Explaining Scientific Collaboration: a General Functional Account

Thomas Boyer-Kassem* and Cyrille Imbert†

October, 2018

Abstract

For two centuries, collaborative research has become increasingly widespread. Various explanations of this trend have been proposed. Here, we offer a novel functional explanation of it. It differs from accounts like that of Wray (2002) by the precise socio-epistemic mechanism that grounds the beneficialness of collaboration. Boyer-Kassem and Imbert (2015) show how minor differences in the step-efficiency of collaborative groups can make them much more successful in particular configurations. We investigate this model further, derive robust social patterns concerning the general successfulness of collaborative groups, and argue that these patterns can be used to defend a general functional account.

*MAPP (EA 2626), Univ. Poitiers, France. `thomas.boyer.kassem@univ-poitiers.fr`
†CNRS, Archives Poincaré, France. `cyrille.imbert@univ-lorraine.fr`

1 Introduction

For two centuries, co-authoring papers has become increasingly widespread in academia (Price, 1963, Beaver and Rosen, 1979), especially in the last few decades. Since the 1950s, the percentage of co-authored papers has grown at a common rhythm 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 (Wuchty et alii, 2007).

Various explanations of this collaborative trend have been proposed: for example, it may be caused by scientific specialization, it may increase the productivity or reliability of researchers, or be promoted by the rules of credit attribution. Here, we aim at offering a new functional explanation of this trend by showing that collaboration exists because it increases the successfulness of scientists. The present explanation differs from accounts like that of Wray (2002) by the social and epistemic mechanism that grounds the beneficialness of collaboration. We analyze further an existing model that shows how minor differences in the step-efficiency of collaborative groups at passing the steps of a project can make them much more successful in particular configurations (Boyer-Kassem and Imbert, 2015) and show how it can be used to build a general and robust functional explanation of collaboration.

We introduce the model in section 2. After presenting functional explanations (section 3), we show how the model can be used to derive robust social patterns of the successfulness of collaborative groups (section 4), and argue that these patterns can refine and strengthen functional explanations of collaboration like the one defended by Wray (sections 5 and 6).

2 Boyer-Kassem and Imbert's Model: Main Results and Explanatory Lacunas

Boyer-Kassem and Imbert (2015) investigate a model in which n agents struggle over the completion of a research project composed of l sequential steps. At each time interval, agents have independent probabilities p of passing a step. When an agent reaches the end of the project, she wins all the scientific credit and the race stops (this is the priority rule). Agents can organize themselves into collaborative groups for the whole project, meaning that they only share information, i.e. step discoveries — clearly, there are more favorable hypotheses associated with collaborating, like having new ideas or double-checking (see below). Within a group, agents make progress together, and equally share final rewards. Thus, a group of k agents (hereafter k-group) passes a step with probability $p_g(k,p) = 1 - (1-p)^k$. In forthcoming illustrations, the value of l is set to 10 and that of p to 0.5, which is not particularly favorable for groups (ibidem, 674). If collaboration is beneficial with these hypotheses, it will be even more so with more favorable or realistic ones. A community of n agents (hereafter, n-community) can be organized in various k-groups. For example, a 3-community can correspond to configurations (1-1-1), (2-1) or (3). The individual successfulness of an agent in a k-group in a particular configuration is defined as the average individual reward divided by time. It has been obtained for all configurations up to $n = 10$, on millions of runs.

Note that this 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. The main finding is that minor differences in the efficiency at 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 (5-4) (resp. (2-1)) configuration, whereas the difference in step efficiency between the 5 (resp. 2) and the 4-group (resp. 1-group) is 3% (resp. 50%), the difference in individual successfulness is 25% (resp. 700%). The scope of these results actually goes beyond the initial hypotheses in terms of information sharing. Formally speaking, the model is a race between (collective) agents i with probabilities p_i of passing steps. *Whatever the origin* of the differences in p_i, they are greatly amplified by the sequential race. In other words, any factor, whether epistemic or not, that implies an increase in p_i of a k-group (e.g. if a collaborator is an expert concerning specific steps, if increased resources improve step-efficiency, etc.) makes this group as successful as a larger group — hence the generality of this mechanism.

