

# Explaining Scientific Collaboration: a General Functional Account

Thomas Boyer-Kassem\* and Cyrille Imbert†

(Working paper – July 23, 2019)

## Abstract

For two centuries, collaborative research has kept developing. Various explanations of this trend have been proposed, e.g. based on the better reliability of teams, scientific specialization, or access to resources. Here, we offer a novel functional explanation of scientific collaboration that encompasses existing explanations. Our argument is grounded in the study of a sequential competitive model of scientific research. We derive from this model robust patterns about the successfulness of collaborative groups which feed this functional account — the existence and the increase of scientists' collaborative behavior can be explained functionally by appeal to the increase of individual successfulness it brings.

---

\*Département de Philosophie, MAPP, Université de Poitiers, France.  
[thomas.boyer.kassem@univ-poitiers.fr](mailto:thomas.boyer.kassem@univ-poitiers.fr)

†CNRS, Archives Poincaré, France. [cyrille.imbert@univ-lorraine.fr](mailto:cyrille.imbert@univ-lorraine.fr). The authors contributed equally to the paper.

# 1 Introduction

For two centuries, co-authoring papers has become increasingly widespread in academia (Price, 1963, Beaver and Rosen, 1979, Wuchty et alii, 2007), especially in the last few decades. Since the 1950s, the percentage of co-authored papers has grown at a common pace for science and engineering, social sciences, and patents; the mean size of collaborative teams has also increased, and even more so in science and engineering. No such increase is visible for the art and humanities. Various explanations of this collaborative trend have been proposed (Boyer-Kassem and Imbert 2015, p. 668-670). For example, it has been explained by scientific specialization (Muldoon 2017), by the better reliability of research work done by groups (Thagard 1997, 251) or by the need to access resources (Wray 2002). These explanatory factors are of different types, some being epistemic and others non-epistemic. They are usually thought of as distinct, competitive explanatory factors. By contrast, we show that many of these factors can be included under a common scheme of an encompassing explanation. This new explanation 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 both scientists' collaborative behavior and its increase can be functionally explained by the increase in the individual successfulness of collaborative scientists. That is, scientists collaborate because they are more successful — they are more productive and get more credit — 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 formal model of scientific research introduced by Boyer-Kassem and Imbert (2015) that highlights the potentially significant beneficialness of collaboration. In this model, a scientific inquiry is pictured as a multi-step process, which requires finding successive solutions to a sequence of intermediate problems. With this simple hypothesis, it is shown that minor differences in the efficiency to pass steps can make collaborative groups much more successful in competitive contexts. Here, we derive additional results from that model, and use them to argue for a robust functional explanation of collaboration. We highlight the generality of this explanation and the extent to which it encompasses existing ones. We argue that several of these can indeed be embedded within Boyer-Kassem and Imbert's model, if it is suitably reinterpreted, i.e. that the various factors traditionally used to explain collaboration (e.g. specialization, increased reliability, etc.) can be one of the sources of the differences in step efficiency that play the crucial role in the model.

The analysis of scientific research problems as composed of several steps is a core ingredient of what makes the explanation work, and the basis of its generality. The emphasis on the effects of sequentiality in the philosophical modeling of scientific research is recent, starting with a model in Boyer (2014) on intermediate results, generalized in Heesen (2017), and reused in Boyer-Kassem and Imbert (2015) with an application to collaboration. Drawing attention to sequentiality in scientific research can be seen as a desidealization of former descriptions that took a research problem to be an unanalyzed unit, like in pioneering modeling work in Kitcher (1990) and Strevens (2003), and taken by default by other philosophers afterwards, including Wray (2002) when considering the explanation of collaboration. This paper shows that the sequential nature of research problems cannot be idealized away without a loss of explanatory power and conversely that the richer, multi-step view paves the way for a more robust and encompassing explanation of scientific collaboration.

