

The Incentive to Share in the Intermediate Results Game

Remco Heesen^{*†}

August 11, 2017

Abstract

I discuss a game-theoretic model in which scientists compete to finish the intermediate stages of some research project. Banerjee et al. (2014) have previously shown that if the credit awarded for intermediate results is proportional to their difficulty, then the strategy profile in which scientists share each intermediate stage as soon as they complete it is a Nash equilibrium. I show that the equilibrium is both unique and strict. Thus rational credit-maximizing scientists have an incentive to share their intermediate results, as long as this is sufficiently rewarded.

1 Introduction

This technical report provides proofs for the mathematical theorems mentioned in my article “Communism and the Incentive to Share in Science”. I begin by briefly motivating the game-theoretical model and the questions it aims to address.

As mentioned in the article, Strevens (2017) gives a “Hobbesian vindication” of the *communist norm*, the social norm in science that mandates scientists to share their research findings widely. There is need for a Hobbesian vindication, according

^{*}Faculty of Philosophy, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DA, UK.
Email: rdh51@cam.ac.uk.

[†]This technical report is a companion to my article “Communism and the Incentive to Share in Science”, published in the journal *Philosophy of Science*, vol. 84, no. 4 (October 2017), under DOI 10.1086/693875. The article is available from the publisher and from my personal website. Work on both documents was partially supported by the National Science Foundation under grant SES 1254291 and by an Early Career Fellowship from the Leverhulme Trust and the Isaac Newton Trust.

to Strevens, because scientists would not be inclined to share their findings in the absence of the communist norm. In particular, because getting rewarded for scientific achievements depends upon publications, there is “a powerful incentive not to share your results before you have extracted every last publication from them” (Strevens 2017). Other authors have also claimed the existence of such a tension between the norm of sharing and the desire to receive credit for one’s work (e.g., Dasgupta and David 1994, Arzberger et al. 2004, Resnik 2006, Borgman 2012, Soranno et al. 2015).

In contrast, two recent papers have used game-theoretic models to suggest that sharing may be individually incentivized despite the potential for future discoveries, as long as partial or intermediate discoveries are rewarded with credit roughly proportional to their difficulty (Boyer 2014, Banerjee et al. 2014). This is directly contrary to the view argued in the quotes above if one thinks that the real reward structure of science awards credit proportional to difficulty; if one thinks that this is not the case these papers may be viewed as providing a policy recommendation for how to incentivize sharing.

Boyer (2014) studies a model which is, by his own admission, highly idealized. Hence it cannot by itself support general claims about the incentives faced by scientists regarding intermediate results sharing. Banerjee et al. (2014) address this worry by providing a model that relaxes Boyer’s assumptions that there are only two scientists, that the scientists are equally productive, that different intermediate results are equally hard to achieve, that intermediate results can only be achieved in one order, and that scientists share either all or no intermediate results (note that while I cast them as addressing a worry about Boyer, to my knowledge Banerjee et al. were not aware of Boyer’s work, and vice versa).

Banerjee et al. (2014) show that sharing is a Nash equilibrium in their model if sufficient credit is given for intermediate results. They do not show that the equilibrium is unique, except in a special case with Stackelberg agents, where a significant proportion of the scientists commits to sharing in advance.

This leaves us with a potential *equilibrium selection problem*. Since there may be other Nash equilibria, it is unclear whether or under what circumstances rational credit-maximizing scientists can be expected to share.

This paper addresses this issue by showing that in many cases no equilibrium selection problem exists: under slightly strengthened assumptions the equilibrium identified by Banerjee et al. is unique. I also show that the equilibrium is strict, and that the game is a weakly better reply game in the sense of Huttegger (2013). The significance of the latter result is that boundedly rational scientists are likely to find the equilibrium.

Section 2 describes Banerjee et al.’s model, which I call the *Intermediate Results*

Game. In section 3 I suggest a distinction between a version of the Intermediate Results Game with perfect information (scientists know when other scientists have completed but not published intermediate results) and one with imperfect information. I show that the sharing equilibrium is the backwards induction solution of the Intermediate Results Game with perfect information and that other equilibria of the Game differ from the backwards induction solution only off the equilibrium path.

Section 4 shows that sharing is the unique and strict equilibrium of the Intermediate Results Game with imperfect information. Section 5 extends the results of sections 3 and 4 to the case of a directed acyclic network determining the order in which intermediate results can be completed (as opposed to a single sequence), again under slightly strengthened assumptions compared to those for which Banerjee et al. show the existence of the sharing equilibrium.

Section 6 shows that the Intermediate Results Game with imperfect information of sections 4 and 5 is a weakly better reply game. Section 7 shows that the results of the preceding sections also hold in a version of the Game in which credit is measured per unit of time. Two appendices provide the proofs.

2 The Intermediate Results Game

Consider a research project that can be divided into $k \geq 1$ intermediate stages. The stages can only be completed sequentially: stages 1 through $j - 1$ must be completed before work on stage j can be started (this assumption will be relaxed in section 5).

There are $n \geq 2$ scientists (or research groups) working on the research project. Their productivity is modeled by (*nonstationary*) *Poisson processes*. Poisson processes are used to model the occurrence of events at random intervals. Here, the occurrence of an event is interpreted as some scientist successfully completing an intermediate stage. John C. Huber has shown in a series of papers that scientists' productivity is accurately modeled by a Poisson process (Huber 1998a,b, Huber and Wagner-Döbler 2001a,b, Huber 2001).

Each scientist i has a productivity rate $\lambda_{ij} > 0$ while working on stage j . This captures the speed at which she works: working at rate λ_{ij} , she would expect to complete λ_{ij} stages per unit of time. The expected time to complete stage j is $1/\lambda_{ij}$.

Since productivity may vary by scientist and by stage, the Intermediate Results Game allows for differences in inherent difficulty between stages, differences in inherent aptitude between scientists, as well as differences in specialization among scientists (making some stages easier and some harder for different scientists, although the extent to which scientists may have different specializations is restricted by the assumption of Proportional Credit to be introduced below).

Define $\sigma_{j_1 \dots j_n}$ to be the total productivity of the scientists when scientist i is working on stage j_i . I use σ_j as shorthand for the total productivity in the important special case in which all scientists are working on stage j , i.e.,

$$\sigma_{j_1 \dots j_n} = \sum_{i=1}^n \lambda_{ij_i} \quad \text{and} \quad \sigma_j = \sum_{i=1}^n \lambda_{ij}.$$

In a Poisson process the waiting time between two events follows an exponential distribution. What this means here is that, if T_{ij} denotes the time it takes scientist i to complete stage j : $\Pr(T_{ij} > t) = \exp\{-t\lambda_{ij}\}$. This distribution has some formal features that I will use (Norris 1998, section 2.3).

First, it is *memoryless*. This means that if at a given time the waiting time has not ended yet, the distribution of the remaining waiting time is equal to the original distribution of the waiting time:

$$\Pr(T_{ij} > s + t \mid T_{ij} > s) = \Pr(T_{ij} > t).$$

Second, the minimum of n independent exponential random variables with parameters λ_{ij_i} ($i = 1, \dots, n$) is itself exponentially distributed with parameter $\sigma_{j_1 \dots j_n}$. In other words, we can equivalently view the scientists' productivity as one Poisson process with parameter $\sigma_{j_1 \dots j_n}$ or n independent Poisson processes with parameters λ_{ij_i} . Third, the probability that scientist i is the first one to finish the stage she is working on is $\lambda_{ij_i} / \sigma_{j_1 \dots j_n}$.

Whenever some intermediate stage j is completed, the scientist who completed it chooses whether to share the result or not. If the scientist chooses not to share (strategy H) she starts working on stage $j + 1$. Other than that nothing happens until the next time some scientist completes a stage.

If the scientist chooses to share (strategy E) she gets $c_j > 0$ units of *credit* (or utility) for the stage she just completed as well as $c_{j'}$ units of credit for each stage j' she has previously completed that had not yet been shared. All scientists who had not yet solved stage j learn its solution. These scientists all start working on stage $j + 1$.

When the final stage is completed by some scientist she automatically shares it, gets c_k units of credit for the last stage, plus c_j for any stage j for which credit has not been claimed yet, and the Game ends. The Intermediate Results Game is zero-sum: at the end the total amount of credit divided among the scientists is $C = \sum_{j=1}^k c_j$.

As in any game-theoretic model, it is assumed that scientists have an interest in maximizing their utility payoff (here, credit). What strategy or strategies maximize utility may in general depend on the strategies chosen by other scientists, so it may

not be obvious what a rational (credit-maximizing) scientist would do. This has resulted in a proliferation of solution concepts for games, a number of which will figure in subsequent sections (e.g., backwards induction, bounded rationality).

