# THE NATURE OF CO-AUTHORSHIP.
## A NOTE ON RECOGNITION SHARING AND SCIENTIFIC ARGUMENTATION

**Jesús Zamora-Bonilla**
**UNED, Madrid**

**Forthcoming in *Synthèse***

**ABSTRACT**

Co-authorship of papers is very common in most areas of science, and it has increased as the complexity of research has strengthened the need for scientific collaboration. But the fact that papers have more than an author tends to complicate the attribution of merit to individual scientists. I argue that collaboration does not necessarily entail co-authorship, but that in many cases the latter is an option that individual authors might not choose, at least in principle: each author might publish in a separate way her own contribution to the collaborative project in which she has taken part, or papers could explicitly state what the contribution of each individual author has been. I ask, hence, why it is that scientists prefer to 'pool' their contributions instead of keeping them separate, if what they pursue in their professional careers (besides epistemic goals) is individual recognition. My answer is based on the view of the scientific paper as a piece of argumentation, following an inferentialist approach to scientific knowledge. A few empirical predictions from the model presented here are suggested in the conclusions.

**KEYWORDS**

Co-authorship, scientific collaboration, reputation, social epistemology, inferentialism, institutions, intellectual property, argumentation.

## 1. INTRODUCTION

Co-authorship of scientific papers has been extensively studied from the sociological and the bibliometric points of view, as well as from research policy studies and research ethics; [1] it has received, however, considerably less attention from philosophers of science and epistemologists, perhaps under the common (though not often expressed) assumption that these disciplines are more concerned with the content, quality, validity and nature of the *items* of scientific knowledge, than with something apparently so accidental as to whether the proof of a theorem or the performance of an experiment is carried out by one single person or by a group. Curiously, this general lack of attention in philosophy of science to the phenomenon of co-authorship co-exists, nevertheless, with a progressive recognition by many philosophers and epistemologists that scientific knowledge, and knowledge in general, can be seen as a kind of social fact (independently of whether they think that this entails that science is 'not objective', or that it is objective *in spite* of being social, or that it is objective *thanks* to their being social). This recent 'socialisation' trend in the epistemology of science[2] may have led, however, to the (certainly right) conclusion that a scientific paper does not become 'more' social by having been created by a group of people rather than by a single individual, and hence, the study of the social aspects of scientific knowledge have concentrated on other, more 'community-wide' types of social relations, rather than on the specific face-to-face interactions that are constitutive of co-authorship.

Some of the few exceptions of this trend are the papers by Wray (2002, 2006), Thagard (2006), Fallis (2006), and perhaps Rolin (2010) (who does not explicitly refer to co-authorship, but studies 'research groups' as engaged in the production of singular claims, and hence we can interpret each of these groups as the team of researchers writing a single paper).[3] These papers, however, don't even consider the question of why scientists *write* articles together; certainly, they deal with some closely related questions: why scientists engage in collaborative research (Wray, Thagard, Fallis), and

[1] See, e.g., Beaver (2001), Biagioli (1999), Birnholtz (2006), Chomalov *et al*. (2002), Glaenzel and Schubert (2004), and Wuchty *et al.* (2007).

[2] For a recent survey of social epistemology, see Goldman and Whitcomb (2010). Some other prominent examples in the application of social thinking to philosophy of science are Steve Fuller, David Hull, Philip Kitcher, Helen Longino, Joseph Rouse or Miriam Solomon.

[3] Wagenknecht (2010) also presents an interesting empirical study of the relation of epistemic trust and joint commitment among collaborators.

