

Is a bird in the hand worth two in the bush? Or, whether scientists should publish intermediate results

Thomas Boyer

► **To cite this version:**

Thomas Boyer. Is a bird in the hand worth two in the bush? Or, whether scientists should publish intermediate results. *Synthese*, Springer Verlag (Germany), 2014, 191 (1), pp.17-35. 10.1007/s11229-012-0242-4 . halshs-01053681

HAL Id: halshs-01053681

<https://halshs.archives-ouvertes.fr/halshs-01053681>

Submitted on 5 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Be careful!

This version of the paper is different from the one actually published, and there are still some slight mistakes.

The last version of the paper can be accessed from <http://thomasboyerkassem.yolasite.com/>

Is a bird in the hand worth two in the bush? Or, whether scientists should publish intermediate results.

Thomas Boyer

Abstract — A part of the scientific literature consists of intermediate results within a longer project. Scientists often publish a first result in the course of their work, while aware that they should soon achieve a more advanced result from this preliminary result. Should they follow the proverb “a bird in the hand is worth two in the bush”, and publish any intermediate result they get? This is the normative question addressed in this paper. My aim is to clarify, to refine, and to assess informal arguments about the choice whether to publish intermediate results. To this end, I adopt a rational decision framework, supposing some utility or preferences, and I propose a formal model. The best publishing strategy turns out to depend on the research situation. In some simple circumstances, even selfish and short-minded scientists should publish their intermediate results, and should thus behave like their altruistic peers, i. e. like society would like them to behave. In other research situations, with inhomogeneous reward or difficulty profiles, the best strategy is opposite. These results suggest qualified philosophical morals.

1 Introduction

A part of the scientific literature consists of intermediate results within a longer project. Scientists often publish a first result in the course of their work, while aware that they should soon achieve a more advanced result from this preliminary result. Consider for example Schawlow and Townes in Summer 1958: they had just designed together the principle of functioning of what would be known as the laser, and they intended then to construct the first laser, relying on their plans¹. Should they publish now their idea of a laser, or rather profit by their advance on possible opponents and start the experimental work? They chose to publish their (theoretical) idea of the laser in a scientific journal, and they also filed a patent application. On the one hand, it was a valuable choice: the patent was granted and their paper was published, making them soon famous. Furthermore, another scientist called Gould was to file a similar patent application a few months later, in March 1959; it was refused, because of the priority of Schawlow and Townes. If the latter had not submitted their idea in 1958, the patent would certainly have been granted to Gould. In this sense, publishing their intermediate result was a worthy strategy. On the other hand, history showed that their publication had also some drawbacks for them. Schawlow and Townes’s publication triggered a race among many American laboratories for the first experimental realization of the laser. In addition to bringing them competitors, their publication helped these competitors in crucial ways. So much so that it was another scientist, Maiman, and not Schawlow and Townes, who was able to build the first working laser, in May 1960, and to publish the result. If Schawlow and Townes had not published their theoretical idea, there are chances that they would have been the first ones to build

¹The historical material of this example is drawn from Bromberg (1991).

a laser. Their publication of the first theoretical step certainly prevented them from being first to reach the second experimental step.

Was Schawlow and Townes's publishing choice optimal? For the scientific community as a whole, it seems that it was, since the theoretical result could be known sooner to all, and the first laser built quicker. But for Schawlow and Townes, it is not easy to tell, as the above arguments seem to pull in opposite directions. In the general case, where does the balance go? That is, should a scientist follow the proverb "a bird in the hand is worth two in the bush", and publish any intermediate result she gets? This is the question addressed in this paper. My aim is to clarify it, to provide a framework to investigate it and finally to draw some philosophical morals. I limit myself in this paper to the normative side of the problem, and not the descriptive one. In other words, I ask whether scientists *should* publish intermediate results, and not whether (or why) they actually *do* publish them. Although it seems to me that they indeed do more often than not, I don't argue for that here.

The question of whether intermediate results should be published falls within the field of social epistemology of science, which studies the collective dimension of knowledge and the design of institutions regarding their ability to promote epistemic progress. Its aim can be expressed as "identify[ing] the properties of epistemically well-designed social systems"². I can see at least three ways in which the question investigated here can be directly related to traditional questions in social epistemology of science. First, the question can be re-expressed as concerning a possible "mismatch between the demands of individual rationality and those of collective (or community) rationality"³. For example, a scientist might be better off not publishing her intermediate results so as to maximize her own scientific production, whereas the community would prefer that she published, for the sake of the progress of collective knowledge. If so, the epistemic progress of society would be slowed just because of the intermediate nature of some results. One of the aims of the paper is to establish whether this concern is justified. Another possible mismatch comes from the fact that non-epistemic motivations may conflict with epistemic ones. For instance, a scientist might get a better academic position if she published many papers instead of a bigger single one, even if doing so didn't bring any better scientific knowledge. Thirdly, the question about publishing intermediate results can be related to the reward system in science. According to the priority rule, only the first scientist who publishes a given result is rewarded. Is it a good rule for science? From the viewpoint of society, one argument in its favor is that it "gives scientists an incentive to publish their research as soon as possible, making the benefits promptly available to society"⁴. This seems fairly right for a scientist who reaches a final result — keeping it for herself brings her nothing. But suppose now the result is intermediate, i. e. it will be reused for the next step of research. Is the incentive to publish it still worth the risk of helping an opponent? In other words, the priority rule might lead scientists to distinguish between intermediate and final results, as far as publication is concerned. The intermediate feature of a result might be

²Kitcher (1993), p. 303.