The most prominent solution concept is the (*Nash*) *equilibrium*. An equilibrium is a strategy profile (i.e., an assignment of a strategy to each scientist) such that no scientist can increase her expected credit by changing her strategy unilaterally (i.e., assuming the other scientists' strategies are unchanged). When an equilibrium is played, each scientist is arguably acting rationally, as she cannot improve her expected credit through her own action. Equilibrium analysis will play a central role in subsequent sections.

3 A Backwards Induction Analysis

The previous section described a game-theoretic model of scientists working on a project that requires some number of intermediate stages to be completed. In the simplest version of the Intermediate Results Game there are two scientists ($n = 2$) and the research project has two stages ($k = 2$). The extensive form of this Game is given in figure 3.1.

At the root node (marked “ N ”) Nature decides which of the two scientists is the first one to complete the first stage of the project. As indicated, Nature picks scientist 1 with probability λ_{11}/σ_1 and scientist 2 with probability λ_{21}/σ_1 (recall that λ_{11} is scientist 1's productivity on stage 1, λ_{21} scientist 2's productivity on stage 1, and σ_1 the sum of these numbers).

Suppose Nature picks scientist 1. This leads to a decision node marked “1”, indicating that scientist 1 is the one to make a decision at this node. If scientist 1 shares the result (strategy E), she collects c_1 units of credit. Both scientists now know the solution to stage 1 of the project, so they start working on stage 2.

Nature again decides which of the two scientists completes the second stage first (with scientist 1's productivity now λ_{12} , scientist 2's λ_{22} and σ_2 the sum). In either case the Game ends. If Nature picks scientist 1, she gets credit for completing both stages of the project and scientist 2 gets nothing (as indicated by the payoff pair $(C, 0)$ in the figure). If Nature picks scientist 2, she gets c_2 units of credit, and since scientist 1 had already claimed credit for the first stage, she ends up with c_1 .

What if scientist 1 chooses not to publish her solution to the first stage of the project (strategy H at the node marked “1”)? Then scientist 1 does not collect c_1 units of credit, and scientist 2 does not learn the solution to stage 1. So now scientist 1 starts working on stage 2, while scientist 2 continues to work on stage 1.

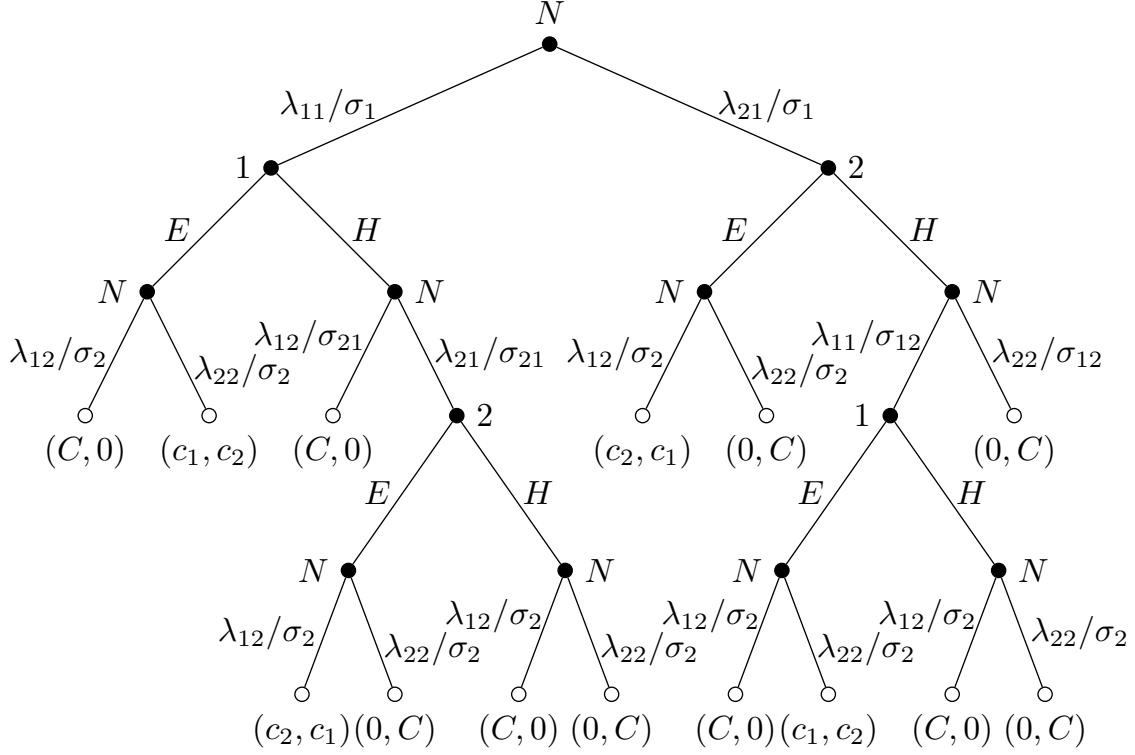


Figure 3.1: Extensive form of the Intermediate Results Game with $n = 2$ and $k = 2$.

Once again Nature decides which of the two scientists finishes the stage she is working on first (due to the memorylessness of the exponential distribution, scientist 2 is not more likely to finish fast despite having already spent some time working on stage 1; cf. section 2), with scientist 1 working at rate λ_{12} , scientist 2 working at rate λ_{21} , and σ_{21} denoting the sum of these rates. If Nature picks scientist 1, she completes the project. The Game ends and scientist 1 gets C units of credit.

If Nature picks scientist 2, she now has a decision to make (at the node marked “2”). She can claim c_1 units of credit by playing strategy E , or defer by playing H . In either case, both scientists can now work on stage 2.

Nature makes its final decision by picking a scientist who completes the second stage first. That scientist gets C units of credit if scientist 2 chose strategy H , whereas if scientist 2 chose E she gets c_1 for sure and the scientist picked by Nature gets c_2 .

The right-hand side of the figure (associated with Nature picking scientist 2 at the root node) works similarly.

It is implicitly assumed in figure 3.1 that scientist 1 knows when scientist 2

completes a stage, even when she keeps the result secret. If this were not assumed, the two decision nodes for scientist 1 would be indistinguishable to her. This is different from Banerjee et al.’s assumption that “agents only know their own progress, and what is shared with them by others” (Banerjee et al. 2014, 156).

This suggests subtly different versions of the Game: an Intermediate Results Game with perfect information, in which scientists know when other scientists have completed but not shared intermediate results, and an Intermediate Results Game with imperfect information, in which scientists only know their own progress and what has been publicly shared. In this section I consider the former version of the Game, referred to as $G_{n,k}^p$ (with n the number of scientists and k the number of stages, so figure 3.1 shows $G_{2,2}^p$), which more realistically models scientific fields where pre-registration of studies is common or small communities where everyone knows what everyone is working on. The next section discusses the latter version of the Game, referred to as $G_{n,k}^m$, which more realistically models cases where it is relatively easy to keep intermediate results secret.

If the first scientist to complete stage 1 in figure 3.1 plays H , and the other scientist completes stage 1 before the first scientist finishes stage 2, it is rational for the other scientist to play E : this makes it certain that she will get c_1 units of credit, without reducing either her probability of completing the second stage or her payoff if she does so.

This is a *backwards induction* argument: if a certain node is reached, then it is rational for the scientist who has to make a decision at that node to choose x ; therefore, other scientists may assume that if that node is reached, x will be played. Applying this argument to the terminal decision nodes in figure 3.1 leads to a truncated game tree, as shown in figure 3.2.

Here it is assumed that the second scientist to complete stage 1 always plays strategy E . The expected payoff of that strategy for the scientist who just completed stage 1 is a certain c_1 units of credit, plus a further c_2 units of credit if she is first to complete stage 2. The payoff for the other scientist is c_2 times the probability that she is first to complete stage 2.