Still, these results do not explain scientific collaboration by themselves. First, collaboration is beneficial for particular k-groups in particular configurations only: a 2-group 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 certain configurations. Second, 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 presents generic social mechanisms with explanatory import, one needs to describe at a general level the effects of these mechanisms and provide some general, invariant pattern between collaboration and beneficialness. This is what we do in section 4. A final serious worry is that the beneficialness of a state by no means explains why it exists, nor perseveres in being. A link needs to be made between the beneficialness of collaboration and its existence over time. We suggest that this connection can be accounted for functionally.

3 Functional Explanations and Collaboration

We review in this section how functional explanations work and how they can be used in the present case. We follow Wray's choice to use Kincaid's account because it is simple, widely accepted, and that nothing substantial hinges on this choice. Functional explanations explain the existence of a feature by one of its effects, usually its usefulness or beneficialness. As such, they can be sloppy and badly flawed. The usefulness of the nose to carry glasses does not explain that humans have one. Nevertheless, if stringent conditions are met, it is usually considered that functional explanations can be satisfactory, typically within biology. Even Elster, who otherwise favors methodological individualism, agrees that functional explanations can be acceptable in the social science (Elster, 1983). According to Kincaid (1996, 105-114), P is functionally explained by E, i.e. P exists "in order to promote <effect E>" if:

(1) P causes E,

(2) P persists because it causes E,

(3) P is causally prior to E.

Then, a functional explanation of collaboration should have the following form:

(1c) Scientists' collaborative behavior causes the increase of their individual successfulness.

(2c) Scientists' collaborative behavior persists (or develops) because it causes a higher individual successfulness.

(3c) Collaborative behavior is causally prior to this increased individual successfulness that is rooted in collaborative behavior.

We agree with Wray (2002, 161) that it is implausible to consider that the high successfulness of scientists is the initial cause of collaboration since many scientists have been successful (and continue to be in some fields) without collaborating. In the same time, there can be various contingent reasons why some researchers have decided to engage in some collaboration. So, what calls for an explanation is the fact that collaboration is widespread and persistent, not its occasional existence.

4 Collaboration Causes Successfulness

We now argue that the above model provides strong evidence in favor of (1c). To explain the general collaborative patterns described above, the causal relation between collaboration and successfulness needs to be general and robust. Hence, one needs to 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 theorem (see the appendix for the proof).

Theorem. When m groups of equal size k merge, the individual successfulness of agents increases.

In other words, as soon as several k-groups of the same size exist, they would improve the individual successfulness of their members by merging. A corollary is that single individuals always have interest in collaborating. However, this theorem only covers a small subset of possible configurations, and cannot provide a general vindication for the causality claim (1c). Further, agents might only use it if they are aware of it and are in a position to identify groups of equal-size competitors, which cannot be assumed in general.

To overcome these difficulties, we now assess agents' successfulness irrespective of what they know about other competitors: we consider the average successfulness of k-groups over all possible configurations for each community size. For example, we average the individual successfulness of 4-groups in configurations (4-1-1-1); (4-2-1) and (4-3)[1]. In order to study the robustness of the causal relation between collaboration and successfulness, we investigate in the next paragraphs how much collaborating remains beneficial under variations of key parameters of the competition context.

Successfulness and community size. Figure 1 shows the average successfulness within k-groups for communities of various sizes. First, the successfulness of loners brutally collapses and is much lower than that of other k-groups as soon as $n > 2$. This confirms that except when nobody collaborates, or in very small communities, loners are outraced. Second, for all group sizes, individual successfulness decreases for larger communities, as can be expected when the number of competing groups and their size increases. Nevertheless, the successfulness of k-groups remains high and stable up to some community size s larger than k till they are eventually outperformed by larger groups or till growing bigger would mean over-collaborating (see (Boyer-Kassem and Imbert, 2015, 679-80) for an analysis of over-collaboration in large groups). Third, the larger the groups are, the longer and flatter this initial plate of successfulness is and the less steep the decrease in successfulness is. Fourth,

[1]There is no clear rationale about how to weigh configurations. From a combinatorial viewpoint, configuration (1,1,1,1,1,1,1) has one realization and (3,2,1,1) several ones. But from an empirical viewpoint, when scientists hardly collaborate, configuration (1,1,1,1,1,1,1) is usual and (3,2,1,1) extremely rare. We have privileged simplicity and chosen to give equal weight to all configurations.