The paper is organized as follows. We start by presenting functional explanations in general and the one we propose for scientific collaboration in particular (Section 2). We introduce the model by Boyer-Kassem and Imbert in Section 3. In Sections 4 and 5, we show how new results from that model can be used to establish functional claims about collaboration. Finally, in Section 6, we discuss the generality of this explanation.

## 2 Functional Explanations and Collaboration

We review in this section what sound functional explanations are and how they can be used in the present case. Functional explanations explain the existence of a feature by one of its effects, usually its usefulness or beneficialness. For example, “The state exists in order to promote the interests of the ruling class” is a marxist functional explanation (Kincaid 1996, p. 104), or “golf clubs are functional in enabling business people, bankers, and various professionals like lawyers and accountants to get to know one another, establish networks and reinforce their mutual confidence” is an anthropological one (Pettit 1996, p. 296). As is well-known, pointing at a useful effect of  $P$  is not enough to provide a *sound* functional explanation of it. For instance, the usefulness of the nose to carry glasses cannot explain why humans have one. Nevertheless, if stringent conditions are met, it is usually considered that functional explanations can be satisfactory, typically within biology and in the social sciences.<sup>1</sup>

How characterize suitable functional explanations? Like Wray (2002), we choose to use Kincaid’s account because it is simple and widely accepted. According to Kincaid (1996, 110-111; 2006, 214-215),  $A$  is functionally explained by  $B$ , or  $A$  exists “in order to  $B$ ” if:

---

<sup>1</sup>Even Elster, who otherwise favors methodological individualism, agrees that functional explanations can be acceptable in the social sciences (Elster 1983).

- (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.” Note that this account of functional explanations manages to rule out the spurious explanation of the existence of the nose by appealing to its usefulness to carry glasses.

To provide a functional explanation of scientific collaboration, we need the concept of a scientist’s *individual successfulness* when working on a research project. Following Boyer-Kassem and Imbert (2015, p. 672), we mean by this term “the amount of reward an individual researcher gets (alone or by collaborating) divided by the time spent” on the research project. That reward is traditionally scientific credit, and “this is the quantity that a selfish scientist wants to maximize” (ibid.). Individual successfulness can be seen as the productivity of the scientist.

We are now in a position to state the two claims we defend in this paper. The first one deals with the *existence* of scientific collaboration:

- (1) scientists’ collaborative behavior can be functionally explained by appeal to the increase of individual successfulness it brings; in other words, scientists collaborate “in order to” increase their individual successfulness.

Our second thesis deals with the *increase* of scientific collaboration that has been observed for two centuries or so:

- (2) the increase of scientists’ collaborative behavior can be functionally explained by appeal to the increase of individual successfulness it brings; in other words, scientists collaborate more over time “in order to” increase their individual successfulness.

Following Kincaid’s analysis, we shall establish these theses by arguing for the following claims:

- (i-c) (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. We agree with Wray (2002, 161) that it is implausible to consider that the high successfulness of scientists is the initial cause of collaboration. Many scientists have been successful (and continue to be in some fields) without collaborating. Also, collaboration is recent, while there have always been successful scientists, so if successfulness was the initial cause of collaboration, one would need to explain why for so long it did not generate collaboration, and it seems implausible. So, we consider that (iii-c) is established. What remains to be argued for are claims (i-c) and (ii-c), which is done in Sections 4 and 5. For that, we first present the model which is central to our analysis in the next section.

### 3 The Sequential Competitive Model: Main Results and Explanatory Lacunas

In the model introduced by Boyer-Kassem and Imbert (2015), a community of  $n$  agents 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. What constitutes a step may depend on one's level of analysis, and one may choose to adopt a finer- or coarser-grained model with different values for  $l$ . In the extreme case of  $l = 1$  step, the model has no particular interest compared to other ones like Kitcher (1990) or Strevens (2003); when  $l$  is of a few units, interesting effects appear in the model, with robust properties not depending on the specific value of  $l$ . In the remainder of the paper, the value of  $l = 10$  is taken for illustrative purposes.