Now consider the decision scientist 1 has to make if she completes stage 1 first. If she plays strategy E , her payoff is c_1 for sure plus an additional c_2 with probability λ_{12}/σ_2 , so her expected payoff is

$$c_1 + c_2 \frac{\lambda_{12}}{\sigma_2} = c_1 \frac{\lambda_{12}}{\sigma_{21}} + c_1 \frac{\lambda_{21}}{\sigma_{21}} + c_2 \frac{\lambda_{12}}{\sigma_2}.$$

If she plays strategy H instead, her payoff is C with probability λ_{12}/σ_{21} and $c_2\lambda_{12}/\sigma_2$




Figure 3.2: Truncated game tree for $G_{2,2}^p$.

with probability λ_{21}/σ_{21} . So in this case her expected payoff is

$$c_1 \frac{\lambda_{12}}{\sigma_{21}} + c_2 \frac{\lambda_{12}}{\sigma_{21}} + c_2 \frac{\lambda_{12}}{\sigma_2} \frac{\lambda_{21}}{\sigma_{21}} = c_1 \frac{\lambda_{12}}{\sigma_{21}} + c_2 \frac{\lambda_{12}}{\sigma_2} + c_2 \frac{\lambda_{12}}{\sigma_{21}} \frac{\lambda_{22}}{\sigma_2}.$$

It follows that the expected payoff of E is at least as high as the expected payoff of H for scientist 1 if and only if

$$c_1 \frac{\lambda_{21}}{\sigma_{21}} \geq c_2 \frac{\lambda_{12}}{\sigma_{21}} \frac{\lambda_{22}}{\sigma_2} \quad \text{or equivalently} \quad \frac{c_1 \lambda_{21}}{c_2 \lambda_{22}} \geq \frac{\lambda_{12}}{\sigma_2}.$$

Similarly, if scientist 2 completes stage 1 first, it is rational for her to play strategy E if and only if

$$\frac{c_1 \lambda_{11}}{c_2 \lambda_{12}} \geq \frac{\lambda_{22}}{\sigma_2}.$$

In general, these inequalities need not be satisfied. For example, if the credit reward for the two stages is equal ($c_1 = c_2$) but the second stage can be completed twice as quickly ($\lambda_{12} = 2\lambda_{11}$ and $\lambda_{22} = 2\lambda_{21}$) then it is only rational to play strategy E if the other scientist is faster at solving stage 2—the more productive scientist has an

incentive not to share.

If the credit rewards are equal but the second stage can be completed 50 % quicker than the first ($\lambda_{12} = \frac{3}{2}\lambda_{11}$ and $\lambda_{22} = \frac{3}{2}\lambda_{21}$) then the more productive scientist has an incentive not to share if she is at least twice as productive as the other scientist. But sharing is rational for both scientists whenever the scientists' productivity rates are less than a factor two different.

So whether it is rational to share may depend on the productivity rates of the scientists. In particular, the more productive scientist has the most incentive to keep results secret, as these examples illustrate (cf. Banerjee et al. 2014, corollary 2.2). However, this is only worth doing if the potential gains (the chance of getting credit for later stages) are big enough.

Note that if $c_1\lambda_{11} \geq c_2\lambda_{12}$ and $c_1\lambda_{21} \geq c_2\lambda_{22}$ the inequalities given above are always satisfied. That is, if the credit awarded for the first stage is at least as high as the credit given for the second stage (relative to the difficulty of the two stages), then both scientists have an incentive to share their intermediate results, regardless of their productivity rates.

I call this reward structure Proportional Credit, and I will show that it incentivizes sharing in a wide range of cases.

Assumption 3.1 (Proportional Credit). *The productivity parameters and the credit rewards stand in the following relation: for every scientist i and for any pair of stages $j < j'$,*

$$c_j\lambda_{ij} \geq c_{j'}\lambda_{ij'}.$$

The name of this assumption refers to the special case where the two sides of the above expression are equal: in this case credit is given for each stage exactly in proportion to its difficulty. This is arguably the most interesting case, but as I will show the incentive to share exists not only in this special case, but also whenever earlier stages are worth relatively more credit.

If Proportional Credit is satisfied, the backwards induction solution of $G_{2,2}^p$ is for both scientists to play E at both of their decision nodes. Like any backwards induction solution, this is an equilibrium.

This Game has other equilibria. If both scientists play strategy E if they are the first to solve stage 1 then the bottom decision nodes in figure 3.1 are never reached. If one or both scientists play strategy H or a mixed strategy at their bottom decision node the resulting assignment of strategies may still be an equilibrium. But these equilibria are behaviorally indistinguishable from the one identified by backwards induction: they differ only in that some scientists make different choices at decision nodes that will not actually be reached in the Game.

It can be shown that a similar analysis goes through when the number of scientists or the number of stages is changed, as stated in the following theorem.

Theorem 3.2. *Let $n \geq 2$, $k \geq 1$ and assume Proportional Credit.*

- (a) $G_{n,k}^p$ has a (unique) backwards induction solution in which all scientists play strategy E at every decision node.
- (b) $G_{n,k}^p$ has no equilibria (in pure or mixed strategies) that are behaviorally distinct from the backwards induction solution.

Part (b) of this theorem is a consequence of theorem A.4, which is proved in appendix A. Part (a) requires a separate proof, which is given in appendix B.

4 The Intermediate Results Game with Imperfect Information

In this section I analyze a version of the Game in which scientists do not know if other scientists have any unpublished results. Figure 4.1 shows the extensive form of the Intermediate Results Game with imperfect information in its simplest form ($G_{2,2}^m$). The only difference compared to figure 3.1 is the appearance of the dashed lines between decision nodes. These indicate so-called *information sets*: sets of decision nodes that the scientist who has to make a decision cannot distinguish between (i.e., she must play the same strategy at each node in the set).

As a result the number of (pure) strategies is reduced. Previously, a scientist had four possible strategies: she could play either E or H at either of her decision nodes. Now each scientist has just one information set, and two possible strategies: E or H .

Table 4.1: Expected Credit in $G_{2,2}^m$

	E	H
E	$\left(\frac{c_1 \lambda_{11}}{\sigma_1} + \frac{c_2 \lambda_{12}}{\sigma_2}, \frac{c_1 \lambda_{21}}{\sigma_1} + \frac{c_2 \lambda_{22}}{\sigma_2} \right)$	$\left(\frac{c_1 \sigma_2 + c_2 \lambda_{12}}{\sigma_1 \sigma_2 / \lambda_{11}} \frac{\sigma_{12} + \lambda_{21}}{\sigma_{12}}, \frac{C \lambda_{21} \lambda_{22}}{\sigma_1 \sigma_{12}} + \frac{c_2 \lambda_{11} \lambda_{22}}{\sigma_1 \sigma_2} \frac{\sigma_{12} + \lambda_{21}}{\sigma_{12}} \right)$
H	$\left(\frac{C \lambda_{11} \lambda_{12}}{\sigma_1 \sigma_{21}} + \frac{c_2 \lambda_{21} \lambda_{12}}{\sigma_1 \sigma_2} \frac{\sigma_{21} + \lambda_{11}}{\sigma_{21}}, \frac{c_1 \sigma_2 + c_2 \lambda_{22}}{\sigma_1 \sigma_2 / \lambda_{21}} \frac{\sigma_{21} + \lambda_{11}}{\sigma_{21}} \right)$	$\left(\frac{C \lambda_{11} \lambda_{12}}{\sigma_1 \sigma_{21} \sigma_2} \left(\sigma_2 + \frac{\sigma_1 + \sigma_2}{\sigma_{12} / \lambda_{21}} \right), \frac{C \lambda_{21} \lambda_{22}}{\sigma_1 \sigma_{12} \sigma_2} \left(\sigma_2 + \frac{\sigma_1 + \sigma_2}{\sigma_{21} / \lambda_{11}} \right) \right)$

Note.—Scientist 1's strategy as the rows, and scientist 2's strategy as the columns.




Figure 4.1: Extensive form of $G_{2,2}^m$.

Table 4.1 gives the expected credit for each scientist as a function of the scientists' choice of strategy. With some algebra it can be shown that the strategy profile (E, E) on which both scientists share is an equilibrium if and only if

$$\frac{c_1 \lambda_{11}}{c_2 \lambda_{12}} \geq \frac{\lambda_{22}}{\sigma_2} \quad \text{and} \quad \frac{c_1 \lambda_{21}}{c_2 \lambda_{22}} \geq \frac{\lambda_{12}}{\sigma_2}.$$

These inequalities are identical to those found in section 3 and are those established by Banerjee et al. (2014, theorem 2.1) as sufficient conditions for a sharing equilibrium. Moreover, if the above inequalities are strict then the equilibrium is both *unique* and *strict* (an equilibrium is strict if unilaterally deviating from the equilibrium strictly decreases a scientist's expected credit). Note that the strict version of both inequalities is satisfied under Proportional Credit.

In the general version of the Game (with n and k possibly greater than 2) each scientist has to formulate a strategy (E or H) for each information set. At an information set, the scientist knows which stage was the last one to be completed

and shared by some scientist, and how many stages she has since completed herself. However, she does not know how many stages have been completed but not shared by other scientists. As a result, the number of possible strategies is smaller than in the game of perfect information of section 3. It turns out that under Proportional Credit the general version of the Game also has only one equilibrium.

Theorem 4.1. *Let $n \geq 2$, $k \geq 1$ and assume Proportional Credit.*

- (a) $G_{n,k}^m$ has an equilibrium in which all scientists play strategy E at every information set (this follows from Banerjee et al. 2014, theorem 2.1).
- (b) $G_{n,k}^m$ has no other equilibria (in pure or mixed strategies).
- (c) The equilibrium is strict.

