alpha+\beta\}})$, which also contributes to the increase of productivity for groups of equal size that start collaborating.

References

- Abbasi, Alireza, Jörn Altmann, and Liaquat Hossain, 2011, “Identifying the effects of co-authorship networks on the performance of scholars: A correlation and regression analysis of performance measures and social network analysis measures”, *Journal of Informetrics*, 5(4), 594-607
- Agrawal Ajay and Goldfarb Avi, 2008, “Restructuring Research: Communication Costs and the Democratization of University Innovation”, *American Economic Review*, 98(4):1578-1590
- Beaver, Donald deB. and Rosen, Richard (1978), “Studies in Scientific Collaboration: Part I”, *Scientometrics* 1(2): 65-84
- Boyer, Thomas (2014), “Is a bird in the hand worth two in the bush? Or, whether scientists should publish intermediate results”, *Synthese* 191: 17-35
- Boyer-Kassem, Thomas, and Cyrille Imbert (2015), “Scientific Collaboration: Do Two Heads Need to Be More than Twice Better than One?”, *Philosophy of Science* 82 (4): 667–88
- Bozeman, Barry and Corley, Elizabeth (2004), “Scientists’ Collaboration Strategies: implications for Scientific and Technical Human Capital”, *Research Policy*, 33: 599-616
- Catalini, Christian, Fons-Rosen, Christian and Gaule, Patrick, 2016, “Did Cheaper Flights Change the Direction of Science?”, SSRN Scholarly Paper, ID 2774366, Rochester, NY, Social Science Research Network
- Ding, Waverly, Levin, Sharon, Stephan, Paula and Winkler, Anne (2010), “The Impact of Information Technology on Academic Scientists Productivity and Collaboration Patterns”, *Management Science*, 56(9):1439-1461
- Elster, Jon (1983), *Explaining Technical Change: A Case Study in the Philosophy of Science*, New York: Cambridge University Press
- Fallis, Don (2006), “The Epistemic Costs and Benefits of Collaboration”, *Southern Journal of Philosophy* 44 S: 197–208
- Frey Daniel, and Šešelja Dunja (2019), “Robustness and Idealizations in Agent-Based Models of Scientific Interaction”, *The British Journal for the Philosophy of Science*. <https://doi.org/10.1093/bjps/axy039>

- Heesen, Remco (2017), “Communism and the Incentive to Share in Science”, *Philosophy of Science* 84: 698–716
- INSERM (2005), “Les indicateurs bibliométriques à l’INSERM”, https://www.eva2.inserm.fr/EVA/jsp/Bibliometrie/Doc/Indicateurs/Indicateurs_bibliometriques_Inserm.pdf
- Katz, J. Sylvan and Martin, Ben R. (1997), “What is Research Collaboration?”, *Research Policy* 84: 1–18
- Kincaid, Harold (1996), “Functionalism defended”, in *Philosophical Foundations of the Social Sciences*, Cambridge University Press, p. 101-141
- Kincaid, Harold (2006), “Functional Explanations”, in Stephen Turner and Mark Risjord (volume editors), *Handbook of the Philosophy of Science. Volume 15: Philosophy of Anthropology and Sociology* (Handbook editors: Dov M. Gabbay, Paul Thagard and John Woods.), p. 205-239. Elsevier BV
- Kitcher, Philip (1990), “The Division of Cognitive Labor”, *The Journal of Philosophy* 87(1): 5-22
- Kitcher, Philip (1993), *The Advancement of Science*. New York: Oxford University Press
- Kuhn, Thomas S. (1962), *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press
- Melin, Göran, (2000), “Pragmatism and self-organization: Research collaboration on the individual level” *Research Policy*, 29, 31–40
- Laughlin, Patrick (2011), *Group Problem Solving*, Princeton, Princeton University Press
- Lee, Soohoo, Bozeman, Barry (2005), “The Impact of Research Collaboration on Scientific Productivity”, *Social Studies of Science* 35: 673–702
- Li Eldon, Liao C. Hsiang and Hsiuju Yen (2013), “Co-authorship networks and research impact: A social capital perspective”, *Research Policy* 42: 9 1515–1530
- Merton, Robert K. (1968), “The Matthew Effect in Science: The Reward and Communication Systems of Science Are Considered”, *Science*, 159 (3810): 56–63
- Muldoon, Ryan (2017), “Diversity, Rationality, and the Division of Cognitive Labor”, in Boyer-Kassem, T., Mayo-Wilson, C. and Weisberg, M. (eds.), *Scientific Collaboration and Collective Knowledge*, New York: Oxford University Press
- O’Connor Cailin (2019), “The natural selection of conservative science”, *Studies in History and Philosophy of Science Part A*, 76: 24–29
- Parish, J. Austin, Kevin W. Boyack and John P. Ioannidis (2018), “Dynamics of co-authorship and productivity across different fields of scientific research”, *PloS One*, 13 (1)
- Price, Derek John de Solla (1963), *Little Science, Big Science*, New York,

- Columbia University Press
- Smaldino Paul E., McElreath Richard (2016), “The natural selection of bad science”, *Royal Society Open Science*, 3(9):160384
- Strevens, Michael (2003), “The Role of the Priority Rule in Science”, *The Journal of Philosophy* 100(2): 55-79
- Thagard, Paul (1997), “Collaborative Knowledge”, *Nouûs* 31(2): 242—261
- (2006), “How to Collaborate: Procedural Knowledge in the Cooperative Development of Science”, *The Southern Journal of Philosophy*, XLIV: 177—196
- Muldoon, Ryan and Michael Weisberg, 2011, “Robustness and idealization in models of cognitive labor”, *Synthese* 183 (2): 161-174
- Wray, K. Brad (2002), “The Epistemic Significance of Collaborative Research”, *Philosophy of Science* 69 (1): 150-168
- Wu, Lingfei, Dashun Wang, James A. Evans (2019), “Large Teams Have Developed Science and Technology; Small Teams Have Disrupted It”, *Nature* 566:378–382
- Wuchty, Stefan, Jones, Benjamin F. and Uzzi, Brian (2007), “The Increasing Dominance of Teams in Production of Knowledge”, *Science* 316(5827): 1036-1039
- Zeng Xiao Han T. et al. (2016) “Differences in Collaboration Patterns across Discipline, Career Stage, and Gender”, *PLOS Biology*, 14 (11)
- Zollman, Kevin J. S. (2010), “The Epistemic Benefit of Transient Diversity”, *Erkenntnis* 72:17–35
- Zollman, Kevin J.S (2018), “The Credit Economy and the Economic Rationality of Science”, *The Journal of Philosophy*, 115 (1) 5–33
- Zuckerman, Harriet (1967), “Nobel Laureates in Science: Patterns of Productivity, Collaboration, and Authorship”, *American Sociological Review* 32 (3): 391-403.