



Cowles Foundation

for Research in Economics
at Yale University

Cowles Foundation Discussion Paper No. 1695

COLLABORATING

Alessandro Bonatti and Johannes Hörner

April 2009

An author index to the working papers in the
Cowles Foundation Discussion Paper Series is located at:
<http://cowles.econ.yale.edu/P/au/index.htm>

This paper can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1396261>

Collaborating*

Alessandro Bonatti and Johannes Hörner

April 26, 2009

Abstract

This paper examines moral hazard in teams over time. Agents are collectively engaged in an uncertain project, and their individual efforts are unobserved. Free-riding leads not only to a reduction in effort, but also to procrastination. The collaboration dwindles over time, but never ceases as long as the project has not succeeded. In fact, the delay until the project succeeds, if it ever does, increases with the number of agents. We show why deadlines, but not necessarily better monitoring, help to mitigate moral hazard.

1 Introduction

Cooperation evolves over time. A lack of tangible results often breeds mistrust, and mistrust leads to lower levels of commitment. Agents grow suspicious that other team members are not pulling their weight in the common enterprise and scale back their own involvement in response. Is this bleak scenario the fate of every team project? What can be done to avert such a scenario? This paper develops a formal framework to address these questions.

In modern times, collaborations have proven crucial in the production of knowledge at the

*We would like to thank Dirk Bergemann, Martin Cripps, Glenn Ellison, Meg Meyer, David Miller, Motty Perry, Sven Rady, Larry Samuelson, Roland Strausz and Xavier Vives for useful discussions, and Nicolas Klein for excellent comments. We thank Glenn Ellison for providing us with the data that was used in Ellison (2002).

level of individual researchers and institutions.¹ Teamwork is a special case of private provision of public goods: when it is hard to quantify the contribution of each member on the advancement of a project, free-riding arises. What sets teamwork apart from most existing models of provision of public goods, though, is the uncertainty regarding the intrinsic feasibility of the project. As of today, and despite the best efforts of minds as brilliant as Newton's, the legendary "philosopher's stone" that would convert base metals into gold has not been found. Applications of the private provision of public goods under uncertainty can be found not only in the direct production of knowledge, but also in its financing. Consider, for instance, a public good of unknown value, such as voluntary contributions (donations) to medical research. Other applications include donations to charity, and investment in research and development (R&D) by firms engaged in a joint venture.

With very few exceptions, most of the work on public goods deals with situations that are either static or involve complete information. This work has provided invaluable insights into the underprovision of the public good in quantity terms. In contrast, given the scope for learning, our interest lies with the dynamics of this provision.

The key features of our model are the following: (i) *Benefits are public, costs are private*: the profit, or value, from completing the project is common to all agents. All it takes to complete the project is one breakthrough, but making a breakthrough requires costly effort. (ii) *Success is uncertain*: some projects are doomed to failure, no matter how much effort is put into them. As for the other projects, the flow probability of a breakthrough increases as the combined effort of the agents increases. Achieving a breakthrough is the only way to ascertain the project's type. (iii) *Effort is hidden*: the choice of effort exerted by an agent is unobserved by the other agents. As long as there is no breakthrough, agents receive no hard evidence whatsoever; they simply become (weakly) more pessimistic about the prospects of the project as time goes on. This captures the idea that output is observable, but effort is not. We shall contrast our findings with the case in which effort is observable. At the end of the paper, we also discuss the intermediate

¹Although Dasgupta (1988) credits Klein and Lie with the first collaborative research in academia resulting in joint publication, for their article on group theory in 1870, collaborative papers can be traced back to 1665 and the joint paper by Hooke, Oldenburg, Cassini and Boyle. This is not to say that collaborations did not take place throughout scientific history, as the well-known collaborations of Kepler and Brahe, Lavoisier and Laplace, or Gauss and Weber illustrate.

case in which a project involves several observable steps.

Our main findings are the following:

- *Agents procrastinate:* as is to be expected, agents slack off, i.e., there is underprovision of effort overall. Agents do not only exert too little effort, but they also do so too late. In the hope that the effort of others will suffice, they work less than they should early on, postponing effort to later dates. Nevertheless, due to growing pessimism, the effort expended dwindles over time, but the plug on the project is never pulled. Although the overall effort expended is independent of the size of the team, the more agents are involved in the project, the later the project gets completed on average, if ever.
- *Deadlines are beneficial:* if agents have enough resolve to fix themselves a deadline, it is optimal to do so. This is the case despite the fact that agents pace themselves so that, if the deadline is hit, the project is abandoned at a point at which the project is still worthwhile. If agents could re-negotiate at this time, they would. But the deadline gives agents incentives to exert effort once it looms close enough. The deadline is desirable, because the reduction in wasteful delay more than offsets the value that is forfeited if the deadline is reached. In this sense, the delay is more costly than the underprovision of effort.
- *Better monitoring need not reduce delay:* when effort is observed, there are multiple equilibria. Depending on the equilibrium, delay might be greater or smaller than under non-observability. In the unique symmetric Markovian equilibrium, delay is actually greater. This is because individual efforts are strategic substitutes. The prospects of the team improve if it is found that an agent has slacked off, because this mitigates the growing pessimism in the team. Therefore, an observable reduction in current effort encourages later effort by other members, and this depresses equilibrium effort. Hence, better monitoring need not alleviate moral hazard. Nevertheless, there are also non-Markovian, “grim-trigger” equilibria for which delay is smaller.