This theorem is a consequence of theorem A.4, which is proved in appendix A.

5 A Network of Stages

This section relaxes the assumption that the stages can be completed in only one order. Here I assume instead that the k stages are arranged in a directed acyclic network with the stages represented as edges. Work on a given stage j can only be started if all stages (edges) ending at the node at which j begins have been completed. (Equivalently, stages can be represented as nodes, with each edge indicating that its beginning node is a prerequisite for starting its ending node.) If g is any such network describing the order in which $k \geq 1$ stages can be completed by $n \geq 2$ scientists, then $G_{n,g}^p$ denotes this more general version of the Intermediate Results Game with perfect information, and $G_{n,g}^m$ denotes the more general version of the Intermediate Results Game with imperfect information.

Define $r_{ij} = c_j \lambda_{ij}$ to be the *reward rate* for scientist i while working on stage j . The following assumption is a variation of the monotonicity assumption made by Banerjee et al. (2014, 158).

Assumption 5.1 (Monotonicity). *There is a strict ordering of the stages by their reward rate, this ordering is the same for all scientists, and a given stage's reward rate is always lower than any of its prerequisites. Formally,*

- (a) For any pair of stages j and j' , either $r_{ij} < r_{ij'}$ for all scientists i , or $r_{ij'} < r_{ij}$ for all scientists i .

(b) If completing stage j is required for starting stage j' , then $r_{ij'} < r_{ij}$ for all scientists i .

Monotonicity imposes a natural ordering on the stages. Making use of this fact, from here on I assume without loss of generality that the stages are relabeled such that $j < j'$ if and only if $r_{ij'} < r_{ij}$ for all scientists i . Note that, since a stage's reward rate is lower than its prerequisites, completing the stages in the order $1, 2, \dots, k$ is consistent with the restrictions imposed by the network g .

At any given time in the Game, let J_i denote the stages available for scientist i to work on (i.e., stages that are unsolved but all prerequisites have been solved by i or solved and shared by another scientist). Let j_i^* denote the available stage that has the highest reward rate for scientist i . Given Monotonicity and the relabeling specified above, it follows that $j_i^* = \min J_i$.

In general, it need not be the case that the expected duration of the Game is minimized if all scientists immediately share any intermediate stages they complete (unlike before). For example, suppose there are two scientists and two stages that can be completed independently. Suppose that scientist 1 works faster on stage 1 and scientist 2 works faster on stage 2; e.g., $\lambda_{11} = \lambda_{22} = 2$ and $\lambda_{12} = \lambda_{21} = 1$. Then expected duration is minimized (at $7/12$) if scientist 1 begins by working on stage 1 and scientist 2 begins by working on stage 2 (and both share if they complete it). But if the reward for completing stage 1 is more than twice the reward for completing stage 2, the unique equilibrium is for both scientists to begin working on stage 1 (and share if they complete it), which yields an expected duration of $2/3 > 7/12$.

However, Banerjee et al. (2014) show that there is a special case of the Game in which sharing does minimize expected duration. Say that the *separable aptitudes* assumption is satisfied if there are parameters a_i ("aptitude", depending only on the scientist) and s_j ("simplicity", depending only on the stage) such that for all i and j , $\lambda_{ij} = a_i s_j$. If separable aptitudes holds, and there exists an equilibrium in which all scientists immediately share any intermediate stages they complete, then expected duration is minimized in this equilibrium (Banerjee et al. 2014, theorem 4.3).

My argument for a unique equilibrium in this section requires expected duration to be minimized. But it does not otherwise depend on Banerjee et al.'s separable aptitudes assumption. So I simply assume that the profile of interest minimizes expected duration (with Banerjee et al.'s result guaranteeing that there are at least some cases where this assumption is satisfied).

Stating the formal version of this assumption requires some notation. Let s_i^E denote the strategy in which scientist i works on stage $\min J_i$ (i.e., the unshared stage with the lowest label) at any given time and plays strategy E whenever she completes a stage. Let S^E denote the strategy profile in which every scientist i plays

strategy s_i^E .

Let $W_j(S)$ denote the waiting time until some scientist shares the solution to stage j , assuming strategy profile S is being played. $W_j(S)$ is a random variable because the time it takes individual scientists to solve stages is random, and it depends on the strategy profile being played because even if a scientist solves a stage at time t she may not share it at that time depending on her strategy. It follows that $\max_j W_j(S)$ is the duration of the Game (note that, in general, stages may be shared out of order).

Assumption 5.2 (Minimal Time). *The expected completion time is minimized under S^E . That is, for all strategy profiles S ,*

$$\mathbb{E} \left(\max_j W_j(S) \right) \geq \mathbb{E} \left(\max_j W_j(S^E) \right).$$

Under these assumptions, theorems 3.2 and 4.1 can be generalized to the context of a directed acyclic network describing the order of the stages.

Theorem 5.3. *Let $n \geq 2$ and let g be a network for $k \geq 1$ stages. Assume Monotonicity and Minimal Time.*

- (a) *The strategy profile S^E is an equilibrium for $G_{n,g}^p$ and $G_{n,g}^m$.*
- (b) *In $G_{n,g}^p$, the equilibrium S^E is unique up to deviations off the equilibrium path.*
- (c) *In $G_{n,g}^m$, the equilibrium S^E is unique and strict.*

Part (a) of theorem 5.3 is not a completely novel result, as Banerjee et al. (2014, theorem 3.2) prove this for $G_{n,g}^m$ under slightly different conditions. They do not require Minimal Time and their statement of the Monotonicity assumption is subtly different. Roughly speaking though, my result shows that the sharing equilibrium found by Banerjee et al. is unique whenever the equilibrium minimizes the expected completion time of the research project.

6 A Boundedly Rational Perspective

In this section I show that the versions of the Intermediate Results Game with imperfect information ($G_{n,k}^m$ and $G_{n,g}^m$) are weakly better reply games. This means that boundedly rational scientists are likely to learn to play the equilibrium strategy, i.e., will learn to share their intermediate results. I show this in detail for the learning rule probe and adjust.

Huttegger (2013) defines the following concepts. Let G be a game. A *weakly better reply path* is a sequence of strategy profiles (S^1, \dots, S^ℓ) for G such that for any $j < \ell$, profile S^j differs from profile S^{j+1} only in one scientist's strategy, say scientist i , and $u_i(S^{j+1}) \geq u_i(S^j)$, i.e., scientist i changes to a strategy that is a (weakly) better reply to the other scientists' strategies. G is a *weakly better reply game* if there exists a weakly better reply path from any strategy profile S to a strict equilibrium.

Theorem 6.1. *Let $n \geq 2$ and let $k \geq 1$.*

- (a) *Assuming Proportional Credit, $G_{n,k}^m$ is a weakly better reply game.*
- (b) *Let g be a network for k stages and assume Monotonicity and Minimal Time. Then $G_{n,g}^m$ is a weakly better reply game.*

Proof. The result is a corollary of theorem A.4. The unique and strict equilibrium in both games is S^E . Let $S \neq S^E$ be any strategy profile. By theorem A.4 there exists a scientist i who is not playing the equilibrium strategy in S but for whom switching to the equilibrium strategy is a better reply. Let S' be the strategy profile that differs from S only in that scientist i has switched to the equilibrium strategy. If $S' \neq S^E$ the same reasoning can be applied again. This generates a weakly better reply path of maximum length n from S to S^E . \square

Now suppose a group of n scientists are repeatedly playing some game G and adjusting their strategy in light of previous payoffs. A scientist using *probe and adjust* follows a simple procedure: on each round, play the same strategy as the round before with probability $1 - \varepsilon$, or probe a new strategy with some “small” probability $\varepsilon > 0$. In case of a probe, pick a new strategy uniformly at random from all possible strategies. After playing this strategy for one round, evaluate the probe: if the payoff for the probing round is higher than the payoff in the previous round, keep the probed strategy (at least until the next probe); if the payoff is lower, return to the old strategy; if payoffs are equal, return to the old strategy with probability $q \in (0, 1)$ and retain the probe with probability $1 - q$ (note that this is not quite the same as asking whether the probed strategy is a better reply to the other scientists' strategies, due to the possibility of simultaneous probes).

Suppose all scientists use probe and adjust to determine their strategy in repeated plays of a weakly better reply game G . Assume moreover that all scientists use the same values of ε and q (this assumption can be relaxed, see Huttegger et al. 2014, 837–838). Let S^t be the profile of strategies played on round t . Then the following result holds.