how they combine their individual points of view once they have decided to present in a unified way the results of their research (Rolin); but it is conceivable that, even if those research processes are collaborative, each scientist would nevertheless publish *individually* her own conclusions or her own portion of the research, even as separate parts of a single paper. Why, instead of this, do the multiple authors of a scientific paper prefer that the conclusions they derive from their research process are *taken* by their other colleagues as a *'corporate claim'*? After all, as Wray rightly (2006) signals, and as many scientists and science managers have largely complained about, co-authorship tends (besides all its possible advantages) to blur the epistemic responsibility of the individual authors, and it is this individual responsibility what constitutes the basis of merit allocation in science. This is a serious problem *per se*, even if co-authorship were always accomplished in a non-problematic way from the ethical point of view (which, as the proliferation of disciplinary norms about the topic testifies, is far from being always so). So, why not to use in scientific publications something like the system of 'movie credits', indicating at the end of the paper who did exactly what (even if some things were necessarily done by more than one person)? I'll try to offer in the rest of the paper an explanation of why co-authorship (in the sense of *collective commitment to the claims of a paper, instead of separate assignments of epistemic responsibility to each different co-author*) can be a rational strategy for scientists. I will do that by framing the problem within the inferentialist approach to scientific knowledge I have been elaborating in a series of other papers (section 2). In section 3 I shall relate the problem, as described in section 2, with the economic study of the demarcation of property rights and the limits of corporate economic agents. Section 4 will describe several types of scientific co-authorship, and will identify the one to which my argument (elaborated mainly in section 5) refers in the clearest way. Section 6 will offer some concluding philosophical comments.

## 2. INFERENTIALISM AS A THEORY OF KNOWLEDGE.

Inferentialism, being a theory closely related to conceptual role semantics, has been mainly developed within philosophy of language and philosophy of logic, its most systematic exposition being Brandom (1994). Its most relevant claim, for the purposes of this paper, is that for a mental state, a series of noises, a number of marks, or a configuration of magnetic fields *to count as a propositional claim*, it is necessary that an agent takes *responsibility* for the inferential links to which the claim belongs; i.e., one must be able of providing premises from which the claim follows, as reasons for the acceptance of the claim, and she also must be able of inferring what other facts follow from the claim and should be accepted if the former was. To have a belief would consist, hence, in being committed to subject the content of your belief to 'the game of giving and asking for reasons' (in Sellars' terms), a game whose rules are *inferential norms*, i.e., norms telling what things follow from other things or are incompatible with them. I have employed in the last sentence the ambiguous word 'things' to point to the fact that, according to Brandom, inferential norms do not only connect claims-to-claims, but the 'game of reasoning' may include also non-verbal inputs and outputs: under appropriate circumstances *perception* can justify the undertaking of doxastic commitments (i.e., facts that one accepts), and claims can justify the attribution of practical commitments (i.e., *actions* one has to do or has to avoid). I have argued elsewhere[4] that this essential connection of claims to perception and action justify the attempt to transform inferentialism into a general theory of rationality, i.e., one including both knowledge and practical action within its scope. According to this approach, the possession of knowledge would consist in the undertaking of claims that can be justified with the help of the inferential norms accepted by the members of the relevant epistemic community; one agent's belief counts as knowledge for a community just in the case that she or other relevant members are able of tracking a *justificatory* chain of inferential links which starts from claims that are already accepted by the community and ends in that belief, and hence, to treat a claim *as knowledge* consists in taking it as something that can be *legitimately* (i.e., following acceptable norms of justification) transmitted *as a commitment* to the other members of the community (i.e, there is a legitimate way of making others being committed to the claim), whereas a 'mere belief' would consist in a commitments an individual has, but such that she cannot find a legitimate way of transmitting it to others as a commitment. The crucial element in this process is, of course, the system of inferential rules accepted by the community members: epistemology would consist in the task of finding out what inferential rules are better for the goals the agents (goals that will essentially include epistemic ones, but probably also other motivations), and it is surely an activity in which the agents themselves can be actively engaged every time they find that

[4] Zamora Bonilla (2011).

following a particular set of rules leads them to having commitments that are problematic *for them* in some way or another, for example, if they clash with other incompatible commitments they have, or if they command them to perform some actions they have some other legitimate reasons not to want to perform.[5]