Some insights into the relevance of these findings can be gleaned from the data collected by Ellison (2002). For papers that get eventually published, Ellison (2002) shows that the number of coauthors has a positive effect on the time lag between the submission and acceptance of the

paper. This agrees with our finding that increasing the team size increases the expected delay. However, a closer look at the data used by Ellison (2002) reveals that the time lag between submission and acceptance is decreasing in the number of coauthors when these coauthors are located in the same department, which presumably facilitates monitoring. This is evidence suggesting that the Markovian equilibrium outcome in the observable case does not describe actual behavior.

The value of deadline raises the issue of mechanism design. We derive the optimal dynamic (budget-balanced) compensation scheme, as well as the optimal wage scheme for a principal who owns the project's returns.

We then investigate the role of synergies. As Alchian and Demsetz (1972) have already pointed out, it is likely that an important factor in the formation of teams is the team members' potential to work synergistically. That is, total output might not be separable in the agents' efforts. To account for this phenomenon, we examine two extensions of the baseline model. In the first extension, workers may have similar skills, but their efforts get combined in a non-separable way. For concreteness, we consider the case in which the arrival rate of success in case the project is good is a constant elasticity function of their individual efforts. In the second extension, we focus on two agents and consider the case in which agents possess different skills, so that one worker might be able to succeed where another could not.

These two kinds of synergy lead to different kinds of behavior. In the first case, there are multiple equilibria that are all inefficient. In all equilibria, all agents put in some effort. These equilibria illustrate the distinction between the amount and the allocation of effort: the symmetric equilibrium is characterized by free-riding and significant delay, but it also is the one in which the probability of a breakthrough is greatest. In the second case, the agent that is viewed as most likely to succeed starts by exerting effort by himself, at a level that is socially efficient. As time passes by and no breakthrough occurs, the difference in the likelihoods of success across agents levels off. When both agents are equally likely to succeed, they both start to exert effort, but these effort levels are low because of free-riding. Unless a breakthrough occurs, both agents keep on exerting effort forever, albeit at rapidly declining levels.

Finally, we consider the case in which completing a project involves several tasks. Tasks

are independent. We are particularly interested in understanding how the type of the tasks affects the structure and efficiency of equilibria. Following the literature on social psychology, we distinguish between (i) additive tasks, in which payoffs are additively separable in the tasks, (ii) conjunctive tasks, in which both tasks must be completed, and (iii) disjunctive tasks, in which the project is completed as soon as there is a breakthrough in one task. Efficiency requires tasks to be worked on simultaneously when they are additive, but sequentially if they are conjunctive. However, there are equilibria in which agents specialize and work simultaneously on different tasks when they are conjunctive.

This paper is related to several strands of literature. First, our model can be viewed as a model of experimentation. There is a growing literature in economics on experimentation in teams. For instance, Bolton and Harris (1999) and Keller, Rady and Cripps (2005) study a two-armed bandit problem in which different agents may choose different arms. While free-riding plays an important role in these papers as well, effort is always observable. Rosenberg, Solan and Vieille (2007), Hopenhayn and Squintani (2008) and Murto and Välimäki (2008) consider the case in which while the outcome of each agent's action is unobservable, their actions are observable. This is precisely the opposite of what is assumed in this paper; here, actions are not observed, but outcomes are. Bergemann and Hege (2005) study a principal-agent relationship with an information structure similar to the one considered here. All these models provide valuable insights into how much total experimentation is socially desirable, and how much can be expected in equilibrium. As will be clear, these questions admit trivial answers in our model, which is therefore not well-suited to address those types of question.

Second, our model ties into the literature on free-riding in groups, starting with Olson (1965) and Alchian and Demsetz (1972), and further studied in Holmström (1982), Legros and Matthews (1993), and Winter (2004). In a sequential setting, Strausz (1999) describes an optimal sharing rule. More precisely, ours is a dynamic version of moral hazard in teams with uncertain output. The static version was introduced by Williams and Radner (1988) and also studied by Ma, Moore and Turnbull (1988).

Third, our paper is related to the literature on dynamic contributions to public goods. Games with observable contributions are examined in Admati and Perry (1991), Compte and Jehiel

(2004), Fershtman and Nitzan (1991), Lockwood and Thomas (2002), and Marx and Matthews (2000). Fershtman and Nitzan (1991) compare open- and closed-loop equilibria in a set-up with complete information and find that observability exacerbates free-riding. In Bag and Roy (2008), Bliss and Nalebuff (1984), and Gradstein (1992), agents have independently drawn and privately known values for the public good. This type of private information is briefly discussed in the conclusion. Applications to partnerships include Levin and Tadelis (2005), and Hamilton, Nickerson and Owan (2003). Also related is the literature in management on alliances, including, for instance, Doz (1996), Gulati (1995) and Gulati and Singh (1998).