when n is much larger than k, the successfulness of k-groups increases with k. However, this increase is a moderate one and small groups still do reasonably well, which is somewhat unexpected, given the general amplification effect — but see the analysis of figure 3 below for more refined analyses. Typically, in 10-communities, 2-groups do badly but remain somewhat viable since their average successfulness remains between 1/3 to 1/2 of that of 3 or 4-groups. Overall, not collaborating is in general not a viable strategy. Collaborating moderately ($k = 2$ or 3) can be very rewarding when there are few competitors (e.g. in small research communities, or on ground-breaking questions that are only known to a handful of scientits). Smalls groups remain viable but tend to be outraced when communities become significantly larger (typically, concerning questions belonging to normal science that many researchers are likely to tackle). Thus, moderately collaborating is a viable but more risky strategy when uncertainty prevails about the number and size of competing groups. Finally, while large collaborative groups rarely get exceptionally high gains, they are extremely safe, with moderate differences in successfulness between them or when the community size increases.

Successfulness and group size. Figure 2 shows the variation of individual successfulness with group size for various community sizes. First, for $n > 2$, the successfulness curve has a one-peaked (discrete) form, the maximum of which grows with the community size. Second, these one-peaked curves are not symmetric: the increase in successfulness is steep (but less so for larger groups), the decrease is gradual (idem). Large groups predate resources so groups need to grow big quickly to get some share and because returns can be increasing (Boyer-Kassem and Imbert 2015, 678), the increase in successfulness is steep. The decrease after the peak is slow because large group are hard to predate but over-collaborating can become suboptimal when the increase in gain by predation no longer makes up for the need to share between more people). These results are not trivial because at the configuration level, the successfulness of groups is contextual. They are important, too. A one-peaked profile is usually *assumed* in the literature about coalitions. Here, it emerges from a micro-model, and gets its justification from it. Overall, these patterns show again that agents have a large incentive to collaborate substantially, whatever the competing environment.

Successfulness in more or less collaborative communities. Figure 3 finally shows how the successfulness of k-group members varies with the degree of collaboration in their competition environment.[2] Here again, what matters

[2]Here, the degree of collaboration in each configuration is assessed by computing the average size of k-groups. For each k, we then compute the average successfulness of a member of a k-group over configurations having a degree of collaboration within intervals $[1, 1.5]$ (represented at coordinate "1.25" on the x-axis), $[1.25, 1.75]$, $[1.5, 2]$ ... $[3.5, 4]$. We

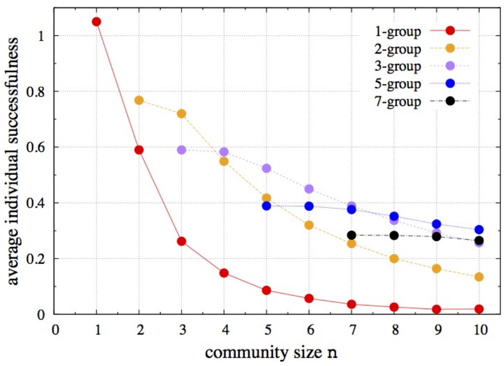

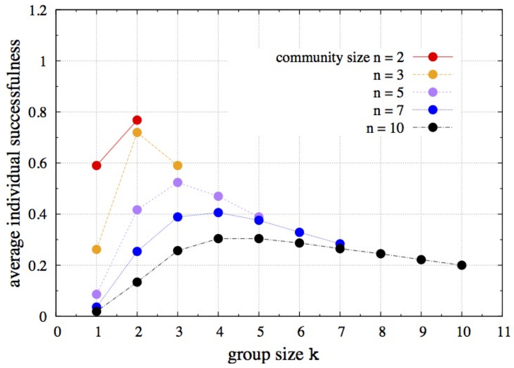

Figure 1: Variation of individual successfulness with community size.