Theorem 6.2 (Huttegger (2013)). *For any probability $p < 1$, if the probe probability $\varepsilon > 0$ is sufficiently small, then the profile S^t is a strict equilibrium of G for all sufficiently large t with probability at least p .*

Since $G_{n,k}^m$ and $G_{n,g}^m$ are weakly better reply games by theorem 6.1, theorem 6.2 applies to them. The (unique) strict equilibria of these games have all scientists share their intermediate results as soon as they complete them. So if, on a given round, the scientists are playing the equilibrium profile, they may be said to have learned to share their intermediate results. Theorem 6.2 says that the probability of this happening can be made arbitrarily high by choosing a small enough probe probability and a long enough waiting time. Moreover, the theorem says that once the scientists learn to share their intermediate results they continue to do so on most subsequent rounds.

Because the equilibrium is both strict and unique, various other learning rules and evolutionary dynamics will display similar behavior: scientists will learn to share their intermediate results and continue to do so with high probability. Examples include fictitious play, the best-response dynamics, and the replicator dynamics.

7 Measuring Credit Per Unit Time

As Boyer-Kassem and Imbert (2015, section 4) have argued, in order to determine what rational credit-maximizing scientists would do it may be better to assume that scientists are maximizing expected credit per unit of time (rather than total expected credit from the project). This is because after scientists finish the present research project, they presumably start working on a new one with its own expected credit reward, and hence spending more time on the present project carries an opportunity cost. This section shows that the results presented so far hold also in a version of the Intermediate Results Game in which payoff is measured in expected credit per unit of time.

Let $G_{n,k}^{pt}$ denote the adapted version of the Game with perfect information and let $G_{n,k}^{mt}$ denote the adapted version of the Game with imperfect information. These games have the same strategy spaces as $G_{n,k}^p$ and $G_{n,k}^m$ respectively, differing only in that the payoff functions measure credit per unit of time. (In this section I restrict attention to the case where stages can only be completed sequentially.)

Theorem 7.1. *Let $n \geq 2$, $k \geq 1$ and assume Proportional Credit.*

- (a) $G_{n,k}^{pt}$ has a (unique) backwards induction solution in which all scientists play strategy E at every decision node.

(b) $G_{n,k}^{pt}$ has no equilibria (in pure or mixed strategies) that are behaviorally distinct from the backwards induction solution.

Theorem 7.2. *Let $n \geq 2$, $k \geq 1$ and assume Proportional Credit.*

(a) $G_{n,k}^{mt}$ has an equilibrium in which all scientists play strategy E at every information set.

(b) $G_{n,k}^{mt}$ has no other equilibria (in pure or mixed strategies).

(c) The equilibrium is strict.

Theorems 7.1.b and 7.2 are consequences of theorem A.5, which is proved in appendix A. Theorem 7.1.a is proved in appendix B.

The same short proof given in section 6 can be used to conclude from theorem A.5 that $G_{n,k}^{mt}$ is a weakly better reply game. Hence scientists using probe and adjust will find the sharing equilibrium with high probability by theorem 6.2.

This shows that, when credit is measured per unit of time, there is an incentive to share in the Intermediate Results Game under the same conditions for which I showed it to exist in the case where total credit at the end of the Game is the key quantity.

8 Conclusion

Despite the claims of Strevens (2017) and others it turns out that there is a range of circumstances under which the sharing of intermediate results is incentivized for credit-maximizing scientists. In particular, a sufficient condition appears to be that intermediate results are rewarded with at least as much credit as the results that depend on them, relative to their difficulty.

Banerjee et al. (2014) had already shown the existence of a sharing equilibrium in these circumstances. I have shown that this equilibrium is unique, a strong result for a game-theoretic model. Moreover, I have shown that a similar equilibrium exists (and is unique up to deviations off the equilibrium path) for a version of the Game with perfect information and for a version of the Game in which credit is measured per unit of time. Finally, I have shown that boundedly rational scientists will also learn to share their intermediate results in these games, because they are weakly better reply games.

A A Unique Nash Equilibrium

This appendix proves theorem A.4, which shows that in any strategy profile in which not all scientists immediately share any intermediate results they complete, some scientist can improve her expected payoff by switching to a sharing strategy. This result holds for each version of the Game under slightly different conditions, but as I will show these conditions make it such that the same proof works in all cases.

For the duration of this section, let $n \geq 2$ be the number of scientists, let $k \geq 1$ be the number of stages, and let g be a directed acyclic network with k stages. Consider the games $G_{n,k}^p$, $G_{n,k}^{pt}$, $G_{n,k}^m$ and $G_{n,k}^{mt}$ in which the k stages have to be completed sequentially and the games $G_{n,g}^p$ and $G_{n,g}^m$ in which g describes the order in which the k stages can be completed.

As is commonly done in game theory, I use $u_i(s_i, s_{-i})$ to denote the payoff (expected units of credit at the end of the game) to scientist i if s_i gives her strategy and s_{-i} gives the strategies of all scientists other than i (call this an “incomplete strategy profile”). I use this notation interchangeably with $u_i(S)$, the payoff to scientist i given a complete strategy profile S .

I will abuse notation somewhat to make the proof work for the different versions of the Game. For any scientist i , let s_i^E denote the equilibrium strategy (or rather the putative equilibrium strategy—part of what I will show is that it is indeed an equilibrium strategy). So in $G_{n,k}^p$ and $G_{n,k}^{pt}$, s_i^E is the strategy that plays strategy E at every decision node; in $G_{n,k}^m$ and $G_{n,k}^{mt}$ this strategy plays E at every information set; in $G_{n,g}^p$ this strategy always works on stage $\min J_i$ and plays E at every decision node; and in $G_{n,g}^m$ this strategy always works on stage $\min J_i$ and plays E at every information set.

Let s_{-i}^E denote the incomplete strategy profile in which every scientist $i' \neq i$ plays strategy $s_{i'}^E$, and let $S^E = (s_i^E, s_{-i}^E)$ be the (putative) equilibrium. The first lemma gives an explicit formula for scientist i 's payoff in the profile S^E . It does not depend on any specific assumptions.

Lemma A.1. *In the games $G_{n,k}^p$, $G_{n,k}^m$, $G_{n,g}^p$, and $G_{n,g}^m$ the payoff to scientist i under the strategy profile S^E is*

$$u_i(S^E) = \sum_{j=1}^k \frac{r_{ij}}{\sigma_j}.$$

Proof. Under strategy profile S^E , the scientists work on the stages in the order they are labeled (starting with stage 1, ending with stage k), sharing each result as soon as they complete it. Under these circumstances, a scientist i can be viewed as a

nonstationary reward process producing payoff at a rate of r_{ij} units of payoff per unit of time if j is the current stage.

Under S^E , the time it takes the scientists to solve and share stage j is exponentially distributed with parameter σ_j (cf. section 2). So the expected time spent on stage j is $1/\sigma_j$. \square

The second lemma shows that, if S^E minimizes the expected completion time of the overall research project, it also minimizes the waiting time until the first j stages have been shared.

Lemma A.2. *Let $S \neq S^E$ be some arbitrary strategy profile in $G_{n,k}^p$, $G_{n,k}^{pt}$, $G_{n,k}^m$, $G_{n,k}^{mt}$, $G_{n,g}^p$, or $G_{n,g}^m$. Depending on the version of the Game, add the following additional assumptions:*

- For $G_{n,g}^p$ and $G_{n,g}^m$, assume Minimal Time.
- For $G_{n,k}^p$, $G_{n,k}^{pt}$, and $G_{n,g}^p$, assume that S involves a deviation on the equilibrium path relative to S^E .

Then for any j ,

$$\mathbb{E} \left(\max_{j' \leq j} W_{j'}(S) \right) \geq \mathbb{E} \left(\max_{j' \leq j} W_{j'}(S^E) \right) = \sum_{j'=1}^j \frac{1}{\sigma_{j'}}.$$

Moreover, there exists a value of j for which the above inequality is strict. In particular, for $G_{n,k}^p$, $G_{n,k}^{pt}$, $G_{n,k}^m$, and $G_{n,k}^{mt}$, strict inequality holds for $j = k$.

Proof. As noted above, under S^E the scientists complete and share the stages in the order they are labeled, and each stage j is expected to take $1/\sigma_j$, so

$$\mathbb{E} \left(W_j(S^E) \right) - \mathbb{E} \left(W_{j-1}(S^E) \right) = \frac{1}{\sigma_j},$$

and

$$\mathbb{E} \left(\max_{j' \leq j} W_{j'}(S^E) \right) = \mathbb{E} \left(W_j(S^E) \right) = \sum_{j'=1}^j \frac{1}{\sigma_{j'}}.$$

This establishes the equality.