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, the value of  $p$  is set to 0.5.<sup>2</sup> 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).

What it means for agents to be in a collaborative group, in this model, is that they share information, namely the passing of steps: when an agent has passed a step, her collaborators pass it with her for free during the same temporal interval, and all start trying to pass the next step (see Section 6 for

---

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

a more general interpretation). Thus, a collaborative group of size  $k$  passes a step with probability  $1 - (1 - p)^k$ . The larger the group, the quicker it will make progress on average.

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. Among a collaborative group, the reward is equally split between agents. Divided by the time spent, this gives the individual successfulness, which is what individuals are supposed to be interested in.

Is this a realistic model of scientific collaboration? Obviously not; in real life, agents may be more expert of one step and less of another, or collaborators may have ideas that loners won't have and pass steps even quicker. But these differences overall point in the same direction: real collaborations will tend to be even more productive than in the model, i.e. give higher individual successfulness (Boyer-Kassem and Imbert 2015, p. 672). So, if collaboration is beneficial with these hypotheses (as will be argued), it should be even more so with more realistic ones. Collaborative costs are discounted but this would simply make collaboration less beneficial without threatening the general conclusions. Note also that the model is not aimed at quantifying the actual successfulness of collaborative agents, but at analyzing the differential successfulness of agents, depending on their collaborative behavior.

Because claims (i-c) and (ii-c) bear on individual successfulness (see Section 2), 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, collaborating is beneficial for particular groups in particular collaboration configurations only: a group of 2 is very successful in configuration (2-1-1-1-1) but not in (7-2). Thus, the model mostly provides possibility results about what can be the case in some configurations. Third, the explanandum is a general social feature of modern science, not some collaborative behavior in some particular case, so the explanans must also involve general statements

about the link between collaboration and beneficialness. Then, if the model is to provide generic social mechanisms with explanatory import, the effects of these mechanisms must be described at a general level and general, invariant patterns between collaboration and beneficialness must be provided. This is what we do in the next sections.

## 4 Collaboration Causes Successfulness

In this Section, we use the above model to derive new results and provide strong evidence in favor of (i-c), namely that “(the increase of) scientists’ collaborative behavior causes an increase of their individual successfulness.” This causal relation between collaboration and successfulness needs to be general and robust, and should go beyond the description of the beneficialness of collaboration in particular situations.

A first route is to find general results about when it is beneficial for individuals to collaborate, such as the following proposition (proof in the appendix):

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

In other words, 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. Note that this proposition covers both cases in (i-c): the existence of scientists’ collaborative behavior (with the corollary), and its increase (in the general case, with the merging of collaborative groups). However, many configurations don’t include groups of equal sizes, so only a small subset of possible configurations are covered, and the proposition cannot provide a general vindication for the causality claim (i-c).

A second route to investigate the general link between collaboration and success is to analyze what happens on average over all configurations, and this is what we do in the remainder of this Section. We study in turn the influence of various parameters — community size, collaborative group size, and average level of collaboration in the community — on the individual successfulness, so as to establish the robustness of the causal claim.<sup>3</sup>

---

<sup>3</sup>Technically, we compute the average individual successfulness of an agent in a collaborative group of  $k$  for 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? Some have more combinatorial realizations than others, but this is not necessarily what matters. Ultimately, the right weights depend on (the set of) empirical situation(s) one is interested in. Because we don’t have clues about the actual frequency of configurations in scientific contexts — which

**Successfulness and community size.** Consider the relation between individual successfulness and community size. Figure 1 shows the average individual successfulness as a function of community size  $n$ , for agents who are 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. This could be 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 for someone in a group of size  $k$  is below a curve of some smaller group. This means that collaborating too much is not the best strategy.<sup>4</sup> Third, when groups are larger, this initial plate 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 the analysis of 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 known to a handful of scientists. When communities become significantly larger (typically, concerning normal science problems that many researchers are likely to tackle), small groups remain viable but tend to be outraced.<sup>5</sup> Thus, collaborating moderately can be a risky strategy when uncertainty prevails about other competing groups. Finally, being in large collaborative groups rarely gives very high successfulness, but it is a reasonably safe situation, with moderate differences depending on the specific size of the group, or on the community size. This provides strong support to the claim that, for various community sizes, collaboration (as opposed to no collaboration at all) causes individual successfulness, and that collaborating (moderately) more causes more successfulness, and more robustly so.