³Kitcher (1990), p. 6.

⁴Strevens (2003), p. 59-60, summarizing Dasgupta and David (1994). Note that "It was the sociologist Robert Merton who established the priority rule as a characteristic of the social organization of science worthy of study", according to Strevens (2003), p. 57. Strevens's paper is a classical reference on understanding the role of the priority rule in science.

enough to dissuade its publication. One of the aims of the paper is to establish whether this concern is justified, and whether the reward system should be improved.

The paper is organized as follows. In section 2, I define what an intermediate result is and I illustrate it with historical examples. Then in section 3, I propose a formal model of the problem of deciding which publishing strategy is the best. Finally in section 4, I use the model to identify conditions under which a scientist should or shouldn't publish an intermediate result.

2 Definition and examples of intermediate results

I define a *result* (or a *work*) as a piece of knowledge produced by scientists, which is large enough to be accepted for publication, under current standards. A result is thus publishable, but not necessarily published. The requirement that it can be accepted for publication intends to rule out too small outcomes (the first successful computation, or the first experimental data), so as to consider them only at the size of what exists in standard scientific journals⁵. I come back on this assumption in section 4.4.

So as to define what an *intermediate* result is, I need beforehand to define a “fruitfulness relation”. It is a well-known fact that the production of a new scientific result usually requires — in a way to be specified below — that a scientist relies on some previous results, and she usually cites the published results she reuses⁶. I propose to define a fruitfulness relation between a set of past results and a new result (or, to say that they are fruitful for it) in case the past results have facilitated the achievement of the new result for the scientist. By “facilitated”, I mean that without these past works, it would have been harder for the scientist to achieve the new result (the “harder” condition can be evaluated with some standard working hours). This definition calls for some remarks. First, the judgement that some results are fruitful is made retrospectively about a given historical situation, not in the abstract (but there may be other paths where these results wouldn't have been used at all⁷). Second, stating a fruitfulness relation amounts to identifying some of the facilitating past works, but the list needs not be exhaustive. Finally, fruitfulness is a matter of degree: a work can rely more strongly on some work, and more lightly on another one.

The fruitfulness relation can be used to analyze the research road to some goal: we can highlight a chain of results W_i which are successively fruitful, with the last one reaching the goal. An intermediate result is a particular kind of fruitful result: I define a result W_i as *intermediate* for another result W_{i+1} in case

- (i) at the time the scientist achieves W_i , she intends to reach W_{i+1} ;
- (ii) at the time she achieves W_i , she expects W_i to be strongly fruitful for W_{i+1} ;

⁵It seems that the size of publications has changed through the decades, and a detailed study of its variations would be interesting. Before the creation of scientific journals in the XVIIth century, the unit of publication was the book. Nowadays, journals publish articles with various lengths. For example in physics, the *Physical Review* publishes regular articles, and also “Rapid Communications”, “Brief Reports”, and “Comments”.

⁶For a discussion on how citations indicate intellectual debts, cf. for instance Collins (1974), p. 170.

⁷About historical contingency, see for example Hacking (1999), chap. 3.

(iii) W_i turns out to be strongly fruitful for W_{i+1} .

(i) means that the scientist who gets W_i doesn't consider to have reached her scientific aim: she will not stop this research there, but start researching for W_{i+1} or more. She may have considered W_{i+1} as her goal from the beginning, but it is sufficient that it is so when she reaches W_i . (ii) requires that the scientist is aware that the result she has passed can be very useful for her forthcoming research. (iii) is meant to exclude cases where she misjudges the possible fruitfulness of a result. Note that the concept of an intermediate result depends on what the scientist is aware of and on what she expects at the time she achieves the result. This is wanted because whether to publish is first not a question for the community, but only for the scientist who reaches the result: only if the scientist considers the result as intermediate will the publication question be posed. Anyway, an objective criteria is also present in condition (iii). To summarize, one can say that an intermediate result is a particular kind of fruitful result, which is expected to be so when it is reached.

To help clarify the above definition, I begin with a counter-example. Consider Einstein's two theories of relativity: the special of 1905, and the general of 1912. Has the special relativity been an intermediate result towards the general relativity? No, because conditions (i) and (ii) are not met: in 1905, Einstein thought he had reached a satisfactory and final result as such, with no idea of a more general theory; only in 1907 did he realize that a new theory was needed for gravitation⁸.

I now give an example of a chain of intermediate results. Consider the series of Aspect's papers in quantum mechanics, which led to the famous experiment of Aspect *et al* (1982). This experiment realized the EPR thought experiment, thus answering crucial questions about entanglement and locality in quantum mechanics. The experimental results were presented in several papers: first, in a basic set-up, and a few months later with some gradual refinements⁹. The successive papers can be considered as a chain of intermediate results. The essential condition (i) is clearly met: the detail of the final experiment had been presented and theoretically discussed in a preliminary paper (Aspect, 1976).