Regarding the production of scientific knowledge, inferentialism leads to what I have called the 'persuasion model' (Zamora Bonilla, 2006, 2007), which combines the ideas I have just sketched with well known facts in the sociology of science. According to this model, researchers pursue, amongst other things, *recognition for having made important discoveries*, i.e., a scientist wants that her colleagues explicitly employ in their own researches the results advanced by the her,[6] for recognition comes mainly through citations, and the relevance of a discovery is normally associated to how useful it is in helping to make other discoveries. Note that my claim is not that scientists pursue recognition 'instead' of truth, but that, in practice, both goals constitute for them a single end, for it is not 'mere' recognition what researchers pursue, but the recognition *of* having contributed to the furthering of the epistemic goals incorporated in the justification rules they have to employ in order to persuade their colleagues of the validity and relevance of their discoveries. A comparison may be useful here: we can model the behaviour of professional athletes by assuming that they pursue the goal of being recognised *as* good athletes; they don't choose to, say, jump as high as they can where nobody sees them, but, instead, they choose to do it in a public and socially regimented way, that warrants both fair competition and public recording, and it can be argued that the records obtained under this regimentation will be better, on average, than the ones obtained in a purely amateur and private way. As the norms of athletics are designed to promote the attainment of the highest possible records through public competition, the norms of science can similarly be conceived as 'designed' (or evolved, or a combination of both things) in order to promote the attainment of results of the

[5] One common complaint about this view of knowledge is that it does not explicitly refer to the idea of truth or objectivity; actually, it does it *implicitly*, as far as being committed to true claims, and to inferential norms that are truth-preserving and truth-tracking, is an epistemic goal for the *agents* themselves, though it is up to them (i.e., up to each particular community defined by some shared system of inferential norms) the task of finding out which ones those norms are, and what is their connection with epistemic and other types of goals, as well as what is the relative importance of the former as compared to that of the later. See Zamora Bonilla (2010) for a justification of the role of rationality and objectivity in the inferentialist approach to scientific knowledge.

[6] See Hull (1988) for a convincing argument about this point.

highest possible epistemic quality *through* the fairest possible distribution of scientific merit amongst researchers.

According to the inferentialist approach, the most important part of the 'social regimentation' of scientific competition is the system of norms that determine what is considered as a valid scientific argument, i.e., what are *good reasons* for the acceptance of a particular claim. My suggestion is to model the scientific paper as an *argument*, in which the author attempts to establish a *conclusion*, that will be her main claim, and that is supported by *premises*, some of which are results of the author's own observations, experiments, etc., and some are taken from other pre-existing papers, i.e., some of *your* premises will be the conclusions of *other* authors' papers. The inferential norms accepted in your research community will tell what relations between claims make it legitimate to take some claims as *supporting* another, they will also tell when a claim has received so much support that its acceptance as a premise is *permitted*, and even when it is *compulsory* to accept it.[7] In order to reach a high score in this game, a scientist will be forced to find out the best arguments she can in favour of the claims for which she wants to be recognised, and the norms governing what a 'good argument' is will lead her to select some claims of other scientists as premises, increasing in this way the recognition got by these researchers.[8]

One central feature of this model is that claims are recognised according to authorship; i.e., one is rewarded for those claims[9] of which she is the legitimate author, and hence it is important that the colleagues, as well as science managers, are able of attributing in an efficient and fair way what claims and how much merit correspond to each individual researcher. Why, then, to make the application of scoring rules more

[7] See Zamora Bonilla (2002).

[8] The argumentative character of the scientific paper, and my talk about 'premises' and 'conclusions', must not be confused with the stronger claim (which I do not endorse) that all scientific papers have something like a 'deductive' structure: deductive arguments are just one type of admissible arguments in science. What is important in the model defended here is the possibility of distinguishing within each paper the claims for which the authors want to be *recognised*, and the claims they employ as *supporting* the former. Take also into account that the fortune of a claim will hardly depend only on the support it receives from its author's own works, but also from arguments and discussions provided by other scientists (e.g., Cavendish's measure of the gravitation constant added support to the principles of Newton's mechanics; Minkowski's geometrical study of space-time added support to Einstein's special relativity, etc.).