---

is different from the frequency of publishing groups —, we privilege simplicity and give equal weights to all configurations.

<sup>4</sup>See Boyer-Kassem and Imbert (2015, 679-80) for an analysis of over-collaboration in large groups.

<sup>5</sup>This is particularly in line with recent findings by Wu et al (2019) that large teams develop science while small teams disrupt it.

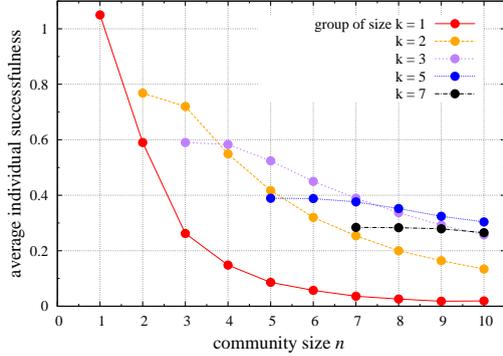


Figure 1: Variation of individual successfulness with community size.

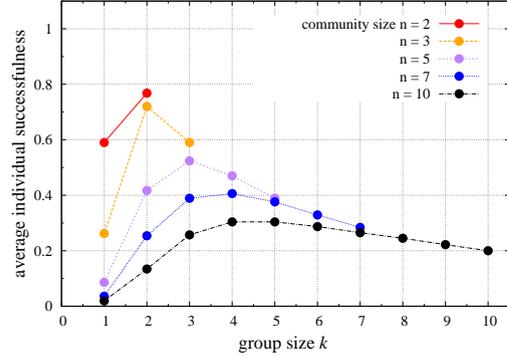


Figure 2: Variation of individual successfulness with the size of groups.

**Successfulness and group size.** Figure 2 shows how an individual’s successfulness changes with the size of the collaborative group she is in, for various community sizes. Note first that, for  $n > 2$ , the successfulness curve has a one-peaked (discrete) form. This means that there is an optimal collaborative group size, and it is not the community size — adding a new collaborator beyond this size means that the gains she brings is not compensated by the equal share she takes. 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 postulating an *ad hoc* mechanism that would build the possibility of overcollaborating in the model. This also 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 less harmful and risky than under-collaborating. For stable large networks that collaborate a lot, adding or removing some collaborators does not change substantially their (high) successfulness. So, the successfulness of these groups are 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 environment.<sup>6</sup> Consider how the various curves stand

<sup>6</sup>For each collaboration configuration, we compute the mean group size, for instance 1.75 for collaboration configuration (3-2-1-1). The average successfulness is computed for each mean group size over a neighboring interval for communities up to size 10. More precisely, for each group size  $k$ , we compute the average successfulness over configurations having

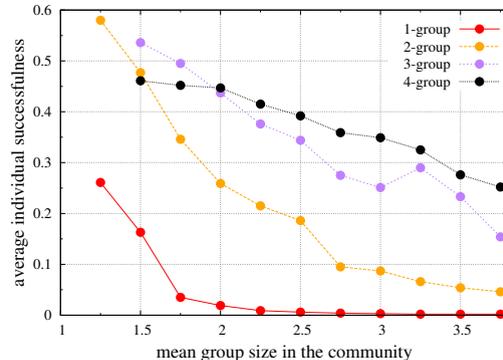


Figure 3: Variation of successfulness with the degree of collaboration in communities.

relatively 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 2. 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. Overall, 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.