3 A sequential model of research

3.1 A formal model

Suppose a scientist reaches an intermediate result: should she publish it now or wait for the final result? Arguments for both choices have been considered in the introduction, and some others could be put forward¹⁰. These arguments pull in opposite directions, in such a way that no obvious strategy seems to come out in general. I can see at least two methodological ways to tackle with this problem, so as to decide what the best publishing strategy is. A first possibility is to refine the description of the particular research situation

⁸Cf. Pais (1982), p. 178.

⁹Aspect, Grangier and Roger (1981, 1982) and Aspect, Dalibard and Roger (1982).

¹⁰For instance: as scientists are pressured to have a great number of publications (the implicit "publish-or-perish" rule), several small papers are better than a big one. Conversely, top scientific journals like *Science* or *Nature* seem not to accept for publication a result of which an intermediate step has already been published. So, a scientist might not publish an intermediate step if she can expect to publish the whole result in such prestigious journals.

under study, in carrying detailed case-studies so as to advise the scientist more surely. Instead of focusing on specific conditions, a second possibility is to select the common and most important parameters and to identify general normative trends. By idealizing or simplifying the research situations, one should become able to decide which strategy is the best. Starting from a simple model, it should be possible to describe more faithfully a given real situation. In this paper, I shall adopt this second methodology: I will propose an idealization (that is, a model) of some of the arguments about the best publishing strategy.

Most of these arguments involve weighting a gain versus a loss or a risk. So as to clarify this informal reasoning, I will define some preferences for the scientists, some probabilities and rewards, and so on. In other words, I adopt the framework of rational decision theory and of game theory. In doing so, I will not force any pre-existing economic model or theory upon the study of intermediate results¹¹. In formalizing some arguments, the modest hope is that it will help to clarify them, and propose a fruitful framework to discuss them. Such a modelling approach has been developing in recent years, in the field of social epistemology of science. Formal models have been proposed for various questions¹², but not yet for investigating the publication of intermediate results.

3.2 Presentation of the model

3.2.1 Outline of the model

The idealized situation I consider is, roughly, the following: two scientists are competing on a research program, consisting of several intermediate results which have to be passed in order. Scientists always carry out some research for the next intermediate result, and they are successful with some probability per unit of time. When a scientist, say *A*, reaches an intermediate result, she can either publish it, or continue her research. After a result is published, the competitor *B* can start from it so as to continue her own research. Thus, when *A* publishes she loses her lead, and *B* becomes in a position to compete with her for the next step. But *A* is also recognized as the first scientist to reach this step.

During the detailed presentation of the model, I will sometimes highlight the fact that its hypotheses are certainly not in favor of early publication. So, when it will be shown in section 4.2 that the best strategy is to publish intermediate results, it should not be thought that it is because of too accomodating hypotheses — the model is indeed a “worse-case” analysis.

3.2.2 A chain of steps

I now turn to the detailed presentation of the model. The research project is a chain of successively intermediate results. The “results” are called synonymously “steps”, and they are numbered 0, 1, 2... (fig. 1). For simplicity, I make the assumption that there is no other route to k than from $k - 1$: the research can be qualified as “sequential”.

¹¹My attempts to reuse and transpose some economic literature about patents to the question of intermediate publication have not been successful.

¹²Important contributions are for instance De Langhe and Greiff (2009), Goldman (2009), Kitcher (1990) and (1993), Strevens (2003) and (2006), Weisberg and Muldoon (2009), Zollman (2007), (2009), (2010a) and (2010b).




Figure 1: The chain of intermediate steps in which the research consists.

Time is discretized in temporal intervals, and each temporal interval includes two possible actions: one about research, and one about publication (cf. rules below). The research game can end in two ways: either there is a last step l on the chain, which is the goal of research (this hypothesis is called “finite chain”), or there is a temporal limitation, with n temporal intervals (hypothesis called “finite time”). These hypotheses model different situations: the finite chain corresponds to a particular scientific goal, whereas the finite time corresponds to a competition with a temporal deadline, like applying for a new position.

3.2.3 Two competitive scientists

The number of scientists is fixed to two, called A and B . These scientists are the agents for the game: they can also be interpreted as teams of scientists, or laboratories — provided that its members take the same decisions regarding publication. Scientists always have a definite position on the chain, corresponding to the results they have achieved, starting from step 0.

3.2.4 Research activity

At each temporal interval, a scientist undertakes some research, so as to pass the next step (she is supposed to be freed from any teaching or administrative duty). No cost is associated with carrying on some research, so the question is whether to publish, not whether to undertake some research. During one temporal interval, her research is successful with probability p , the same for all scientists at all time. p being small means that the steps are hard to pass, it being close to 1 means that they are very easy.