[9] It is not essential that all scientist's contributions have the linguistic form of a proposition; they can also be experimental or abstract technics, classifications, and even questions, or, of course, arguments themselves.

difficult by sharing the authorship (i.e,. the scientific ‘property’) of a claim amongst a group of authors, instead of clearly stating what part of the merit corresponds to each author, what has been the contribution of each one to the construction of the argument? As we shall see in the remaining sections, though this seems to pose a delicate question for my inferentialist model, in the end it will show to be a nice case for illuminating the nature of scientific arguments and scientific knowledge. Note that this is not only a problem from the point of view of inferentialism, but for all theories about the motivations of scientists in which individual recognition is a relevant goal, no matter how minor it is regarding other goals, because, in principle, nothing of epistemic nature would be lost if scientists just opted systematically for stating in an explicit way the contribution of each individual author in cases of collective research (as it is actually done in some cases). Lastly, we can argue that the mere existence in many scientific disciplines of explicit codes of conduct regarding the handling of co-authorship is sufficient to conclude that the distribution of scientific merit in those cases is not only an internal problem of a particular philosophical or sociological paradigm, but a serious matter of concern for flesh and bone scientists as well.

## 3. AN ECONOMIC DIVERSION: THE NATURE OF THE FIRM.

Let me leave aside for a moment our talk about scientists, papers and arguments, and bring up an apparently unconnected topic: Ronald Coase’s economic theory about the nature of the firm and the allocation of property rights. This theory also starts up from what is a seemingly paradox; in principle, nothing could seem more essential to the capitalist system than private firms and the free market, but the truth is that the private enterprise actually constitutes a *complete suspension* of the market mechanism, for, while firms compete against each other in buying and selling their products to other firms, consumers or institutions, and this competitive process *is* the free market, no doubt the market rules cease to be valid *within* the firm, which is internally governed not by ‘the free encounter of demand and supply’, but by sheer *authority* relationships between managers and employees, bosses and workers. The boss doesn’t go to the assembly line and asks for how much cents are each worker willing to fix the next screw, and then buys the next item of screw-fixing to the worker who accepts a lower price. What the workers have really sold is their time and competence, that will be used by the managers (within the limits of the contract and the laws) according to what the latter decide to order at every moment. But, as Coase argued, if it were true that free market competition guarantees that the best possible product is sold at the lowest possible price, it would be irrational to produce something by means of a system of authority relations, instead of leaving *every* part of the process to the free market. After all, corporations buy some things in the market, use them and transform them into others within their own physical and institutional limits, and sell these products in the market again; and the crucial fact is that *every* one of the processes that take place within the firm, could be *bought* instead, i.e., could be 'externalised' (for example, instead of having employees devoted to making tyres, a car manufacturer can buy the tyres to other firms). The question is, if the free market is so efficient, why do firms not buy in the market *everything* they need? Why do they not externalise their *whole* production process? Which is equivalent, obviously, to ask the following: why do firms exist at all, instead of the capitalist economy being simply an aggregate of 'atomic' market exchanges?

Ronald Coase, in his 1937 epoch making paper "The Nature of the Firm", pointed to *transaction costs* as a reasonable explanation of this paradox. Though market exchange tends, thanks to competition, to minimise the costs of production of goods and services, the process of getting producers and users in touch, the dissemination of information about the items, and even the process of deciding what to purchase and to whom, are costly in themselves; the concept of 'transaction cost' covers all these heterogeneous types of costs not associated with production itself. And, as it happens, when taking into account not only the cost of producing a good by an external agent, but also the costs of finding out the best offer, negotiating the purchase, and so on, the firm can discover that carrying out by itself the production of that good is cheaper than buying it. This theory, with all its developments in the decades following the publication of Coase's paper, not only explains the existence of firms, but also their limits, i.e, their *size*. For when transaction costs for some goods or services decrease, it will be better for a firm to externalise their production, and hence become 'smaller' (i.e, needing less workforce, or contributing with less added value). In a later paper (Coase, 1960), Ronald Coase provided a similar analysis for the explanation of the limits of property rights in general: in the absence of transaction costs, how are property rights allocated would be irrelevant, in the sense that the agents would always exchange those rights in such a way that the one who can make the most profitable economic use of one right ends buying it; but when the exchange of rights is costly in itself, then some distributions of rights become economically more efficient than others.