There is a vast literature on free-riding, also known as social loafing, in social psychology. See, for instance, Latané, Williams and Harkins (1979), or Karau and Williams (1993). Levi (2007) provides a survey of group dynamics and team theory. The stage theory, developed by Tuckman and Jensen (1977) and the theory by McGrath (1991) are two of the better known theories regarding the development of project teams—the patterning of change and continuity in team structure and behavior over time.

2 The Set-up

There are n agents engaged in a common project. The project has a probability $\bar{p} < 1$ of being a good project, and this is commonly known by the agents. It is a bad project otherwise.

Agents continuously choose whether to exert effort or not over the infinite horizon \mathbb{R}_+ . Effort is costly, and the instantaneous cost to agent $i = 1, \dots, n$ of exerting effort $u_i \in \mathbb{R}_+$ is $c_i(u_i)$, for some function $c_i(\cdot)$ that is differentiable and strictly increasing. In most of the paper, we assume that $c_i(u_i) = c_i \cdot u_i$, for some constant $c_i > 0$, and that the choice is restricted to the unit interval, i.e. $u_i \in [0, 1]$. The effort choice is, and remains, unobserved.

Effort is necessary for a breakthrough to occur. More precisely, a breakthrough occurs with instantaneous probability equal to $f(u_1, \dots, u_n)$, if the project is good, and to zero if the project is bad. That is, if agents were to exert a constant effort u_i over some interval of time, then the delay until they found out that the project is successful would be distributed exponentially over that time interval with parameter $f(u_1, \dots, u_n)$. The function f is differentiable and strictly

increasing in each of its arguments. In the baseline model, we assume that f is additively separable and linear in effort choices, so that $f(u_1, \dots, u_n) = \sum_{i=1, \dots, n} \lambda_i u_i$, for some $\lambda_i > 0$, $i = 1, \dots, n$.

The game ends if a breakthrough occurs. Let $\tau \in \mathbb{R}_+ \cup \{+\infty\}$ denote the random time at which the breakthrough occurs ($\tau = +\infty$ if it never does). We interpret such a breakthrough as the successful completion of the project. A successful project is worth a net present value of 1 to each of the agents.² As long as no breakthrough occurs, agents reap no benefits from the project. Agents are impatient, and discount future benefits and costs at a common discount rate r .

If agents exert effort (u_1, \dots, u_n) , and a breakthrough arrives at time $t < \infty$, the average discounted payoff to agent i is thus

$$r \left(e^{-rt} - \int_0^t e^{-rs} c_i(u_{i,s}) ds \right),$$

while if a breakthrough never arrives ($t = \infty$), his payoff is simply $-r \int_0^\infty e^{-rs} c_i(u_{i,s}) ds$. The agent's objective is to choose his effort so as to maximize his expected payoff.

To be more precise, a (pure) strategy for agent i is a measurable function $u_i : \mathbb{R}_+ \rightarrow [0, 1]$, with the interpretation that $u_{i,t}$ is the instantaneous effort exerted by agent i at time t , conditional on no breakthrough having occurred. Given a strategy profile $u := (u_1, \dots, u_n)$, it follows from Bayes' rule that the belief held in common by the agents that the project is good (hereafter, the common belief), p , is given by the solution to the familiar differential equation

$$\dot{p}_t = -p_t(1 - p_t)f(u_t),$$

with $p_0 = \bar{p}$.³ Given that the probability that the project is good at time t is p_t , and that the

²We keep the value of the project fixed throughout (independently of n), in line with our interpretation of success as a public good. Note that the finding that delay increases in n will only be exacerbated if the value of success to an agent is decreasing in the size of the team.

³To see this, note that, given p_t , the belief at time $t + dt$ is

$$p_{t+dt} = \frac{p_t e^{-f(u_t)dt}}{1 - p_t + p_t e^{-f(u_t)dt}},$$

instantaneous probability of a breakthrough conditional on this event is $f(u_t)$, the instantaneous probability assigned by the agent to a breakthrough occurring is $p_t f(u_t)$. It follows that the expected instantaneous reward to agent i at time t is given by $p_t f(u_t) - c_i(u_{i,t})$. Given that the probability that a breakthrough has not occurred by time t is given by $\exp\{-\int_0^t p_s f(u_s) ds\}$, it follows that the average (expected) payoff that agent i seeks to maximize is given by

$$r \int_0^\infty (p_t f(u_t) - c_i(u_{i,t})) e^{-\int_0^t (p_s f(u_s) + r) ds} dt.$$

Given that there is a positive probability that the game lasts forever, and that agent i 's information set at any time t is trivial, strategies that are part of a Nash equilibrium are also sequentially rational on the equilibrium path; hence, our objective is to identify the Nash equilibria of this game. (We shall nevertheless briefly describe off-the-equilibrium-path behavior as well.)

3 The Benchmark Model

We begin the analysis with the special case in which agents are symmetric, and both the instantaneous probability and the cost functions are linear in effort:

$$f(u_1, \dots, u_n) = \sum_{i=1}^n \lambda_i u_i, \quad c_i(u_i) = c_i u_i, \quad u_i \in [0, 1], \quad \lambda_i = \lambda, \quad c_i = c, \quad \text{all } i.$$

Equivalently, we may define the normalized cost $\alpha := c/\lambda$, and redefine u_i , so that each agent chooses the control variable $u_i : \mathbb{R}_+ \rightarrow [0, \lambda]$ so as to maximize

$$V_i(\bar{p}) := r \int_0^\infty \left(p_t \sum_i u_{i,t} - \alpha u_{i,t} \right) e^{-\int_0^t (p_s \sum_i u_{i,s} + r) ds} dt, \quad (1)$$

subject to

$$\dot{p}_t = -p_t(1 - p_t) \sum_i u_{i,t}, \quad p_0 = \bar{p}.$$

by Bayes' rule. Subtracting p_t on both sides, dividing by dt and taking limits gives the result.