3.2.5 Publication

At the end of each temporal interval, a scientist can decide to publish. Her strategy is totally free: she can publish only the step she has just passed, or her last two steps together, or a step she has passed long ago, etc. For a scientist, the consequence that her competitor publishes is that she is moved ahead on the chain, if she had not yet reached this step. When a scientist decides to publish, the journal is supposed to accept her submission because she has passed the step. Note that a scientist can never know exactly where her competitor is: a published step is only a minimum reached position. Note also that in this model, the scientist doesn’t keep anything to herself in addition to what she publishes, for a given step — like some details of the experiment or of the computation. The publication of a step enables her competitor to catch up with her in all respect (having passed the previous step gives no advantage whatsoever). As this idealization quite exaggerates the disadvantageous effect of a publication, it is rather unfavorable to the strategy of publishing intermediate results.

3.2.6 Value v of the reward

Each published step has a value v , which is given to the first scientist who publishes it, because of the priority rule. In case both scientists publish the same step during the same temporal interval, both get $v/2$ ¹³.

The reward v can receive two interpretations. The first one is epistemic: the reward directly corresponds to the progress the scientist has brought to society's knowledge¹⁴. The condition for the reward is only that she makes her progress publicly known in a journal, so that it becomes available to anyone. A scientist's interest can be epistemic and selfish, if she cares only about her own production. The second interpretation of v is non-epistemic: it represents fame, authority, esteem, money, prizes, and so on. When a scientist publishes a result, she becomes more famous and her reputation increases with some amount v . A central simplifying assumption here is that the non-epistemic value of the reward corresponds to its epistemic value. Thus, whether scientists care about epistemic or non-epistemic reward is equivalent, and the model will not study a possible discrepancy between them. I will only consider the "individual" strategy, be it epistemic or not, to be contrasted with the best strategy according to society (which is epistemic only).

3.2.7 End of the game

During the game, scientists gather rewards. I suppose that their aim is to get the maximum amount of reward at the end of the game¹⁵ (in the finite time hypothesis, the end occurs with the n^{th} time interval; in the finite chain hypothesis, it occurs when a scientist reaches the last step). I make the assumption that the scientist is risk-neutral¹⁶ and that she wants to maximize her *average* reward. I also suppose that scientists are rational and that they will choose their publishing strategy in accordance with this final aim. The end of the game is the only moment of comparison between the two scientists: if n or l varies, the model enables to investigate "long run" (large n or l) versus "short run" strategies (small n or l).

This model assumes that the probabilities p to pass a step, and the reward v are the same for all the steps. I call this hypothesis the "homogeneous chain" hypothesis (modifications will be considered later).

¹³There are two ways to justify for it: either one of the scientists is considered to be actually the fastest to write to the journal and to get the whole reward, with probability 1/2; or simultaneity is interpreted in a strict way: both scientists are acknowledged as being independent discoverers, and both receive half of the reward — for example, as the amount of money in a Nobel Prize is divided into the laureates.

¹⁴I don't want to argue for any particular conception of scientific progress here. My notion of progress is supposed to be compatible with almost any conception discussed in the literature.

¹⁵It must be noted that the hypothesis according to which scientists are interested in the sum of v is not particularly favorable to the strategy to publish. It doesn't make a difference between a scientist who publishes 3 steps at one time, or the 3 steps one by one. For instance, counting the number of publications would favor the last option, and the strategy to publish in general.

¹⁶A *risk-neutral* agent has no preference if she is presented with the following choice: either receive v with certainty, or receive $2v$ or 0 each with probability 0.5. A *risk-averse* agent prefers the first scenario, a *risk-seeking* agent prefers the second one.

4 Results of the model

I now turn to the results which can be obtained from this model.

4.1 The best strategy for society

The interest of society has been considered to be only epistemic. In the finite time hypothesis, it consists in the highest published step. In the finite chain hypothesis, it lies in the fact that the last step is reached in the shortest time. In both cases, it is easy to see that it demands that scientists publish an intermediate result as soon as they get it, because it avoids that the other scientist undertakes some research for a step which has actually already been passed. The question investigated in the next sections is whether the best strategy for scientists as individual rational players is also to adopt such a strategy.

4.2 The best strategy for an individual scientist

As for the best strategy for an individual scientist, two theorems can be proved, which correspond to the hypothesis of finite time or of finite chain¹⁷:

Theorem 1: In the homogeneous chain (in terms of rewards and difficulties) with the finite time hypothesis, the best strategy for a scientist is to publish immediately every intermediate step she has passed.

Theorem 2: In the homogeneous chain with the finite chain hypothesis, the best strategy for a scientist is to publish immediately every intermediate step she has passed.

The hypothesis made in the model, stating that the scientist is risk neutral, is not required in the proof and can be relaxed. For these results, all that is needed is an homogeneous chain, with constant probabilities and rewards. There is no condition on the specific value of the size of the steps, nor on their difficulties. The theorems show that there is no difference either between the short and long-run best strategy (the result doesn't depend on l or n). The best strategy for a scientist doesn't depend on her competitor's strategy.