What has this story about private firms to do with our argument about co-authorship? I suggest that, though there is not an exact analogy between the two cases, there are two important similarities. In the first place, in both cases we have a distribution mechanism (the free market, and the priority rule for scientific recognition, respectively) that seem to work apparently well, but that are 'blurred' by some institutional arrangements (private firms, and co-authorship) that intuitively seem to go against the working of those mechanisms. In the second place, as I'll try to show in the remaining sections, the case of co-authorship can be convincingly explained by reference to the benefits brought about by a wise distribution of (intellectual) property rights.

## 4. REASONS TO WRITE TOGETHER.

Let's examine some of the reasons a group of researchers may have to present the conclusions of their research as a collective claim, instead of as an aggregate of individual, separable contributions. Of course, they might just be lying, in the sense that one or more of the co-authors have not contributed in any relevant way to the paper's content, but are included for other illegitimate reasons: exchange of favours, academic promotion of a friend, or just sheer abuse of power. These types of co-authorship are rightly condemned by all scientific codes of conduct, but their commission is easily explained when there is a high probability of the misconduct passing unnoticed. A particular form of cheating would consist in a mere exchange of papers: you write one paper by your own, I do the same with another paper, but we signed *both* papers as if they had been the result of collaborative research between us. This is obviously another example of scientific misbehaviour, but I doubt it is as frequent as the former cases are,[10] not because it is more likely to be detected, but because it is probably not as profitable for authors at it might seem at first sight. If the merit an individual gets from co-authoring two papers with another individual were equal to the merit from single-authoring one paper, then they wouldn't gain anything from cheating in this way; and if we take into account the chances of being discovered and punished, then this strategy

[10] Of course, because of the nature of these actions, we don't have empirical data enough to assess actual frequencies.

becomes clearly inefficient. It would only be 'rational' if getting two co-authored papers gave substantially *more* recognition to an individual than having one single-author paper, and so, the fact that this kind of behaviour is not too common seems to point to the conclusion that scientific merit *decreases more than linearly* with the number of co-authors, i.e., single authoring one paper gives more recognition than co-authoring *n* papers with *n–1* colleagues .[11]

Legitimate co-authorship corresponds, obviously, to the cases in which there is real collaboration between the co-authors. But we can divide this possibility into two different cases. First, a group of collaborating scientists may write a single collective paper simply because there is no clear enough way of separating the contributions of each researcher to the argument supporting the paper's conclusions. Since ideas are discussed amongst all the members of the group, continuously examining them and proposing new ones on the face of the arising problems, often it can certainly be the case that all the elements of the paper contain a significant contribution from every individual member of the team. After all, even a single authored paper that in its final version is the paradigm of clarity and precision can be the result of a messy process of trial and error and of dozens of muddy drafts, a process the author may be unable of reconstructing after the work is done. I assume that many co-authored papers will fall into this category, in particular those that are written by teams (usually couples) that work together for many years and that produce a large amount of collective research (one can think, for example, in pairs like the one formed by psychologists Amos Tversky and Daniel Kahneman). Co-authorship in such cases is really unavoidable, even psychologically enriching to authors, and there would be nothing left to explain about its rationality.

Second, there is the possibility that a group of researchers collaborate and write a coherent article, but that in principle they could 'factor out' the contributions made by each member. This is particularly the case when collaboration is mainly justified by the technical and cognitive division of labour amongst people with different competences or from different areas, but also amongst colleagues in the same field that simply concentrate on a different part of the paper though they discuss together the general argument at the beginning, and its final versions at the end of the process. What does it

[11] Another reason deterring scientists from this form of fraud might be that at least one of the fake co-authors thinks that her own paper deserves more merit that those of the others.

justify in this case that the paper is signed by several co-authors, instead of each one claiming merit *only* for the part to which she has mostly contributed? After all, what some of them are doing is using the conclusions reached by the others as premises in a (sub-)argument. Assuming that the goal of the authors of those intermediate conclusions is to get cited, why not write a separate paper for each of those conclusions, one that will be *cited* by the colleagues that are writing the rest of the argument? Stated in the language of section 3, why are not the fragments of the argument *'externalised'* as far as possible, giving to each of those fragments the possibility of being autonomously recognised? This case, which I will call 'optional co-authorship', is the one that we had to explain.