Observe that the parameter α is the Marshallian threshold: it is equal to the belief at which a myopic agent would stop experimenting, because at this point the instantaneous marginal revenue from effort, p_t , is the same as the marginal cost, α .

From each agent's point of view, this is a standard optimal control problem. Before addressing it, we describe the cooperative solution.

3.1 The Team Problem

If agents behaved cooperatively, they would choose efforts so as to maximize the sum of their individual payoffs, that is,

$$W(\bar{p}) := \sum_{i=1}^n V_i(\bar{p}) = r \int_0^\infty (np_t - \alpha) u_t e^{-\int_0^t (p_s u_s + r) ds} dt,$$

where, with some abuse of notation, $u_t := \sum_i u_{i,t} \in [0, n\lambda]$. The integrand being positive as long as $p_t \geq \alpha/n$, it is clear that it is optimal to set u_t equal to $n\lambda$ as long as $p_t \geq \alpha/n$, and to zero otherwise. The belief p_t is then given by

$$p_t = \frac{\bar{p}}{\bar{p} + (1 - \bar{p})e^{n\lambda t}},$$

as long as the right-hand side exceeds α/n . In short, the team solution specifies that each agent sets his effort as follows:

$$u_{i,t} = \lambda \text{ if } t \leq T_n := (n\lambda)^{-1} \ln \frac{\bar{p}(1 - \alpha/n)}{(1 - \bar{p})\alpha/n}, \text{ and } u_{i,t} = 0 \text{ for } t > T_n.$$

Not surprisingly, the resulting payoff is decreasing in the discount rate r and the normalized cost α , and increasing in the prior \bar{p} , the upper bound λ and the number of agents, n .

Observe that the instantaneous marginal benefit from effort to an agent is equal to p_t , which decreases over time, while the marginal cost is constant and equal to α . Therefore, it will not be possible to provide incentives for selfish agents to exert effort beyond the Marshallian threshold. The wedge between this threshold and the efficient one, α/n , captures the well-known free-riding

effect in teams, which is described eloquently by Alchian and Demsetz (1972), and has since been studied extensively. In a non-cooperative equilibrium, the amount of experimentation is too low.⁴ Here instead, our focus is on how free-riding affects when effort is exerted.

3.2 The Non-Cooperative Solution

As mentioned above, once the common belief drops below the Marshallian threshold, agents do not provide any effort. Therefore, if $\bar{p} \leq \alpha$, there is a unique equilibrium, in which no agent ever works, and we might as well assume throughout that $\bar{p} > \alpha$. Further, we shall assume throughout this section that agents are sufficiently patient. More precisely, the discount rate satisfies

$$\frac{\lambda}{r} \geq \alpha^{-1} - \bar{p}^{-1} > 0.$$

At the end of this section, we briefly discuss what happens when this assumption is violated.

We begin our analysis by restricting attention to symmetric equilibria. The proof of the main result of this section relies on Pontryagin’s principle, but the gist of it is perhaps best understood by the following heuristic argument from dynamic programming.⁵

What is the trade-off between exerting (some) more effort at some instant and exerting it at the next? Consider some time t and assume that players have followed the equilibrium strategies up to t . Given some small time interval $dt > 0$, we consider the marginal gain from effort over the time interval $[t, t + dt]$ vs. the marginal gain over the time interval $[t + dt, t + 2dt]$. Write u_i, p for $u_{i,t}, p_t$, and u'_i, p' for $u_{i,t+dt}, p_{t+dt}$, and let $V_{i,t}$, or V_i , denote the unnormalized continuation payoff of agent i at time t . Given that

$$V_{i,t} = (p(u_i + u_{-i}) - \alpha u_i)dt + (1 - rdt)(1 - p(u_i + u_{-i})dt)V_{i,t+dt},$$

⁴This means that there is neither an “encouragement effect” in our set-up, unlike in some papers on experimentation (see, for instance, Bolton and Harris (1999)), nor any effect of patience on the experimentation threshold. This is because a breakthrough yields a unique lump sum to all agents simultaneously, rather than conditionally independent sequences of lump sum payoffs.