At first sight, the results of the theorems can seem to be counter-intuitive: they state that there exist simple situations in which selfish scientists who care about their own fame or reward should publish as soon as possible, that is, should help their competitors in disclosing their own results. This comes from the fact that the incentive to publish is bigger than the help given to competitors. One moral to keep home is that it would be wrong to think that selfishness is a sufficient reason for scientists to hide their intermediate results in any case, and cease any cooperation with others.

A comparison with section 4.1 shows an agreement between the individual and the collective rationalities, as they have the same best strategy. This means that, for homogeneous research situations (in terms of difficulties and rewards), selfishness should not hamper the collective epistemic progress, and that society should not fear individual

¹⁷The proofs can be found in the appendix, sections A and B respectively. They rely on a backward induction proof, which shows some structural similarities between this model and the centipede game, for example.

interests. As the incentive to publish intermediate steps is sufficient, society needs not reconsider the reward system of science, which seems appropriate here.

If these theorems are normative, note that the hypothesis of the model are not themselves normative: I am not arguing that scientists should care about their individual reward only, or that a reward should come from a publication only. I just claim that, making almost the worst hypotheses (scientists are selfish, interested in non-epistemic rewards, which come only from publication), scientists should publish their intermediate results anyway. If they should do so when they are selfish, they should all the more do so when they care about the community.

4.3 So, when should researchers not publish intermediate steps? Two modifications of the model.

In the model considered so far, the best strategy is always to publish every passed step, whatever the end of the game. How should the research situation look like, so that the best strategy is not to publish? I consider in turn two modifications, based on the previous model: the steps may have different rewards (the “inhomogeneous reward” case), or some steps may be more difficult than others (the “inhomogeneous difficulty” case).

4.3.1 An inhomogeneous reward

Let’s start with a simple case of an inhomogeneous reward. Consider a two step chain, with the first step worth nothing ($v_1 = 0$) and the second one worth $v_2 = v$. Suppose A passes the first step: publishing it gives her no advantage (no reward), but only drawbacks, because it moves B upward and enables B to compete more easily for the v reward. It suggests that if an intermediate step is not worth enough, or if a final step comes with a big reward, a scientist should not publish the intermediate step. For two steps¹⁸, theorem 3 states a quantitative limit¹⁹:

Theorem 3: In a chain with two steps with different rewards v_1 and v_2 , the necessary and sufficient condition for which a scientist should not publish immediately the intermediate step, but only at the end, is

$$v_2 > \frac{2-p}{1-p}v_1.$$

Note that the value of the threshold depends on the relative values of the rewards. Supposing that p is small (that is, that research is hard, which seems to be a reasonable assumption), the limit approximates to $v_2 > 2v_1$. The theorem states a necessary and sufficient condition: an intermediate step should be published immediately if and only if $v_2 < \frac{2-p}{1-p}v_1$: this reveals the fact that there is some margin in the conditions for validity of theorem 2.

¹⁸The theorem is limited to two steps, for computational reasons. With more steps, computations will be much more complicated, without in all likelihood changing the trend of the result.

¹⁹The proof can be found in the appendix, section C. A technical point to note is that theorem 3 is only proved on average, in contrast with theorems 1 and 2, in section 4.2. For some outcomes of A and B ’s research, publishing the intermediate step might be better. In other words, the validity of theorem 3 might depend on the hypothesis that scientists are risk-neutral (cf. section 3.2).

The interpretation of the theorem is that scientists should not (from their selfish viewpoint) publish an intermediate step if they are not rewarded enough for it, compared to what they can expect from the next step. This gives a limit to the size of what scientists should publish. When research is hard, the limit is that the second step gives twice as much reward as the first step.

Scientific research is full of such situations, where intermediate and preliminary steps don't bring much reward, but are necessary to reach a big and valuable goal. There, the best strategy according to the individual scientists is different from the best strategy according to society. In contrast to the homogeneous chain, scientists' selfishness will hamper epistemic progress, and should be feared by society. Thus, one might consider improving the reward system, giving more incentives for the publication of intermediate results with lower rewards. This is a vast question, to which I shall propose no definitive solution. A suggestion could be to better acknowledge the value of preliminary and intermediate steps, instead of emphasizing only the achievement of the last step (and keeping only "big names" for history's records).

4.3.2 An inhomogeneous difficulty

I now turn to a chain with an inhomogeneous difficulty (but homogeneous reward). Consider a simple case: in a two step chain, the first one is hard to pass (p_1 is close to 0), whereas the second one is passed without difficulty ($p_2 = 1$). Thus, a scientist who passes the first step almost gets a second step "for free", so it seems that she should not publish it now. This simple case suggests that if an intermediate step is followed by a much simpler step, a scientist should not publish it now. The following theorem states a quantitative limit²⁰:

Theorem 4: In a chain with two steps passed with probabilities p_1 and p_2 , the necessary and sufficient condition for which a scientist should not publish immediately the intermediate step, but only at the end, is

$$p_2 > \frac{2 - p_1}{1 - p_1} p_1.$$

Comments are, in part, similar to those on theorem 3. The limit to what scientists should publish is not fixed once for all, but is only comparative. Again, supposing that p_1 is small, the condition approximates to $p_2 > 2p_1$.