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

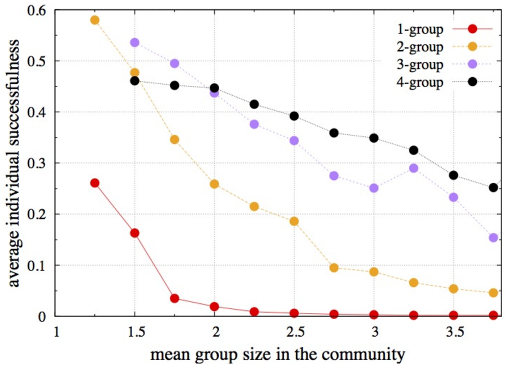

Figure 3: Variation of successfulness with the degree of collaboration in communities.

is less the exact value of the successfulness than the differential successfulness between more or less collaborating individuals. The graph confirms that successfulness depends less on the absolute size of groups than on how much they collaborate in comparison with their competitors. Scientists who collaborate more than average are very successful; those who collaborate as their peers do reasonably well; those that collaborate less than average are outraced by a large margin. This general result is not unexpected given all the above results, but the graph highlights that success for intensively collaborating scientists, and underachievement for under-collaborators can be very large. This is an important finding because if, as we shall see, successful scientists pass over their collaborative habits more than their peers, then the feedback loop provides a mechanism that favors the *increase* of the degree of collaboration by promoting those that collaborate more than others.

have chosen overlapping intervals to smoothen results. The average is computed up to communities of size 10.

Partial conclusion. Overall, the results show that — everything else being equal — collaborating a lot 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 much more successful. Collaborating too much is not a significant problem, under-collaborating is. So, 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, (1c) seems adequately supported.

5 Collaborative Practices Develop Because of the Success of Collaborative Scientists

We have so far argued that collaborative scientists, especially when they collaborate more than others, are more successful. We now need to argue that, because of this differential successfulness, collaborative habits persist and possibly develop in scientific communities (2c). A wide variety of social mechanisms across scientific contexts can contribute to this feedback loop. Accordingly, we shall be content with giving various evidence that strongly suggests that this link is a likely one.

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 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 can be passed over and need not be reinvented by newcomers.

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 the collaboration rate will therefore increase. 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. Within pure science, the connection is less straightforward. But 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, and pass over their collaborative habits.

Note that it is merely needed that the function between the pragmatic rewards of scientists and their success is on average increasing. This remains compatible with the fact that *some* epistemically successful scientists get little resource and *some* unsuccessful scientists get a lot — which seems to be the case. Actually, non-epistemic factors may even tend to over-credit successful scientists, and in particular collaborative ones. First, individual successfulness has been assessed in the model with a conservative estimate. It seems that an agent's publication within a k-group 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). Also, a publication within a 10-group will generally be more visible than one single-authored publication, since more people can promote or publicize collective publications and research topics. Second, 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). 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 [3]. Once winners of the scientific race publish co-authored articles, it becomes easy for others to see that successful scientists are highly collaborative ones. (For instance, if agents of a 3-group are 4 times more successful than a single agent, this means that their groups publishes 12 more articles than this agent). Accordingly, the belief that collaborating is beneficial can be acquired as collaborating becomes usual. 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

[3] Kincaid mentions that "complex combinations of intentional action, unintended consequences of intentional action, and differential survival of social practices might likewise make these conditions [(1)–(3) in our Section 3] true" (Kincaid 1996, 112).

and their members will be more successful, may attract resource, and will keep developing and transmitting their working habits.

In light of the above discussion, we believe that the causal connection between the success of collaborative scientists and the persistence and development of collaborative practices is highly plausible.

6 Discussion

Good functional explanations should be unambiguous about when the causal mechanisms that they rely on are efficient. In the present case, the following conditions can be emphasized.

First, conditions for the application of the priority rule should be met. In particular, (i) it should be possible to single out problems and to state uncontroversially when they are solved. Second, for the model to apply, (ii) scientific problems should be dividable into subtasks, and (iii) the solutions of these subtasks should be communicable. Finally, the model assumes that (iv) the completion of these subtasks should be sequential, but our conclusions still hold if this condition is relaxed. Indeed, if some subtasks can be tackled in parallel then the project can be completed even more quickly by different agents of a group, and collaboration is even more successful. Conditions (i)-(iii) are somewhat met in the formal and empirical sciences, less so in the social science, and almost not in the humanities. For example, as noted by Thagard (1997, 249), 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. This may account for the differences concerning collaborative patterns in these fields.