⁵Applying dynamic programming more generally is possible. However, it is difficult in this setting, because deviations are unobserved. This implies that an agent’s best-reply after a deviation depends both on the commonly held belief about the project, but also on his privately held belief; hence, the value function becomes a function of both beliefs.

it follows, by applying the same expansion to $V_{i,t+dt}$, that

$$V_{i,t} = (p(u_i + u_{-i}) - \alpha u_i)dt + (1 - rdt)(1 - p(u_i + u_{-i})dt) ((p'(u'_i + u'_{-i}) - \alpha u'_i)dt + (1 - rdt)(1 - p'(u'_i + u'_{-i})dt)V_{i,t+2dt}).$$

The first-order effect of effort today and effort tomorrow is approximately the same, as

$$\lim_{dt \rightarrow 0} \frac{dV_i/du_i}{dt} = p - \alpha - pV_i = \lim_{dt \rightarrow 0} \frac{dV_i/du'_i}{dt}.$$

That is, increasing effort has two first-order effects. It increases the instantaneous chance of a breakthrough, which is worth $p - \alpha$, but it also makes it more unlikely that the continuation payoff V_i will be collected. While this provides a necessary condition for the optimal effort to be an interior solution, it does not determine the level of effort. To do so, we must consider the next term in the Taylor expansion, and use the fact that $p' = p - p(1 - p)(u_i + u_{-i})dt$.⁶ Then, up to the third order,

$$\frac{dV_i}{du_i} = (p(1 - V_i) - \alpha)dt + (p^2(u_i + u_{-i})V_i + 2rpV_i - p(u'_i + u'_{-i})(1 - V_i) + \alpha pu'_i)(dt)^2.$$

The second-order terms have intuitive interpretations. The first term corresponds to the fact that the first-order loss of the continuation value V_i in the case of a breakthrough due to the marginal increase in effort must be corrected by the probability that two breakthroughs arrive simultaneously. Next, by exerting more effort, the probability of a breakthrough goes up, which allows a saving to be made on the depreciation of the continuation value due to discounting (since we consider an interval of length $2dt$, this cost saving is proportional to $2r$). On the other hand, the probability of collecting the capital gain $1 - V_i$ tomorrow goes down, but there is a commensurate saving due to the lower probability of having to exert effort tomorrow. Similarly, up to the third order,

$$\frac{dV_i}{du'_i} = (p(1 - V_i) - \alpha)dt + (p'^2(u'_i + u'_{-i})V_i + 2rpV_i - p(u_i + u_{-i})(1 - V_i) + \alpha p(u_i + u_{-i}) - r(p - \alpha))(dt)^2,$$

⁶Here, recall also that $e^{-p(u_i + u_{-i})dt} = 1 - p(u_i + u_{-i})dt + \frac{p^2}{2}(u_i + u_{-i})^2(dt)^2 + o(dt^3)$.

where the first three second-order terms have analogous interpretations. The penultimate term corresponds to the saving in marginal cost from exerting effort later: due to the fact that there is a probability $p(u_i + u_{-i})dt$ that success occurs in the next instant, the cost of effort is spared in this event. The last term corresponds to the marginal product of effort that is dissipated because of discounting. Comparing these two expressions, and assuming continuity of u_i, u_{-i} , and thus of p , if procrastinating does not strictly dominate exerting effort now, it must be that

$$\alpha p u_i \geq \alpha p (u_i + u_{-i}) - r(p - \alpha).$$

Observe that the right-hand side strictly exceeds the left-hand side if u_i is continuous and u_{-i} is bounded away from zero, for $p \rightarrow \alpha$. Therefore, effort must tend to zero as p tends to α . Similarly, it follows that effort must tend to zero as $r \rightarrow 0$. The underlying economic forces are clear: by working a bit more now, the agent's cost saving tomorrow is proportional to his level of effort. By planning to work a bit more tomorrow instead, if that should prove to be necessary, the saving is proportional to the combined current levels of effort, which determine whether it will indeed be necessary to produce that additional effort tomorrow. As a result, agents procrastinate.

Second, assume for the sake of contradiction that agents stop working at some point. Then, considering the penultimate instant, it must be that, up to the second order, $p - \alpha = p(1 - p)(u_i + u_{-i})dt$, and so we may divide both sides of the previous inequality by $u_i + u_{-i} = n u_i$, yielding

$$p(1 - p)r dt \geq \frac{n - 1}{n} \alpha p,$$

which is impossible, as dt is arbitrarily small. Therefore, not only does effort go to zero as p tends to α , but it does so sufficiently fast that the belief never reaches the threshold α , and agents keep on working on the project forever, albeit at negligible rates.

From this derivation, it is easy to guess what the (symmetric) equilibrium value of u_i must be. Namely, given that agent i must be indifferent between exerting effort at different moments

in time, it must be that

$$\alpha p u_i = \alpha p(u_i + u_{-i}) - r(p - \alpha), \text{ or } u_i(p) = \frac{r(\alpha^{-1} - p^{-1})}{n - 1}.$$

Hence, the common belief tends to the Marshallian threshold asymptotically, and total effort, as a function of the belief, is actually decreasing in the number of agents. To understand this last result, observe that the equilibrium reflects the logic of mixed strategies. Because efforts are perfect substitutes, the indifference condition of each agent requires that the total effort by all agents but him be a constant that depends on the belief, but not on the number of agents. This implies that each agent's level of effort must be decreasing in the total number of agents. In turn, this implies that the total effort by all agents for a given belief is the sum of a constant function and of a decreasing function of the number of agents. Therefore, it must be decreasing in the total number of agents.

This simple logic relies on two substitutability assumptions: efforts of different agents are perfect substitutes, and the cost function is linear. Later, we shall discuss relaxations of those two assumptions.