To prove the inequality, suppose for reductio that there exists a stage j such that

$$\mathbb{E} \left(\max_{j' \leq j} W_{j'}(S) \right) < \mathbb{E} \left(\max_{j' \leq j} W_{j'}(S^E) \right).$$

Now define the strategy profile S' as follows. Under S' , each scientist i plays strategy s'_i , which is a combination of s_i (the strategy scientist i plays under S) and s_i^E , as follows:

$$s'_i = \begin{cases} s_i & \text{at any time until } \max_{j' \leq j} W_{j'}(S'), \\ s_i^E & \text{at any time after } \max_{j' \leq j} W_{j'}(S'). \end{cases}$$

Then

$$\begin{aligned} \mathbb{E} \left(\max_{j'} W_{j'}(S') \right) &= \mathbb{E} \left(\max_{j' \leq j} W_{j'}(S') \right) + \mathbb{E} \left(\max_{j'} W_{j'}(S') - \max_{j' \leq j} W_{j'}(S') \right) \\ &= \mathbb{E} \left(\max_{j' \leq j} W_{j'}(S) \right) + \sum_{j'=j+1}^k \frac{1}{\sigma_{j'}} \\ &< \sum_{j'=1}^k \frac{1}{\sigma_{j'}} = \mathbb{E} \left(\max_{j'} W_{j'}(S^E) \right) \end{aligned}$$

This contradicts the fact that S^E minimizes the expected completion time of the research project, which is known to be true in the case in which stages are completed sequentially (cf. Banerjee et al. 2014, 159) and which I have assumed to be true through Minimal Time in the case of a network of stages. So the inequality holds.

It remains to show that the inequality is strict for some value of j . Here I distinguish two cases.

1. Under S , all scientists complete the stages in the order they are labeled.

Since S differs from S^E (on the equilibrium path), it follows that there is a positive probability under S that a situation arises in which a scientist i plays strategy H . If this happens, there is a positive probability that some other scientist finishes and shares the stage that i failed to share before scientist i completes another stage. This increases the expected completion time of the overall project (cf. Banerjee et al. 2014, 159), i.e.,

$$\mathbb{E} \left(\max_{j' \leq k} W_{j'}(S) \right) > \mathbb{E} \left(\max_{j' \leq k} W_{j'}(S^E) \right).$$

In the games $G_{n,k}^p$, $G_{n,k}^{pt}$, $G_{n,k}^m$, and $G_{n,k}^{mt}$, the stages can only be completed in one order. So for these games strict inequality holds in particular for the expected completion time of the overall project.

2. Under S , at least one scientist does not complete the stages in the order they are labeled.

Let j be the lowest stage number such that at least one scientist does not start working on stage j as soon as she knows the solution to stages $1, \dots, j-1$. From the inequality established above,

$$\mathbb{E} \left(\max_{j' \leq j-1} W_{j'}(S) \right) \geq \mathbb{E} \left(\max_{j' \leq j-1} W_{j'}(S^E) \right) = \sum_{j'=1}^{j-1} \frac{1}{\sigma_{j'}}.$$

Since not every scientist immediately starts working on stage j , the expected time until stage j is shared must be strictly greater than $1/\sigma_j$. So

$$\mathbb{E} \left(\max_{j' \leq j} W_{j'}(S) \right) > \sum_{j'=1}^{j-1} \frac{1}{\sigma_{j'}} + \frac{1}{\sigma_j} = \mathbb{E} \left(\max_{j' \leq j} W_{j'}(S^E) \right). \quad \square$$

Note that the preceding lemmas do not require Proportional Credit or Monotonicity; they are true for all (positive) productivity rates and for all (positive) credit rewards. The next lemma shows that when these assumptions are introduced, if not every scientist plays the (putative) equilibrium strategy, scientists who do get a higher payoff than they do in lemma A.1.

Lemma A.3. *Let i be a scientist in $G_{n,k}^p$, $G_{n,k}^m$, $G_{n,g}^p$, or $G_{n,g}^m$, and assume scientist i plays strategy s_i^E . Let s_{-i} denote an incomplete strategy profile such that $S = (s_i^E, s_{-i}) \neq S^E$. Depending on the version of the Game, add the following additional assumptions:*

- For $G_{n,g}^p$ and $G_{n,g}^m$, assume Monotonicity and Minimal Time.
- For $G_{n,k}^p$ and $G_{n,k}^m$, assume Proportional Credit.
- For $G_{n,k}^p$ and $G_{n,g}^p$, assume that S involves a deviation on the equilibrium path relative to S^E .

Then $u_i(S) > u_i(S^E)$.

Proof. Just like in the proof of lemma A.1, it is useful to view scientist i as a nonstationary reward process. When she is working on stage j , she produces payoff at a rate of r_{ij} units of payoff per unit of time. By Proportional Credit or Monotonicity, this rate is non-increasing throughout the Game for scientist i .

So scientist i expects to get a payoff of $r_{ik} > 0$ per unit of time throughout the Game, which has an expected duration of $\mathbb{E}(\max_{j'} W_{j'}(S))$. Moreover, she expects to get an additional $r_{ik-1} - r_{ik}$ per unit of time as long as stage $k-1$ remains unshared,

which is expected to take $\mathbb{E}(\max_{j' \leq k-1} W_{j'}(S))$ time units. And so on. Hence (setting $r_{ik+1} = 0$ for notational convenience)

$$u_i(S) = \sum_{j=1}^k (r_{ij} - r_{ij+1}) \mathbb{E} \left(\max_{j' \leq j} W_{j'}(S) \right).$$

Lemma A.2 says that $\mathbb{E}(\max_{j' \leq j} W_{j'}(S)) \geq \sum_{j'=1}^j 1/\sigma_{j'}$, with strict inequality for some value of j . In the case of $G_{n,g}^p$ and $G_{n,g}^m$, since $r_{ij} - r_{ij+1} > 0$ for all j by Monotonicity, plugging this in yields

$$u_i(S) > \sum_{j=1}^k (r_{ij} - r_{ij+1}) \sum_{j'=1}^j \frac{1}{\sigma_{j'}}.$$

In the case of $G_{n,k}^p$ and $G_{n,k}^m$, the above inequality also holds due to the following facts: $r_{ij} - r_{ij+1} \geq 0$ by Proportional Credit, $\mathbb{E}(\max_{j' \leq k} W_{j'}(S)) > \sum_{j'=1}^k 1/\sigma_{j'}$ by lemma A.2, and $r_{ik} > 0$.

Interchanging the sums then gives the desired result:

$$u_i(S) > \sum_{j'=1}^k \sum_{j=j'}^k (r_{ij} - r_{ij+1}) \frac{1}{\sigma_{j'}} = \sum_{j'=1}^k \frac{r_{ij'}}{\sigma_{j'}} = u_i(S^E). \quad \square$$

The main result follows from lemma A.3. It shows that any strategy profile that differs from S^E (on the equilibrium path) is not an equilibrium.

Theorem A.4. *Let $S \neq S^E$ be some arbitrary strategy profile in $G_{n,k}^p$, $G_{n,k}^m$, $G_{n,g}^p$, or $G_{n,g}^m$. Depending on the version of the Game, add the following additional assumptions:*

- For $G_{n,g}^p$ and $G_{n,g}^m$, assume Monotonicity and Minimal Time.
- For $G_{n,k}^p$ and $G_{n,k}^m$, assume Proportional Credit.
- For $G_{n,k}^p$ and $G_{n,g}^p$, assume that S involves a deviation on the equilibrium path relative to S^E .

Then there exists at least one scientist i playing strategy $s_i \neq s_i^E$ such that she would be strictly better off playing strategy s_i^E :

$$u_i(s_i^E, s_{-i}) > u_i(s_i, s_{-i}).$$

Proof. Recall that the Intermediate Results Game is zero-sum (if payoff is measured in total credit): regardless of strategies, there are C units of credit to be divided, and so if one scientist's payoff increases, another's decreases. Combined with lemmas A.1 and A.3 this yields the theorem. Distinguish three cases:

1. There is only one scientist i playing a (pure or mixed) strategy $s_i \neq s_i^E$.

In this case every scientist i' other than scientist i is playing strategy $s_{i'}^E$ and so by lemma A.3 is getting a payoff greater than $u_{i'}(S^E)$. Because the Game is zero-sum, it follows that $u_i(s_i, s_{-i}) < u_i(S^E)$. By lemma A.1, $u_i(s_i^E, s_{-i}) = u_i(S^E)$, and the result follows.

2. There is at least one scientist i' playing strategy $s_{i'}^E$ and at least two scientists playing some other strategy.

In this case any scientist i' who is playing strategy $s_{i'}^E$ is getting a payoff greater than $u_{i'}(S^E)$ by lemma A.3. Because the Game is zero-sum, at least one of the remaining scientists, say scientist i , must be getting a payoff less than $u_i(S^E)$. But if scientist i changed her strategy to s_i^E , by lemma A.3 she would get a payoff $u_i(s_i^E, s_{-i}) > u_i(S^E)$, establishing the result.