When steps have such decreasing difficulties, the best strategy according to an individual scientist is different from the best strategy according to society. Again, in this inhomogeneous chain, there is a discrepancy which should be feared by society. However, it is of lesser importance than in the previous case: if the intermediate result is not published, the second step should be reached quickly (because it is easier).

²⁰The proof can be found in the appendix, section D. Note that theorem 4 is only proved on average: its validity depends on the hypothesis that scientists are risk-neutral.

4.4 Accounting for the size of publishable steps?

I would like now to discuss the limits of this idealized model, and more particularly the hypothesis according to which the size of publishable results (steps) is given or fixed by the standards of the discipline.

First, consider a homogeneous research situation, with regular steps of reward v and probability p , as in section 4.2. Suppose a new journal is created, which accepts to publish every step of this chain, plus half steps. For this journal, there are twice as much steps, all worth $v/2$. It is a homogeneous chain: theorems 1 and 2 apply, stating that scientists should publish these half steps. This means that, if scientists are given the possibility to publish smaller steps, they should do it — or that they should give themselves this possibility, for instance in launching their own journal. Could the size of a publishable step really be as small as possible? Some neglected effects should be taken into account, like the cost of writing, reviewing, publishing an increasing number of papers, as well as the cost of finding and retrieving the relevant information. All these reasons go against the existence of always smaller steps, and could account for the existence of a minimum threshold for a publishable result.

Consider now an inhomogeneous chain of steps, as in section 4.3. In particular, suppose it is a long chain of steps with inhomogeneous rewards, alternating bigger and smaller ones (with homogeneous difficulty). Assuming that the trends identified in section 4.3 remain valid in more complex situations, reward inhomogeneities should give rise to local thresholds, under which too small steps will not be published when passed. Instead, they will be published together with, say, their next step: uneven steps aggregate for publication into bigger steps of comparable rewards. These bigger steps define a more homogeneous and new chain of steps with (almost) the same reward. It falls under the case of theorem 2: such intermediate steps should be published. The same can be said of a chain of steps with inhomogeneous difficulties: they will aggregate for publication into bigger steps, in such a way that the difficulty of these bigger steps is now comparable. They define a new and more homogeneous chain of steps, which should be published when passed. To sum up, my conjecture is that the net result of section 4.3 on a long chain is to redefine the size of what should be considered for publication, and section 4.2 says that it should indeed be published immediately. Thus, the size of the steps needs not just be fixed externally, but may also be accounted for within the model.

5 Conclusion

These results of the model suggest two qualified philosophical morals. The first one is about the discrepancy between the individual and the collective rationalities, as thematized by Kitcher. For the publication of intermediate steps, there exist simple and plausible research situations in which there is no mismatch between the individual and the collective rationalities (be they epistemic or not). It would be wrong to think that, as soon as some results are identified as intermediate, scientists's selfishness should hamper collective progress. This first moral needs to be qualified: as soon as a final aim brings much more (say, two times) reward than an intermediate step, the mismatch between individual and collective rationalities cannot be avoided. In these situations, selfishness advises not to

publish the intermediate step, which is a brake to epistemic progress, from the viewpoint of the collective.

A second moral to be drawn is about the reward system in science. There exist simple and plausible research situations for which the incentive to publish intermediate steps is sufficient. However, when the final step has a much bigger reward than the intermediate steps, the incentive is not sufficient any more. As a consequence, looking for improvements of the reward in these situations could be a valuable goal for social epistemologists of science.

Several improvements or extensions of this model could be considered in a future work. Detailed case-studies could be developed so as to improve the application of the model to real cases (parameters in the model might take different values according to the subject matter of the research field, or to its theoretical or experimental nature). Instead of considering a chain of steps, one could consider a network of inter-related steps, where a given result can be reached by different paths. An asymmetry could be introduced between the two scientists, one being a better researcher than the other. The number of scientists could be allowed to vary. While I am aware that such improvements could be brought to the model, I consider this paper itself as a first step in this research program, and I follow my own advice in deciding to publish this paper.

Acknowledgements

References

- ASPECT, Alain (1976), “Proposed Experiment to Test the Nonseparability of Quantum Mechanics”, *Physical Review D*, 14 (8), p. 1944-1951.
- ASPECT, Alain, DALIBARD, Jean and ROGER, Gérard (1982), “Experimental Test of Bell’s Inequalities Using Time-Varying Analysers”, *Physical Review Letters* 49 (25), p. 1804–1807.
- ASPECT, Alain, GRANGIER, Philippe and ROGER, Gérard (1981) “Experimental Test of Realistic Local Theories via Bell’s Theorem”, *Physical Review Letters*, 47 (7), p. 460-463.
- (1982), “Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell’s Inequalities”, *Physical Review Letters*, 49 (2), p. 91-94.
- BROMBERG, Joan Lisa (1991), *The Laser in America, 1950-1970*, MIT Press.
- COLLINS, Harry M. (1974), “The TEA Set: Tacit Knowledge and Scientific Networks”, *Science Studies*, Vol. 4, No. 2 (Apr., 1974), p. 165-185.
- DASGUPTA, Partha and DAVID, Paul. A. (1994), “Toward a New Economics of Science”, *Research Policy*, XXIII: 487-521.
- DE LANGHE, Rogier and GREIFF, Matthias (2009), “Standards and the Distribution of Cognitive Labour: A Model of the Dynamics of Scientific Activity”, *Logic Journal of the IGPL*, 18(2): 278-94.
- GOLDMAN, Alvin (2009), “Systems-Oriented Social Epistemology”, in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 3, Oxford: Oxford University Press, p. 189-214.
- KITCHER, Philip (1990), “The division of cognitive labor”, *Journal of Philosophy*, Vol. 87, No. 1 (Jan. 1990), p. 5-22.