To describe the equilibrium, we must specify each agent's level of effort as a function of his information, not just as a function of his belief, which is endogenous. Since effort choices are unobservable, effort is a function of time (and of one's own past effort). The following theorem, which we prove in the appendix, describes the optimal strategy on the equilibrium path.⁷

Theorem 1 *There exists a unique symmetric equilibrium, in which, on the equilibrium path, the level of effort of any agent is given by, for all $t \geq 0$,*

$$u_{i,t}^* = \frac{r}{n - 1} \frac{\alpha^{-1} - 1}{1 + \frac{(1-\bar{p})\alpha}{\bar{p}-\alpha} e^{\frac{n}{n-1}r(\alpha^{-1}-1)t}}. \quad (2)$$

From this formula, it is immediate to derive the following comparative statics.

⁷In the case $\bar{p} = 1$ that was ruled out earlier, the game reduces essentially to the static game. The effort level is constant and, because of free-riding, inefficiently low ($u_{i,t} = \frac{r(\alpha^{-1}-1)}{n-1}$).

Lemma 1

1. Individual instantaneous effort is decreasing in t . For a fixed time t , individual instantaneous effort is increasing in r and \bar{p} , and decreasing in α .
2. As a function of t , individual and total instantaneous effort are first decreasing, and then increasing in n .
3. The agent's average discounted payoff $V_i(\bar{p})$ is increasing in n and \bar{p} , decreasing in α , and independent of r .
4. Aggregate effort (over time) and the probability that a breakthrough will occur eventually are independent of n . Conditional on a breakthrough occurring eventually, its expected time increases with the number of agents.

As discussed above, the total instantaneous effort is decreasing in n for a given belief p , so that, in terms of time t , total effort is also decreasing in n when t is small enough. Yet this means that the belief decreases more slowly when there are more agents; because effort is increasing in the belief, it must be that eventually, the total instantaneous effort is higher in larger teams. Another way to see this is that, since the belief approaches the threshold α asymptotically, independently of n , this implies that the aggregate effort that is exerted over the infinite horizon is independent of n as well. Ultimately, then, larger teams must catch up in terms of effort.




Figure 1: Individual and collective effort

The last conclusion of the lemma is an immediate consequence of this temporal pattern. As mentioned above, the payoff turns out to be independent of the discount rate. For instance, the

expected cost of delay until the random time τ of a breakthrough, conditional on a breakthrough occurring eventually, defined as $1 - \mathbb{E}[e^{-r\tau} | \tau < \infty]$, is independent of the discount rate. While larger teams are slower to succeed, their members are better off, because the payoff is increasing in n . This result is obvious for $n = 1$ relative to $n = 2$, because an agent may always choose to behave as if he were by himself, securing the payoff from a single-agent team. The result is less obvious for larger team sizes. Figures 1 and 2 provide a quantitative illustration of these results.




Figure 2: Payoffs and cost of delay

If an agent deviated, what would his optimal continuation strategy be? Suppose that this deviation is such that, at time t , the aggregate effort of agent i over the interval $[0, t]$ is lower than that which it would have been on the equilibrium path. This means that agent i is more optimistic than the other agents, and his private belief exceeds the common belief of the other agents. Given that agent i would be indifferent between exerting effort or not if he shared the common belief, his optimism leads him to exert maximal effort until the time at which his private belief catches up with the other agents' common belief, at which point he will revert to the common, symmetric strategy. If instead his aggregate effort up to t is greater than what it is supposed to be, then agent i is more pessimistic than the other agents, and he will provide no effort whatsoever until the common belief catches up with his private belief, if it ever does. This

completes the description of the equilibrium strategy.

While the symmetric equilibrium is unique, there exist other, asymmetric, equilibria. Consider, for instance, the case of two agents only ($n = 2$). If an agent were by himself, he would behave as he would in the cooperative set-up. That is, he would exert maximal effort up to time T_1 , at which his belief reaches the level α . Faced with such behavior, the best reply of another (sufficiently patient) agent would be to exert no effort whatsoever; indeed, the value of effort identified in the symmetric equilibrium is precisely the threshold such that, if an agent expected his partner to put in more effort than this, he would find it optimal to put in none himself. So there is an asymmetric equilibrium in which one agent behaves as if he were by himself, and in which the other agent puts no effort whatsoever. It is not difficult, then, to see that there is actually an entire continuum of equilibria, of which we have identified the two extreme points. Each equilibrium is indexed by an agent $i = 1, 2$, and some time $t_1 \leq T_1$, such that, up to time t_1 , agent i chooses maximum effort, while agent $-i$ exerts no effort at all, and from time t_1 onward, given the resulting belief at time t_1 , the two agents behave as in the symmetric equilibrium.

In the appendix, we prove that, as long as, as assumed, $\lambda/r \geq \alpha^{-1} - \bar{p}^{-1}$, this is precisely the set of all equilibria: with n agents, an equilibrium is indexed by a collection of nested subsets of agents, $\{i\} \subset \{i, j\} \subset \{i, j, k\} \subset \dots \subset \{1, \dots, n\}$, at (not necessarily distinct) times $t_1 \leq t_2 \leq \dots \leq t_n$, with $t_1 \leq T_1$, $t_k \in \mathbb{R}_+ \cup \{\infty\}$ for $k \geq 2$ and $t_n = \infty$, such that agent i exerts maximal effort up to t_1 , agents i, j exert effort as in the symmetric equilibrium (i.e., $u_i = u_j = r(\alpha^{-1} - p^{-1})$ given the resulting p) over the interval $(t_1, t_2]$, etc.