3. Every scientist i' is playing some strategy $s_{i'} \neq s_{i'}^E$.

Because the Game is zero-sum, it is impossible for every scientist i' to be getting a greater payoff than $u_{i'}(S^E)$. So there is at least one scientist, say scientist i , such that $u_i(s_i, s_{-i}) \leq u_i(S^E)$. By lemma A.3, $u_i(s_i^E, s_{-i}) > u_i(S^E)$, and the result follows. \square

Theorem A.4 may be used to prove theorems 3.2.b, 4.1, and 5.3.

Proof of theorem 3.2.b. Let S be any strategy profile for the game $G_{n,k}^p$. If S differs from S^E on the equilibrium path (with positive probability, in the case of mixed strategies), then at least one scientist has an incentive to change her strategy by theorem A.4, and so S is not an equilibrium. \square

Proof of theorem 4.1. Let S be any strategy profile (of pure or mixed strategies) for the game $G_{n,k}^m$. If $S \neq S^E$, then at least one scientist has an incentive to change her strategy by theorem A.4, and so S is not an equilibrium.

That S^E is a strict equilibrium also follows from theorem A.4 by considering the special case where $s_{-i} = s_{-i}^E$. This shows that a scientist i who deviates unilaterally makes herself strictly worse off. \square

Proof of theorem 5.3. Consider first the case of imperfect information. Let S be any profile for the game $G_{n,g}^m$. If $S \neq S^E$, then at least one scientist has an incentive to change her strategy by theorem A.4, and so S is not an equilibrium. That S^E is a strict equilibrium also follows from theorem A.4 by considering the special case where $s_{-i} = s_{-i}^E$.

In the case of perfect information, any profile which contains deviations from the equilibrium path is not an equilibrium by theorem A.4, as at least one scientist has an incentive to change her strategy. On any profile that differs from S^E only off the equilibrium path all scientists get the same payoff as on S^E . It follows that S^E is an equilibrium of $G_{n,g}^p$ and any other equilibria differ from S^E only off the equilibrium path. \square

To conclude this appendix, I show that theorem A.4 extends to the case in which credit is measured per unit of time. This yields proofs of theorems 7.1.b and 7.2 analogous to the proofs of theorems 3.2.b and 4.1 given above.

Theorem A.5. *Let $S \neq S^E$ be some arbitrary strategy profile in $G_{n,k}^{pt}$ or $G_{n,k}^{mt}$, and assume Proportional Credit. In the case of $G_{n,k}^{pt}$, add the further assumption that S involves deviations on the equilibrium path. Then there exists at least one scientist i playing strategy $s_i \neq s_i^E$ such that she would be strictly better off playing strategy s_i^E :*

$$u_i(s_i^E, s_{-i}) > u_i(s_i, s_{-i}).$$

Proof. As before, view the scientists as nonstationary reward processes with a reward rate (expected credit per unit of time) depending on the stage they are working on and their choice of strategy. As a result, their expected credit per unit of time from the Game is a weighted average of their reward rate at any given time with the weights being the expected time spent working at that reward rate. Hence the sum of the weights (the denominator of the weighted average) is the expected duration of the Game. It follows that the scientists' expected payoff (average credit per unit of time throughout the Game) is equal to their expected total credit divided by the expected duration of the Game (this is a consequence of the Poisson model of productivity; in general expectation does not distribute over quotients).

By theorem A.4 there exists a scientist i whose total expected credit from the Game is higher under strategy profile (s_i^E, s_{-i}) than under (s_i, s_{-i}) . But (by reasoning similar to that given in the proof of lemma A.2) it is also clear that the expected duration of the Game can only decrease if scientist i switches to strategy s_i^E , i.e.,

$$\mathbb{E} \left(\max_j W_j(s_i^E, s_{-i}) \right) \leq \mathbb{E} \left(\max_j W_j(s_i, s_{-i}) \right).$$

But then it follows immediately that scientist i 's expected credit per unit of time must also be higher under (s_i^E, s_{-i}) :

$$u_i(s_i^E, s_{-i}) > u_i(s_i, s_{-i}). \quad \square$$

B The Backwards Induction Solution

Let $G_{n,k}^p$ be the Intermediate Results Game with perfect information, as described in section 3. This appendix proves theorem 3.2.a, which says that under Proportional Credit, the Game has a unique backwards induction solution in which all scientists play strategy E at every decision node. At the end of the appendix I indicate briefly how the proof given here can be used to show that the same backwards induction solution holds for $G_{n,k}^{pt}$ (the version of the Game in which the scientists aim to maximize credit per unit of time).

Begin by fixing a decision node. Let i^* be the scientist making a decision at this decision node, having just completed stage $j^* < k$. Let $j' < j^*$ denote the highest stage number whose solution has been shared ($j' = 0$ if no stages have been shared yet). For $i \neq i^*$, let j_i denote the stage that scientist i is working on at the time scientist i^* completes stage j^* (so $j' < j_i \leq k$).

To prove the theorem, it suffices to show that it is rational for scientist i^* to play strategy E at this decision node, assuming that every scientist (including herself) plays strategy E at all remaining decision nodes. I write $\mathbb{E}(u_{i^*}(E))$ for the expected payoff to scientist i^* if she plays strategy E at the present decision node, and $\mathbb{E}(u_{i^*}(H))$ if she plays H .

Let a_{i^*} denote the credit scientist i^* has accumulated before the present decision node. Let $\Pr_E(R_j)$ denote the probability that scientist i^* eventually claims credit for stage $j > j'$, assuming she plays strategy E at the present decision node. For convenience write λ_{i^*} for $\lambda_{i^*j^*+1}$, the productivity rate of scientist i^* working on stage $j^* + 1$. Let

$$\sigma_E = \lambda_{i^*} + \sum_{i:j_i \leq j^*} \lambda_{ij^*+1} + \sum_{i:j_i \geq j^*+1} \lambda_{ij_i}$$

denote the total productivity of the scientists immediately after the present decision node, if scientist i^* chooses to play strategy E . Note that $\Pr_E(R_j) = 1$ if $j < j^* + 1$

and $\Pr_E(R_{j^*+1}) = \lambda_{i^*}/\sigma_E$. So

$$\begin{aligned}\mathbb{E}(u_{i^*}(E)) &= a_{i^*} + \sum_{j=j'+1}^k c_j \Pr_E(R_j) \\ &= a_{i^*} + \sum_{j=j'+1}^{j^*} c_j + c_{j^*+1} \frac{\lambda_{i^*}}{\sigma_E} + \sum_{j=j^*+2}^k c_j \Pr_E(R_j).\end{aligned}$$

Similarly, let $\Pr_H(R_j)$ denote the probability that scientist i^* eventually claims credit for stage $j > j'$, assuming she plays strategy H at the present decision node. So

$$\mathbb{E}(u_{i^*}(H)) = a_{i^*} + \sum_{j=j'+1}^k c_j \Pr_H(R_j).$$

Two lemmas provide the crucial inequalities to complete the proof.

Lemma B.1. $\Pr_H(R_j) \leq \Pr_E(R_j)$ for any $j > j^* + 1$.

Proof. Divide the scientists into two groups: let $S^+ = \{i \mid j_i \geq j\}$ be those scientists already working on stage j or higher at the time of the present decision node and let $S^- = \{i \mid j_i < j\}$ be those working on stage $j - 1$ or lower, including i^* . Let A be the event that a scientist in S^+ claims credit for stage j and let \bar{A} be the event that a scientist in S^- claims credit for stage j . I make two claims.

First,

$$\Pr_H(R_j \mid \bar{A}) = \Pr_E(R_j \mid \bar{A}) = \lambda_{i^*j} / \sum_{i \in S^-} \lambda_{ij}.$$

This is because, due to the backwards induction assumption, if a scientist in S^- completes stage $j - 1$ she shares the solution, and so all scientists in S^- start working on stage j at the same time.

Second, $\Pr_H(A) \geq \Pr_E(A)$. This is because choosing strategy H at the present decision node can only increase the expected time it takes the scientists in S^- to get to start working on stage j , thus improving the probability that one of the scientists in S^+ completes whatever stage she is working on before that happens.

From these two claims it follows that

$$\Pr_H(R_j) = \Pr_H(R_j \mid \bar{A}) \Pr_H(\bar{A}) \leq \Pr_E(R_j \mid \bar{A}) \Pr_E(\bar{A}) = \Pr_E(R_j). \quad \square$$

Lemma B.2.

$$\sum_{j=j'+1}^{j^*+1} c_j \Pr_H(R_j) < \sum_{j=j'+1}^{j^*} c_j + c_{j^*+1} \frac{\lambda_{i^*}}{\sigma_E}.$$

Proof. The left-hand side indicates the share of the credit from stages $j' + 1$ through $j^* + 1$ that scientist i^* expects to receive if she plays strategy H at the present decision node. Since there is a total of $\sum_{j=j'+1}^{j^*+1} c_j$ units of credit to be divided, scientist i^* 's share can be no higher than that total, minus any portions the other scientists expect to receive.