Coming back to our causality claim (i-c), the results of the model show that collaborating entails successfulness. This relation is robust under changes in the size of communities or in the exact size of groups. Further, those who collaborate more than average are more successful. Collaborating too much is slightly problematic, under-collaborating is strongly so. Hence, collaborating a lot is a safe working habit, especially in the absence of information about the size and structure of the competing community. In light of this evidence, (i-c) seems adequately supported: (the increase of) scientists' collaborative

---

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"),  $[1.5, 2]$  ...  $[3.5, 4]$ .

behavior causes an increase of their individual successfulness.

## 5 Collaborative Behaviors Exist and Develop Because of the 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 wide variety of social mechanisms across scientific contexts, and we shall be content with giving various evidence that strongly suggests the plausibility of this link.

**Transmission.** Knowing how and when to collaborate is not straightforward. Like other know-how skills, it can be developed by exercising it with people who already possess the relevant procedural knowledge. In this case, people who already collaborate can endorse this role of cultural transmission for colleagues, and above all students (Thagard 2006). Working with students is indeed an efficient way to train them as scientists (Thagard 1997, 248–50), so scientists have incentives to enroll students in their collaborative groups. Then, the cultural transmission of collaborative practice does not require any particular effort on top of that. 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, so the whole process can be gradual and accumulative.

**Transmission opportunities.** We now argue that collaborative scientists, because they are more successful, will more often be in a position to transmit their collaborative habits and that collaboration will persist and even develop.

Within pure science, because scientific success is the official goal of science, successful scientists can be expected to stand better chances to get good positions and grants, develop research programs, direct larger teams and pass over their collaborative habits. And although scientific institutions are imperfect, it seems that this is actually often the case. Further, within applied science, in which collaboration is also widespread (Wuchty 2007), research projects are usually directed at finding profitable applications, which can be patented. Thus, fund providers are directly and strongly interested in hiring and providing resource to successful scientists, who develop such applications. Note that

it is merely needed that the pragmatic rewards of scientists is linked on average to their success, and it remains compatible with the fact that *some* epistemically successful scientists get little resource and *some* unsuccessful scientists get a lot.

In practice, non-epistemic factors may even tend to over-credit successful scientists, and in particular collaborative ones, giving them even more opportunities. First, individual successfulness has been assessed in the model with a conservative estimate. It seems that an agent’s publication within a group of size  $k$  is actually more appreciated than just  $1/k$  of 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 will generally be more visible than one single-authored publication, since more people can promote or publicize collective publications and research topics. Third, sociology of science seems to indicate that scientific credit tends to accrue to a subset of scientists who are perceived as extremely successful — this is the Matthew effect (Merton, 1968). And 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 from this effect and transmit more their working habits. The concentration of credit and resource may further stimulate collaborative behavior with these fortunate scientists.

Other types of mechanisms may contribute to this process, like conscious ones. So far, agents have only been supposed to follow their working habits and sometimes transmit them. But supplementary intentional or imitative processes may also feed this dynamics.<sup>7</sup> 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 their expertise to produce collective articles. Then, these institutions and their members will be more successful, attract resource, and keep developing and transmitting their working habits. In light of the above discussion, the causal connection between the success of collaborative

---

<sup>7</sup>Kincaid 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 2] true” (Kincaid 1996, 112).

scientists and the persistence and development of collaborative practices (ii-c) is highly plausible.

## 6 Discussion: scope of the explanation

What is really the scope of the explanation we have defended? First, we argue that it is a general explanation which encompasses several collaborative mechanisms discussed in the literature. Second, we discuss why the humanities do not seem to fall under its scope.