As mentioned above, different causal pathways may connect the successfulness of collaborative scientists to the persistence and development of collaborative practices. Thus, conditions for the fulfillment of claim (2c) cannot be uniquely specified. But several points are worth mentioning. First, the activity of epistemically successful scientists should be favored by scientific institutions. This can be the case if it is agreed that scientific success, in the form of publications or patents, is valued and promoted. Concerning scientific results that lead to patents, applications and financial gains, this condition is met when public or private funders value such outputs. Concerning pure scientific results, this means that there should be a wide agreement about which results are scientifically good and significant, and there should exist common and accessible publication venues, the value of which is consensual. Again, these conditions are approximately met in the formal and empirical sciences, less so in the social science and, almost not in the humanities in which scholars do not share paradigms, methods or norms about what is scientifically sound and significant, and cultural and linguistic barriers can restrain the existence of unified communities and common publication venues. Second, in contexts in which researchers and projects are regularly evaluated, especially by agents or institutions who are not in a position to asses the scientific value of their work, the existence of a common standard of success in terms of publications (through simple and calibrated publication indicators) may even more favor researchers who are successful, and therefore the development of collaboration. Finally, when resources are crucial to carry out or facilitate research, snowball effects can favor even more successful scientists, and in particular collaborative ones. This resource accessibility condition, which is central in Wray's explanation, is not in ours. But we agree that in such cases, the functional mechanisms that we describe will be even stronger. In this sense, our account encompasses Wray's. This condition about resources may be another reason for the difference in collaborative behavior between the formal or empirical sciences, the social sciences and the humanities.

7 Conclusion

We have argued that collaborating a lot is overall a safe and success-conducing practice. This conclusion is robust for various sizes of groups, communities and degrees of collaboration; everything being equal, those who collaborate more than average do better. Then, to the extent that the successfulness of researchers gives them more opportunities to transmit their research habits, the development of collaborative practices in communities can be functionally explained. We have further emphasized that the conditions for this functional pattern to work are specifically met in the scientific fields in which collaboration is well-developed. Accordingly, it seems reasonable to consider that this functional mechanism is an important element of the explanation of the development of collaboration in modern science.

The explanation of collaboration is probably a multi-factorial issue. Nevertheless, an asset of our general functional explanation is that it highlights the unexpected force of beneficial aspects of collaborative activities and suggests important roles for contextual factors that are associated with the rise of collaboration. As such, it is general and unifying. For instance, the competition model shows how the division of scientific labor, the use of specialized experts (Muldoon 2017), or the increased reliability of collaborative teams (Fallis 2006, 200) can increase the probability that groups pass research steps and have amplified effects in terms of successfulness. Similarly, factors like the need to access resources to carry out or facilitate research can create a snowball effect that favors epistemically successful (collaborative) researchers (Wray 2002). And factors like the globalization of research or professionalization (Beaver, 1979) can be seen as conditions favoring the application of the priority rule and scientific competition.

Finally, while nothing in the model provides an internal limit to the growth of collaboration, one can note that 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.

8 Appendix: Proof of the Theorem

Consider first the simple case where the m k-groups 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 km-group will be at least as quick, and sometimes more, than all k-groups — more precisely: for a given drawing of all random variables corresponding to attempts to pass the steps, for all agents and temporal intervals, the km-group will move at least as quickly as all k-groups. 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. 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 km-group is sometimes quicker than the m k-groups, 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 (1979) "Studies in Scientific Collaboration: Part III", *Scientometrics*, 1(3): 231-245.

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.

INSERM (2005), "Les indicateurs bibliométriques à l'INSERM", https://www.eva2.inserm.fr/EVA/jsp/Bibliometrie/Doc/Indicateurs/Indicateurs_bibliometriqu Inserm.pdf

Kincaid, Harold (1996), *Philosophical Foundations of the Social Sciences*, Cambridge University Press.

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.

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.

Wuchty, Stefan, Jones, Benjamin F. and Uzzi, Brian (2007), "The Increasing Dominance of Teams in Production of Knowledge", *Science* 316(5827): 1036-1039.