— (1993), *The Advancement of Science. Science without legend, objectivity without illusions*, Oxford: Oxford University Press.

HACKING, Ian (1999), *The Social Construction of What?*, Harvard University Press.

PAIS, Abraham (1982), *Subtle is the Lord: The Science and the Life of Albert Einstein*, Oxford University Press.

STREVVENS, Michael (2003), “The Role of the Priority Rule in Science”, *Journal of Philosophy*, 100:55–79.

— (2006), “The Role of the Matthew Effect in Science”, *Studies in History and Philosophy of Science*, Vol. 37, p. 159-173.

WEISBERG, Michael and MULDOON, Ryan (2009), “Epistemic Landscapes and the Division of Cognitive Labor”, *Philosophy of Science*, 76 (April 2009), p. 225–252.

ZAMORA BONILLA, Jesús P. (2012), “The Economics of Scientific Knowledge”, in Mäki, U. (ed.), *Philosophy of Economics*, series *Handbook of the Philosophy of Science*, Elsevier.

ZOLLMAN, Kevin J. S. (2007), “The Communication Structure of Epistemic Communities”, *Philosophy of Science* 74(5), p. 574-587.

— (2009), “Optimal Publishing Strategies”, *Episteme: a Journal of Social Epistemology*, 6(2), p. 185-199.

— (2010a), “Social Structure and the Effects of Conformity”, *Synthese*, 172(3), p. 317-340.

— (2010b), “The Epistemic Benefit of Transient Diversity”, *Erkenntnis*, 72(1), p. 17–35.

Appendix

A Proof of theorem 1 in section 4.2

It is useful to distinguish between two kinds of steps: the steps both A and B have passed but not yet published (call them the “common steps”), and the steps A is the only one to have passed, if any (her “solitary steps”). I am going to show that solitary steps should be published. Then, taking into account the possibility of common steps cannot make publication a worse strategy than it is if there are only solitary steps (because nothing can be gained in not publishing them, and a potential reward can be lost). For a result aiming at proving the superiority of publication, this is just fine.

Consider two general strategies:

- Strategy I: always publish a passed step.
- Strategy II: always publish the common steps, but never the solitary steps (publish them only when reaching the end of the game).

During the proof of the theorem, it will become obvious that mixed strategies, in which A publishes some of her a steps now and the rest at the last moment, would not be better. Proving theorem 1 amounts to showing that strategy I is better than strategy II.

The theorem is proved by induction on the number of remaining temporal intervals, for scientist A . To start with, suppose there is just one temporal interval left. As discussed before, I suppose there are only solitary steps.

Suppose A adopts strategy I. Then she publishes now her a solitary steps. She immediately gets an av reward. Then, A and B start on the same level for the last temporal interval. It is easy to fill in table 1.

$A \setminus B$	success	failure
success	$av + v/2$	$av + v$
failure	av	av

Table 1: A 's reward during the last temporal interval, with strategy I.

Suppose now A adopts strategy II: she doesn't publish now, and waits for the last time. A similar table (table 2) is easily filled in.

$A \setminus B$	success	failure
success	$av + v/2$	$(a + 1)v$
failure	$av - v/2$	av

Table 2: A 's reward during the last temporal interval, with strategy II.

Tables 1 and 2 only have one difference, in the bottom left cell. This shows that strategy I is better than strategy II, and establishes the first case of the induction proof.

Suppose now it has been proved for some $m < n$ that, when m temporal intervals remains, A should publish every step she has passed. Consider A when there remains $m + 1$ temporal intervals. By hypothesis, A should publish at the next step, when there remains only m temporal intervals. Thus, she is in the case already proved²¹, where $m = 1$. Hence, theorem 1 has been proved by induction.

Note that the conclusion that A should always publish has not only been shown on average, but also for every possible outcome of research — for every comparison of the corresponding cells in tables 1 and 2. So, the theorem doesn't depend on the hypothesis that the scientist is risk-neutral.

B Proof of theorem 2 in section 4.2

Theorem 2 is a corollary of theorem 1. As the chain is finite, an unknown x number of temporal intervals will be needed so that (at least) one researcher reaches the end of the chain (x is finite, otherwise whether researchers should publish does not make much sense anymore). A researcher wants to reach first the end of the chain or, in other words, to get the bigger reward after this unknown x number of temporal intervals. Theorem 1 can be applied with the hypothesis of the finite time with x temporal intervals. One concludes that the best strategy is to publish every intermediate step, and theorem 2 is proved.