**A general, encompassing explanation.** The initial model makes a simple hypothesis: collaboration within a group consists in agents independently researching for the next step and sharing it when they have passed it. Isn't it a specific, and unrealistic, hypothesis? Yes, but it anyway enables to reach very general results. Here is why (Boyer-Kassem and Imbert 2015, p. 682). Formally speaking, the net result of this hypothesis in the model is to give various probabilities  $p_i$  of passing steps for groups, depending on the size of group  $i$ . Then, the details of the inner working within the group as described in the initial model (namely, that agents pass steps with the same probability  $p$  and share information) can be forgotten — this is just one possible mechanism, and it needs not be the only one or the true one. What really matters is just that the model describes a race between collective agents who have various probabilities of passing steps. So, the results of the model is that differences in  $p_i$  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 make this group as successful as those depicted in the model who research independently and share information — hence the generality of this mechanism. This really grounds the generality of the model: the information sharing hypothesis can be discarded once the relation between group sizes and probabilities to pass step has been established, which is just what matters. So, readers who are skeptical towards what it means to collaborate in the initial model (namely sharing steps) can substitute it with their favorite mechanisms, like expertise or access to resource, and get the amplification-of-differences result with the general robust pattern.

This is what enables our explanation to encompass several existing ones. Consider first the explanation proposed by Muldoon (2017): because acquiring new skills comes with a large cost for scientists, hard problems which require a wide range of skills will optimally be solved by groups who gather scientists with varied specialties.<sup>8</sup> This can be framed in the model we consider here: a

---

<sup>8</sup>Appealing to specialization to explain collaboration of course has a longer history than

hard problem involves some steps which can only be solved by specific agents, so a larger group of scientists, with suitable specialties, will increase its probability of passing steps — or even allow the group to just pass a step, in a more dramatic way. This is just in line with our model’s feature that larger groups have higher probabilities to pass steps. Then, our results according to which differences in probabilities are greatly 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, which is what the model assumes. 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 will have better instruments, more administrative and teaching assistants, which overall will help researchers to pass the steps more quickly. So, Wray’s explanation 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 them, and this is made possible by its reliance on a formal model in which the origin of differences in probabilities to pass steps can be reinterpreted in various ways. As such, our explanation is general and unifying.

Our explanation points at a growth in the size of collaborative groups. Note that while nothing in the model provides an internal limit to the growth of collaboration, there is a wealth of reasons why collaborating groups cannot develop forever. For example, communities are limited in size, spatially distributed, and collaboration is all the more costly as groups are large. The model could be easily modified to integrate factors that limit the success and development of collaboration.

**Objection: what about the humanities?** According to recent studies

---

Muldoon’s recent paper; see for instance Beaver and Rosen (1978, p. 69) for references of the 1950’s. But 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.

(Uzzi et al 2007), the mean size of teams in the arts and the humanities has almost not varied along the last decades, and is still very close to 1. In other words, researchers in the arts and the humanities just don't collaborate. Yet, research in these fields arguably also involves some successive steps (reading the literature, developing ideas, drafting a paper, and so on), and collaborating would fasten the progress of researchers on the project. So isn't our explanation blatantly refuted by the absence of collaboration in these fields?

Actually no, because several hypotheses of the model (besides the sequential aspect) are not met. First, conditions for the application of the priority rule are not fulfilled. In particular, it should be possible to single out problems and to state uncontroversially when they are solved. Thagard (1997, 249) notes that the humanities do not obviously lend themselves to the division of labor and to teacher/apprentice collaborations. 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 it is quite unlikely that philosophers from one tradition will 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 practice play a prominent role, it is not that easy to communicate results to subtasks, if that notion makes sense at all. Take for instance history of philosophy, and note that almost each specialist of Hegel has her own interpretation of the master's writings. So there is no sense in which two specialists could work together on a first step — read Hegel or some commentators on some point — and communicate to the other one the result of her research — the other specialist would just not agree on her results! On the opposite, these conditions seem to be met in the formal and empirical sciences, somewhat met in the social science, and this enables our explanation to apply to these fields, and to explain a feature which is indeed present.