Clearly, if agents are sufficiently patient (so that $u_{i,t} < \lambda$ for all t), the overall payoff of the team is maximized by the asymmetric equilibrium in which one agent works by himself. So, according to the utilitarian rule, this asymmetric equilibrium is the best equilibrium and the symmetric equilibrium is the worst. However, according to the maximin rule, the ranking is reversed, because the agent who works alone is worse off than in the symmetric equilibrium.

The existence of such an asymmetric equilibrium relies on the strong substitutability conditions that have been assumed so far. As we shall see, if the agents' efforts display complementarities, such an extreme outcome can no longer occur in equilibrium. Nevertheless, the trade-off between efficiency and fairness will persist.

The assumption regarding the discount rate was necessary to ensure that the agent's individual effort characterized in Theorem 1 was less than the maximum effort level λ . If this assumption is not satisfied, it is possible that both agents exert maximal effort simultaneously, at least initially. That is, the unique symmetric equilibrium then has the feature that all agents exert maximal effort up to time t at which, given the resulting belief, the level of effort $u_i(p)$, as defined above the theorem, is equal to λ (since $u_i(p) \rightarrow 0$ as $p \rightarrow \alpha$, this always occurs at some time $t < T_n$). From that point on, agents exert effort at level $u_i(p)$, and the qualitative features of the equilibrium are as before.

So far, agents have been assumed to be identical. If the normalized cost α is the same across agents, but not necessarily the capacity λ_i , there is little change in the analysis. In particular, the symmetric equilibrium remains an equilibrium provided that $\lambda_i/r \geq \alpha^{-1} - \bar{p}^{-1}$ for all agents i . Similarly, if agents have different discount rates, there is an equilibrium in which all agents exert interior effort levels, so that the more patient agent i is, the smaller u_{-i} is.

However, the outcome changes dramatically if the normalized cost differs across agents. Indeed, if $\alpha_i < \min_{j \neq i} \alpha_j$ (and $\lambda_i/r \geq \alpha^{-1} - \bar{p}^{-1}$ for all i), then, in the unique equilibrium of the game, agent i behaves as if he were on his own, by exerting maximal effort up to the point at which the belief p reaches α_i . The proof for the case $n = 2$ is in the appendix. The intuition is straightforward. Given that agent i has incentives to exert effort for any belief $p > \alpha_i$, but no other agent has an incentive to exert any effort as soon as $p < \min_{j \neq i} \alpha_j$, agent i must be the last agent to exert effort.⁸ At that point, he might as well exert maximum effort. However, consider the last agent other than i , say j , to exert any effort, and let t be the time at which he is supposed to stop exerting effort. At time $t - dt$, agent j has a strict preference for procrastinating. After all, that last bit of effort could always be exerted later, and given that agent i will be starting to exert maximum effort in an instant, the probability that he might be able to avoid exerting this effort altogether is high enough for him to defer. This implies that $t = 0$; hence, agent i is always the only one to exert effort. In particular, if entering the collaboration involved any type of additional cost for each agent, agent i would never join the team.

⁸More precisely, it cannot be that $\lim_{t \rightarrow \infty} p_t > \alpha_i$, because, when the combined efforts of all agents $j \neq i$ become negligible, agent i has a strict incentive to exert maximum effort. Hence, $\lim_t p_t = \alpha_i$, so that, if t is the smallest solution to $p_t = \min_{j \neq i} \alpha_j$, agent i must be the only agent exerting effort after t .

Again, this extreme outcome is partly driven by the perfect substitutability in the productivity of the agents' efforts. As we shall see, when efforts are imperfect substitutes, both agents work even when their productivities differ. Nevertheless, it is suggestive that teams involving agents with skills of a similar kind, but dissimilar levels, are unlikely to be successful.

3.3 A Comparison with the Observable Case

We now contrast the previous findings with the corresponding results for the case in which effort is perfectly observable. That is, we assume here that all agents' efforts are observable, and that agent i 's choice as to how much effort to exert at time t (hereafter, his effort choice) may depend on the entire history of effort choices up to time t . Such a continuous-time game involves well-known nontrivial modeling choices. A standard way to sidestep these choices is to focus on Markovian strategies. Here, the obvious state variable is the belief p . Unlike in the unobservable case, this belief is always commonly held among agents, even after histories off the equilibrium path.