Let i be a scientist working on stage $j_i \leq j^*$ at the time of the present decision node. From that time until the time stage $j^* + 1$ is shared, she can be viewed as a nonstationary reward process (cf. lemma A.1) producing payoff at a rate of $c_j \lambda_{ij}$ units of payoff per unit of time, where $j \in \{j_i, j_i + 1, \dots, j^* + 1\}$. It follows from Proportional Credit that her expected credit per unit of time is at least $c_{j^*+1} \lambda_{ij^*+1}$ during this time.

If scientist i^* had chosen strategy E at the present decision node, the expected time until stage $j^* + 1$ is shared would be $1/\sigma_E$. But since scientist i^* has chosen strategy H , scientist i and other scientists like her have to finish stages j_i through j^* first. So the expected time until stage $j^* + 1$ is shared is greater than $1/\sigma_E$. Hence scientist i 's expected credit until stage $j^* + 1$ is shared is strictly greater than $c_{j^*+1} \lambda_{ij^*+1} / \sigma_E$.

Now let i be a scientist working on stage $j_i \geq j^* + 1$ at the time of the present decision node. If scientist i^* chooses strategy H this means that credit for (at minimum) stage j^* and $j^* + 1$ remains unclaimed. The probability that scientist i claims credit for stage $j^* + 1$ is at least $\lambda_{ij_i} / \sigma_E$ (this is the probability that she completes a stage before anyone else assuming all scientists immediately learn the solution to stage j^* ; in reality her chance of claiming stage $j^* + 1$ may be higher because some scientists need to finish stage j^* first). Moreover she has some positive probability of claiming credit for stage j^* . So scientist i 's expected credit from stages $j' + 1$ through $j^* + 1$ is strictly greater than $c_{j^*+1} \lambda_{ij_i} / \sigma_E$.

Putting this all together yields

$$\begin{aligned} \sum_{j=j'+1}^{j^*+1} c_j \Pr_H(R_j) &< \sum_{j=j'+1}^{j^*+1} c_j - \sum_{i:j_i \leq j^*} c_{j^*+1} \frac{\lambda_{ij^*+1}}{\sigma_E} - \sum_{i:j_i \geq j^*+1} c_{j^*+1} \frac{\lambda_{ij_i}}{\sigma_E} \\ &= \sum_{j=j'+1}^{j^*} c_j + c_{j^*+1} \left(1 - \sum_{i:j_i \leq j^*} \frac{\lambda_{ij^*+1}}{\sigma_E} - \sum_{i:j_i \geq j^*+1} \frac{\lambda_{ij_i}}{\sigma_E} \right) \\ &= \sum_{j=j'+1}^{j^*} c_j + c_{j^*+1} \frac{\lambda_{i^*}}{\sigma_E}, \end{aligned}$$

where the inequality is strict because at least one of the sets $\{i : j_i \leq j^*\}$ and $\{i : j_i \geq j^* + 1\}$ is nonempty. \square

Now the overall proof can be completed.

Proof of theorem 3.2.a. As noted above,

$$\mathbb{E}(u_{i^*}(H)) = a_{i^*} + \sum_{j=j'+1}^k c_j \Pr_H(R_j).$$

From lemma B.1 it follows that

$$\mathbb{E}(u_{i^*}(H)) \leq a_{i^*} + \sum_{j=j'+1}^{j^*+1} c_j \Pr_H(R_j) + \sum_{j=j^*+2}^k c_j \Pr_E(R_j).$$

Combining this with lemma B.2 yields

$$\begin{aligned} \mathbb{E}(u_{i^*}(H)) &< a_{i^*} + \sum_{j=j'+1}^{j^*} c_j + c_{j^*+1} \frac{\lambda_{i^*}}{\sigma_E} + \sum_{j=j^*+2}^k c_j \Pr_E(R_j) \\ &= \mathbb{E}(u_{i^*}(E)). \end{aligned}$$

This shows that scientist i^* prefers to play strategy E at the present decision node, and hence the induction goes through. Because the preference for E over H is strict, the solution is unique. \square

Proof of theorem 7.1.a. The average credit per unit of time to scientist i^* under either strategy is equal to the total credit scientist i^* expects to get from the Game divided by the expected duration of the Game (this holds because scientists are modeled as Poisson processes, cf. the proof of theorem A.5). By the proof just given scientist i^* 's expected total credit is higher if she plays strategy E at the present decision node than if she plays strategy H . The expected duration of the Game is at least as high when she chooses strategy H as when she chooses strategy E . So scientist i^* 's average credit per unit of time must also be higher when she chooses strategy E at the present decision node than when she chooses strategy H . \square

References

Arzberger, Peter, Peter Schroeder, Anne Beaulieu, Geof Bowker, Kathleen Casey, Leif Laaksonen, David Moorman, Paul Uhlir, and Paul Wouters. 2004. "Promoting

- Access to Public Research Data for Scientific, Economic, and Social Development.” *Data Science Journal* 3: 135–152. <http://dx.doi.org/10.2481/dsj.3.135>.
- Banerjee, Siddhartha, Ashish Goel, and Anilesh Kollagunta Krishnaswamy. 2014. “Re-incentivizing Discovery: Mechanisms for Partial-progress Sharing in Research.” In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*. New York: ACM, 149–166. <http://dx.doi.org/10.1145/2600057.2602888>.
- Borgman, Christine L. 2012. “The Conundrum of Sharing Research Data.” *Journal of the American Society for Information Science and Technology* 63 (6): 1059–1078. <http://dx.doi.org/10.1002/asi.22634>.
- Boyer, Thomas. 2014. “Is a bird in the hand worth two in the bush? Or, whether scientists should publish intermediate results.” *Synthese* 191 (1): 17–35. <http://dx.doi.org/10.1007/s11229-012-0242-4>.
- 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): 667–688. <http://www.jstor.org/stable/10.1086/682940>.
- Dasgupta, Partha, and Paul A. David. 1994. “Toward a new economics of science.” *Research Policy* 23 (5): 487–521. <http://www.sciencedirect.com/science/article/pii/0048733394010021>.
- Huber, John C. 1998a. “Invention and Inventivity as a Special Kind of Creativity, with Implications for General Creativity.” *The Journal of Creative Behavior* 32 (1): 58–72. <http://dx.doi.org/10.1002/j.2162-6057.1998.tb00806.x>.
- . 1998b. “Invention and Inventivity Is a Random, Poisson Process: A Potential Guide to Analysis of General Creativity.” *Creativity Research Journal* 11 (3): 231–241. http://dx.doi.org/10.1207/s15326934crj1103_3.
- . 2001. “A new method for analyzing scientific productivity.” *Journal of the American Society for Information Science and Technology* 52 (13): 1089–1099. <http://dx.doi.org/10.1002/asi.1173>.
- Huber, John C., and Roland Wagner-Döbler. 2001a. “Scientific production: A statistical analysis of authors in mathematical logic.” *Scientometrics* 50 (2): 323–337. <http://dx.doi.org/10.1023/A%3A1010581925357>.

- . 2001b. “Scientific production: A statistical analysis of authors in physics, 1800–1900.” *Scientometrics* 50 (3): 437–453. <http://dx.doi.org/10.1023/A%3A1010558714879>.
- Huttegger, Simon M. 2013. “Probe and Adjust.” *Biological Theory* 8 (2): 195–200. <http://dx.doi.org/10.1007/s13752-013-0114-2>.
- Huttegger, Simon M., Brian Skyrms, and Kevin J. S. Zollman. 2014. “Probe and Adjust in Information Transfer Games.” *Erkenntnis* 79 (4): 835–853. <http://dx.doi.org/10.1007/s10670-013-9467-y>.
- Norris, James R. 1998. *Markov Chains*. Cambridge: Cambridge University Press. <http://dx.doi.org/10.1017/CB09780511810633>.
- Resnik, David B. 2006. “Openness versus Secrecy in Scientific Research.” *Episteme* 2 (3): 135–147. <http://dx.doi.org/10.3366/epi.2005.2.3.135>.
- Soranno, Patricia A., Kendra S. Cheruvilil, Kevin C. Elliott, and Georgina M. Montgomery. 2015. “It’s Good to Share: Why Environmental Scientists’ Ethics Are Out of Date.” *BioScience* 65 (1): 69–73. <http://bioscience.oxfordjournals.org/content/65/1/69.abstract>.
- Strevens, Michael. 2017. “Scientific Sharing: Communism and the Social Contract.” In *Scientific Collaboration and Collective Knowledge*, edited by Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg. Oxford: Oxford University Press, chapter 1. <https://philpapers.org/rec/STRSSC-2>.