## 5. OWNING SCIENTIFIC CLAIMS.

There is a relatively obvious (but, as I shall show, partial) answer to our last question: perhaps the 'fragments' or 'sub-arguments' of a complete paper are not individually publishable, i.e., perhaps they are not *relevant enough* to deserve publication according to the views of journal editors, and so the authors have no other option but to join their own sub-arguments into a bigger, and more relevant paper. It is important to note that it is a matter of mere *relevance*, not of epistemic *validity*, for we are assuming that the complete paper would be acceptable for publication, and so we can suppose that all its part are scientifically sound. But this answer is only partial because it takes as given the 'threshold of relevance' journal editors have established, and this threshold is just a convention that might have been different. After all, there is no natural line that can serve to 'cut arguments at their joints', no place where we can absolutely assert that an argument begins or ends, since its premises must also receive support from previous premises, and from its conclusions further conclusions might be derived. The question is that, considered from a *global* point of view, any research field can be seen as one single 'long argument', one that is massively entangled, certainly, with premises entering from countless places and nodes ramifying almost without end. The 'state of knowledge' within a discipline at a particular moment would just consist in the state of validation of every single link and node in this gargantuan network. Each scientific paper contributes to this network with just a few nodes and links, but, obviously, always with more than just *one* new node or link. The question is, *what determines whether a particular fragment of the discipline's 'one long argument' is* *relevant enough to gain the right of being published as one 'autonomous' piece of argumentation?*

My guess is that this basically depends on two things. In the first place, the technology that is used to communicate the arguments, and the costs (particularly the cognitive ones) associated to it, can entail that some argument 'sizes' become more appropriate than others for communication and storage of scientific information. In the second place, and more importantly, scientists pursue being recognised for making important discoveries, and the most approximate measure of that importance is the number of times the discovery is employed in further research by other scientists; so, the main goal of a researcher in building up an argument is not only that each step she adds is well grounded, but that at least one of those steps *becomes a premise in other papers written by colleagues*, the more the better. So, whereas from the logical point of view an 'intermediate' conclusion is, considered as a step in an argument, equally a 'conclusion' than the other steps, the author or authors of a paper will not consider its possible 'fragments' (i.e., sub-arguments leading to 'intermediate' conclusions) as the particular *piece* by which they want to be recognised, even if each of these 'fragments' could be given the structure of a 'complete' argument (i.e., of a 'separate', but shorter paper). The authors are pursuing a *specific* conclusion, the one they hope will be more useful for other scientists. It is true that, in principle, each 'fragment' of a single paper might be published independently, and cited at least in the other 'fragments' of the same research, but the conclusions of these 'fragments' would very likely not give their authors as much recognition as the 'big' discovery itself.

Let's consider this situation through an extremely simplified example. Suppose that you and I are collaborating in a research, in which we are trying to find out a good answer to an important scientific problem. We provide a compelling argument showing that Q is that answer. Our argument can be decomposed in two sub-arguments, one showing that P is the case, and another showing that P entails Q. Besides some informal chatter in the bar, during which we have exchanged some ideas, you have basically proved P by your own, and I have done the same with $P \rightarrow Q$. In principle, we might publish our discoveries independently, even appropriately citing the other's result in the references of our own papers, but this would be silly if Q is the claim that our colleagues are waiting to be proved in order to use *it* in their own works.