²¹Indeed, the hypothesis of the end of the game is nothing more than the insurance that everyone will publish, and this is all the role it played in the above proof. Note that if B didn't publish at n remaining intervals, it would be even better for A . The above proof supposed that B would publish, so it deals with the worse case.

C Proof of theorem 3 in section 4.3.1

As before, it is clear that A should publish a common step, and also that she cannot know in practice which steps are solitary or common. I will compare strategy I (cf. section A) with a new strategy:

Strategy III: never publish an intermediate step — only publish everything at the last step.

Theorem 3 is equivalent to stating that strategy III gives a larger average reward than strategy I. As the total reward distributed among scientists is constant at $v_1 + v_2$, it is equivalent to stating table 3.

$A \setminus B$	Strategy I	Strategy III
Strategy I	$(v_1 + v_2)/2 ; (v_1 + v_2)/2$	$< (v_1 + v_2)/2 ; > (v_1 + v_2)/2$
Strategy III	$> (v_1 + v_2)/2 ; < (v_1 + v_2)/2$	$(v_1 + v_2)/2 ; (v_1 + v_2)/2$

Table 3: Expected rewards for A and B .

Let's consider the case of the bottom left cell. Let's call E_4 the average reward that A can expect after the first temporal interval. Table 4 is filled in: if A and B succeed, only B publishes because A follows strategy III; then, they are in a symmetrical position for the last step: A 's expected reward is $v_2/2$. If A fails and B succeeds, only B publishes too, and A 's expected reward is $v_2/2$. If both scientists fail, they remain in the same position, and the reward is still E_4 . If A succeeds and B fails, A 's reward is called E_6 and is computed below.

$A \setminus B$	success	failure
success	$v_2/2$	E_6
failure	$v_2/2$	E_4

Table 4: A 's reward, according to the outcome of the first temporal interval.

$A \setminus B$	success	failure
success	p^2	$p(1-p)$
failure	$p(1-p)$	$(1-p)^2$

Table 5: Probability of the outcomes of A and B 's research.

E_4 is computed by weighting table 4 with table 5:

$$E_4 = p^2 \frac{v_2}{2} + p(1-p) \frac{v_2}{2} + p(1-p)E_6 + (1-p)^2 E_4 \quad (1)$$

The evaluation of E_6 is made by filling in a similar table 6.

$A \setminus B$	success	failure
success	$v_1/2 + v_2$	$v_1 + v_2$
failure	$v_2/2$	E_6

Table 6: A 's expected reward, in case A is at step 1 and B at step 0.

Weighting table 6 with table 5, one computes

$$E_6 = p^2 \left(\frac{v_1}{2} + v_2 \right) + p(1-p) \left(v_1 + v_2 + \frac{v_2}{2} \right) + (1-p)^2 E_6. \quad (2)$$

E_6 is extracted from it and injected in eq. 1:

$$E_4 = \frac{v_2}{2(2-p)} + \frac{1-p}{2(2-p)^2} [(2-p)v_1 + (3-p)v_2] \quad (3)$$

The condition expressed in the bottom left cell of table 3 is that $E_4 > \frac{v_1+v_2}{2}$. With eq. 3, it is equivalent to

$$v_2 > \frac{2-p}{1-p} v_1. \quad (4)$$

D Proof of theorem 4 in section 4.3.2

The proof of theorem 4 is very similar to the proof in section C, and I only give the key differences here. The total distributed reward is $2v$, so the expected rewards should be table 7 and 8.

$A \setminus B$	Strategy I	Strategy III
Strategy I	$v ; v$	$< v ; > v$
Strategy III	$> v ; < v$	$v ; v$

Table 7: Expected rewards for A and B , according to their strategies.

$A \setminus B$	success	failure
success	$v/2$	E_9
failure	$v/2$	E_8

Table 8: A 's expected reward, according to the outcome of the first temporal interval.

E_8 is computed by weighting table 8 with table 5:

$$E_8 = p_1 \frac{v}{2} + p_1(1-p_1)E_9 + (1-p_1)^2 E_8 \quad (5)$$

E_9 is computed from table 9, weighted with table 10:

$$E_9 = \frac{p_1/2 + 2p_2 - p_1p_2}{p_1 + p_2 - p_1p_2} v. \quad (6)$$

$A \setminus B$	success	failure
success	$v/2 + v$	$2v$
failure	$v/2$	E_9

Table 9: Expected reward E_9 for A , in case A is at step 1 and B at step 0.

$A \setminus B$	success	failure
success	$p_1 p_2$	$p_2(1 - p_1)$
failure	$p_1(1 - p_2)$	$(1 - p_1)(1 - p_2)$

Table 10: Probability of A and B 's outcomes of research, when A is at step 1 and B at step 0.

From eq. 6 and eq. 5, one gets

$$E_8 = \frac{v}{2(2 - p_1)} + \frac{1 - p_1}{2 - p_1} \times \frac{p_1/2 + 2p_2 - p_1 p_2}{p_1 + p_2 - p_1 p_2} v \quad (7)$$

The condition expressed in the bottom left cell of table 7 ($E_8 > v$) leads to theorem 4.