Third, the successfulness of scientists should be well identified by institutions for (ii-c) to hold, but this is not particularly the case in the humanities. There, scholars generally do not share paradigms, methods or norms about what is scientifically sound and significant. Cultural or linguistic barriers can restrain the existence of unified communities and common publication venues. Thus, it is generally harder for institutions to identify successful researchers in the humanities than in the sciences.

## 7 Conclusion

We have provided a general, functional explanation of the existence and development of collaboration. Key ingredients in this explanation include the existence of competition among researchers, well-identified goals and results, the use of the priority rule, and research projects being composed of sequential steps. In this framework, we have shown that collaborating a lot is overall a safe and success-conducting practice. This conclusion is robust for various sizes of groups, communities and degrees of collaboration; everything else being equal, those who collaborate more than average do better, which feeds a kind of unintentional arms race. Then, to the extent that the successfulness of researchers gives them more opportunities to transmit their research habits, the existence and the development of collaborative practices in communities can be functionally explained. What is more, our functional explanation encompasses various factors that have been proposed to account for collaboration, and thus can be claimed to have a large scope.

## 8 Appendix: Proof of the Proposition

Consider first the simple case where the  $m$  groups of size  $k$  don't have other competitors. By symmetry, all groups have the same probability  $1/m$  to win the race and get the reward — call this reward  $r$ . So, the individual expected reward is  $r/(km)$ . Suppose now the groups merge and all  $km$  agents collaborate. Each of them will receive the same reward, so their expected individual rewards are  $r/(km)$  too. However, what matters in the model is not the expected reward, but the successfulness, which is this quantity divided by time. Because within a collaboration agents share all the steps they pass, the larger group of size  $km$  will be at least as quick, and sometimes more, than all groups of size  $k$ — more precisely: for a given drawing of all random variables corresponding to attempts to pass the steps, for all agents and temporal intervals, the group of size  $km$  will move at least as quickly as all groups of size  $k$ . So the individual successfulness is at least as high when identical groups merge.

Consider now the case where there are other competitors than the  $m$  groups of size  $k$ . For a given drawing of all random variables, either the winner is one of the  $m$  groups, or another competitor. In the former case, the above reasoning can be made again, and the same conclusion holds. In the latter case, there is nothing to lose, and because the group of size  $km$  is sometimes quicker than the  $m$  groups of size  $k$ , there can be additional cases where it outcompetes the other competitors; then, the individual successfulness increases with the merging. QED.

## 9 References

- Beaver, Donald deB. and Rosen, Richard (1978), “Studies in Scientific Collaboration: Part I”, *Scientometrics* 1(2): 65-84.
- Beaver, Donald deB. and Rosen, Richard (1979) “Studies in Scientific Collaboration: Part III”, *Scientometrics*, 1(3): 231-245.
- 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.
- Elster, Jon (1983), *Explaining Technical Change: A Case Study in the Philosophy of Science*, Studies in Rationality and Social Change, New York: Cambridge University Press.
- Fallis, Don (2006), “The Epistemic Costs and Benefits of Collaboration”, *Southern Journal of Philosophy* 44 S: 197–208.
- 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](https://www.eva2.inserm.fr/EVA/jsp/Bibliometrie/Doc/Indicateurs/Indicateurs_bibliometriques_Inserm.pdf)
- 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.
- Philosophical Foundations of the Social Sciences*, Cambridge University Press.
- Kitcher, Philip (1990), “The Division of Cognitive Labor”, *The Journal of Philosophy* 87(1): 5-22.
- 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.
- Price, Derek John de Solla (1963), *Little Science, Big Science*, New York, Columbia University Press.
- Strevens, Michael (2003), “The Role of the Priority Rule in Science”, *The*

- Journal of Philosophy* 100(2): 55-79.
- Thagard, Paul (1997), “Collaborative Knowledge”, *Nous* 31(2): 242—261.
- (2006), “How to Collaborate: Procedural Knowledge in the Cooperative Development of Science”, *The Southern Journal of Philosophy*, XLIV: 177—196.
- 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.