It might be argued that, if what our colleagues want is the 'big important problem' to be solved, they shouldn't care if there is a single paper in which Q is established, or if they have to use separately the premises P and P$\rightarrow$Q, for, after all, Q follows from that. If they need Q, and you and I have independently proved and published P and P$\rightarrow$Q, our colleagues will cite both papers, won't they? The truth is that not necessarily. For it could be the case that *other* scientist proves a different couple of premises, R and R$\rightarrow$Q, and the colleagues needing Q might choose either to cite our papers or the other one. So, *we decide to publish together our combined argument proving Q in order to collectively own the 'intellectual property' (i.e., the priority right) of the claim Q*, a claim whose ownership we would not get if we had published in a separate way our discoveries of P and P $\rightarrow$ Q.

The decision to engage in optional co-authorship depends, hence, on the weighting of costs and benefits expected by recognition-seeking researchers from two different strategies: becoming the *sole* author of a 'smaller' discovery that may be cited by the colleagues you are collaborating with (and also by an unknown number of others), or waiting for receiving *partial* recognition from being the *co*-author of a 'bigger' discovery (which, besides, is uncertain because you cannot know for sure that your co-authors will be successful in their 'fragments'). The clearer it is for a group of researchers that this second option will result in a higher degree of recognition to all of them than the former, the more likely it will be that they become co-authors.

The connection of this mechanism and our inferentialist model or scientific research is clear: in the first place, the *pursuit of legitimate recognition through chains of acceptable arguments* is an important element of the decision process I have outlined. In the second place, and from my point of view more importantly, this decision process involves the consideration of scientific papers mainly as *arguments*, and the relevant claims of the papers as essentially being *conclusions* for which a justification has to be offered. And thirdly, it is mainly the expectation of recognition through the use *other scientists* will make of the conclusions of a paper, what mainly determines the value of each possible decision about the structure, length and content, and of course about the status of authorship, of the published arguments.

## 6. CONCLUDING REMARKS.

Co-authorship, when it is optional, gets its rationality from what we might call 'the magic of logic'. With this expression I refer to the (trival) fact that, in logic, the whole is more than the sum of its parts. Co-authors of a scientific paper contribute each with some *premises* or inferential *steps* to the argument the paper consists in, but what they collectively get in exchange is the ownership of the paper's *conclusion*. In the toy example given above, being one of the discoverers of Q (the real 'big' discovery) is more valuable than either being the discoverer of P or being the discoverer of P→Q. But this may lead us to ask the following: if the 'magic of logic' is so powerful, why not to enlarge the set of conclusions our paper offers? If someone else shows that from Q other important things follow (say, S), why not to enrol *her* in our team and publish a paper that proves Q&S? And why the authors of the premises we used to establish P (say, N and N → P) do not insist in becoming also part of the team? In all these cases, there is a tension between the importance of what can be proved, the size of the group involved, the time to wait till the discoveries are done, and the uncertainty of future discoveries. The first of these factors invites scientists to become co-authors, but the other factors make that decision costly and uncertain. Another important factor is that, usually, the results of the collaborating group are kept more or less secret till they are collectively published; if we have discovered Q and hope that this will help to make lots of new discoveries, we might try to 'seize' as bigger a fraction of the merit of those new discoveries, by simply not publishing Q yet, but letting the people that may use it enter our team. This would be a similar strategy to the one used by private firms in many industries. Fortunately, in the case of science is even more difficult than in industry to know in advance what the future discoveries will be and who has a real chance of making them. So, in the tension between becoming the member of a really big team that will make (with some probability) really big important discoveries, or just becoming the member of a smaller team (even perhaps a one-person team) that makes with more certainty a not so important discovery that can be published it in a much shorter time, it is reasonable that most scientists prefer the second strategy.

This last reflection can suggest some ways of empirically testing the model proposed here. Those factors that tend to create uncertainty about the future use of a discovery X will lead researchers to decide not to wait for finding out co-authors of possible papers in which X is used to make another discovery Y, and so, will tend to make co-authorship less frequent; so, in those contexts where the uncertainty about the use of discoveries is greater, co-authorship will be rarer. Some of these factors can be:

- the time taken for ideas to spread and circulate,
- the heterogeneity of scientific standards,

- the typical length of the arguments that are considered relevant and independently publishable.

I have no doubt that the empirical testing of these suggested connections would throw more light on the social and intellectual aspects of co-authorship, even if these modest ‘predictions’ became falsified, though casual knowledge of the main areas of academic research make these hypothetical claims more than plausible.

Lastly, it can be argued that my model does not particularly support the inferentialist view of knowledge sketched in section 2, but I think it is based on the assumption that scientific works are essentially *arguments*. Of course, one does not need to be a full-blown philosophical inferentialist in order to accept the latter claim, but it is difficult to defend the argumentative nature of scientific works without being committed to the thesis that acceptable patterns of inference are a crucial element in the constitution of a scientific discipline, nor to that these patterns have considerable normative force for the discipline’s members, and is this part of inferentialism the one I think is essential in studying these topics.

## REFERENCES


Beaver, D. (2001), “Reflections on scientific collaboration (and its study): past, present, and future”, *Scientometrics*, 52.3, 365-377.

Biagioli, M. (1999), “Aporias of Scientific Authorship: Credit and Responsibility in Contemporary Biomedicine”, in M., Biagioli, ed., *The Science Studies Reader*, New York, London, Routledge, pp. 12-31.

Birnholtz, J. (2006). What does it mean to be an author? The intersection of credit, contribution and collaboration in science. *Journal of the American Society for Information Science and Technology,* 57, 1758-1770.

Brandom, R. 1994, *Making it Explicit,* Cambrideg: Harvard University Press.

Chompalov, I., Genuth, J., & Shrum, W. (2002). The organization of scientific collaborations. *Research Policy,* 31, 749-767.

Coase, R., (1937), “The Nature of the Firm”. *Economica*, 4.16: 386–405.

Coase, R., (1960), “The Problem of Social Cost”, *Journal of Law and Economics,* 3, 1-44.

Fallis, D., (2006), “The epistemic costs and benefits of collaboration”, *The Southern Journal of Philosophy*, 44, 197-208.

Glaenzel, W., and A. Schubert (2004), “Analysing scientific networks through co-authorship”, in: Henk F., Moed et al., ed., *Handbook of Quantitative Science and Technology Research. The Use of Publication and Patent Statistics in Studies of S&T Systems,* Dordrecht, Boston, London, Kluwer Academic Publishers, pp. 257-276.

Goldman, A., and D. Whitcomb (2010), *Social Epistemology: Essential Readings*, Oxford, Oxford University Press.

Hull, D., 1988, *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science,* Chicago, The University of Chicago Press.

Rolin, K., (2010), “Group justification in science”, *Episteme*, 7.3, 215-231.

Thagard, P., (2006), “How to collaborate: procedural knowledge in the cooperative development of science”, *The Southern Journal of Philosophy*, 44, 177-196.

Wagenknecht, S. 2010, “Epistemic Trust: An Empirical Study in Natural Science”, communication to the *3rd Congress of the European Philosophy of Science Association*, Athens.

Wray, K. B. (2002), “The epistemic significance of collaborative research”, *Philosophy of Science*, 69.1, 150-168.

Wray, K. B. (2006), “Scientific authorship in the age of collaborative research”, *Studies in History and Philosophy of Science*, 37, 505-514.

Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science*, 316(5827), 1036–1039.

Zamora Bonilla, J. P., 2002, “Scientific Inference and the Pursuit of Fame: A Contractarian Approach”, *Philosophy of Science*, 69, 300-23.

Zamora Bonilla, J. P., 2006, “Science as a Persuasion Game”, *Episteme*, 2, 189-201.

Zamora Bonilla, J. P., 2007, “Science Studies and the Theory of Games”, *Perspectives on Science*, 14, 639-71.

Zamora Bonilla, J. P., 2010, “What games do scientists play? Rationality, objectivity, and the social construction of scientific knowledge.” *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association* (M. Suárez, ed.) Pp. 323-332. Springer. Amsterdam.

Zamora Bonilla, J. P., 2011, “Rationality in the social sciences: bridging the gap.” *The SAGE Handbook of the Philosophy of Social Science* (eds. Ian Jarvie and Jesús Zamora-Bonilla). Pp: 721-738. SAGE. London.