A strategy for agent i , then, is a map $u_i : [0, 1] \rightarrow [0, \lambda]$ from possible beliefs p into an effort choice $u_i(p)$, such that (i) u_i is left-continuous; and (ii) there is a finite partition of $[0, 1]$ into intervals of strictly positive length on each of which u_i is Lipschitz-continuous. By standard results, a profile of Markov strategies $u(\cdot)$ uniquely defines a law of motion for the agents' common belief p , from which the (expected) payoff given any initial belief p can be computed (cf. Presman (1990) or Presman and Sonin (1990)). A Markov equilibrium is a profile of Markov strategies such that, for each agent i , and each belief p , the function u_i maximizes i 's payoff given initial belief p . See, for instance, Keller, Rady and Cripps (2005) for details. Following standard steps, agent i 's continuation payoff given p , $V_i(p)$, must satisfy the optimality equation given by, for all p , and $dt > 0$, to the second order,

$$\begin{aligned} V_i(p) &= \max_{u_i} \{((u_i + u_{-i}) p_t - u_i \alpha) dt + (1 - (r + (u_i + u_{-i}) p_t) dt) V_i(p_{t+dt})\} \\ &= \max_{u_i} \{((u_i + u_{-i}) p_t - u_i \alpha) dt + (1 - (r + (u_i + u_{-i}) p_t) dt) (V_i(p) - (u_i + u_{-i}) p (1 - p) V_i'(p) dt)\}. \end{aligned}$$

Taking limits as $dt \rightarrow 0$ yields

$$0 = \max_{u_i} \{ (u_i + u_{-i})p - u_i\alpha - (r + (u_i + u_{-i})p)V_i(p) - (u_i + u_{-i})p(1-p)V'_i(p) \},$$

assuming, as will be verified, that V is differentiable. We focus here on a symmetric equilibrium in which the effort choice is interior. Given that the maximand is linear in u_i , its coefficient must be zero. That is, dropping the agent's subscript,

$$p - \alpha - pV(p) - p(1-p)V'(p) = 0,$$

and since $V(\alpha) = 0$, the value function is given by

$$V(p) = p - \alpha + \alpha(1-p) \ln \frac{(1-p)\alpha}{(1-\alpha)p}.$$

Plugging back into the optimality equation, and solving for $u := u_i$, all i , we get

$$u(p) = \frac{r}{\alpha(n-1)}V(p) = \frac{r}{\alpha(n-1)} \left(p - \alpha + \alpha(1-p) \ln \frac{(1-p)\alpha}{(1-\alpha)p} \right).$$

It is standard to verify that the resulting u is the unique equilibrium strategy profile provided that \bar{p} is such that $u \leq \lambda$ for all $p < \bar{p}$. In particular, this is satisfied when, as assumed in the unobservable case, $\lambda/r \geq \alpha^{-1} - \bar{p}^{-1}$, which we maintain henceforth. In the model without observability, recall that, in terms of the belief p , the effort is given by $u(p) = \frac{r}{n-1}(\alpha^{-1} - p^{-1})$. As is clear from these formulas, the eventual belief is the same whether effort is observed or not, and so aggregate effort over time is the same in both models. However, delay is not.

Theorem 2 *In the symmetric Markov equilibrium with observable effort, the equilibrium level of effort is strictly lower, for all beliefs, than that in the unobservable case.*

Proof: Subtracting the equilibrium value of effort in the observable case from the value in the unobservable case gives

$$\frac{r}{\alpha(n-1)} \frac{1-p}{p} \left(p - \alpha - \alpha p \ln \frac{(1-p)\alpha}{(1-\alpha)p} \right).$$

This term is positive, as can be seen as follows. Let $f(p)$ be the term in brackets. This function is convex, as $f''(p) = \alpha/(p(1-p)^2)$, and its derivative at $p = \alpha$ is equal to $(1-\alpha)^{-1} > 0$. Hence, f is increasing in p over the range $[\alpha, 1]$, and it is equal to 0 at $p = \alpha$, so it is positive over this range. \square

Thus, fixing a belief, the instantaneous equilibrium level of effort is lower when previous choices are observable, and so is the welfare. This means that delay is greater under observability. While this may be a little surprising, it is an immediate consequence of the fact that effort choices are strategic substitutes. Because effort is increasing in the common belief, and because a reduction in one agent's effort choice leads to a lower rate of decrease in the common belief, such a reduction leads to a greater level of effort by other agents. That is, to some extent, the other agents take up the slack. This depresses the incentives to exert effort and leads to lower equilibrium levels. This cannot happen when effort is unobservable, because an agent cannot induce the other agents into exerting the effort for him. Figure 3 illustrates this relationship. As can be seen from the right panel, a lower level of effort for every value of the belief p does not imply a lower level of effort for every time t : given that the total effort over the infinite horizon is the same in both models, levels of effort are eventually higher in the observable case.

The individual payoff is independent of the number of agents $n \geq 2$ in the team in the observable case. This is a familiar rent-dissipation result: when the size of the team increases, agents waste in additional delay what they save in individual cost. This can be seen directly from the formula for the level of effort, in that the total effort of all agents but i is independent of n . It is worth pointing out that this is *not* true in the unobservable case. This is one example in which the formula that gives the effort as a function of the common belief is misleading in the unobservable case: given p , the total instantaneous effort of all agents but i is independent of n here as well. Yet the value of p is not a function of the player's information only: it is the common belief about the unobserved past total efforts, including i 's effort; hence, it depends on the number of agents. As we have seen, welfare is actually increasing in the number of agents in the unobservable case. The comparison is illustrated in the left panel of Figure 3.

There also other, asymmetric Markovian equilibria in the observable case similar to those described in Keller, Rady and Cripps (2005). These equilibria may involve many points at

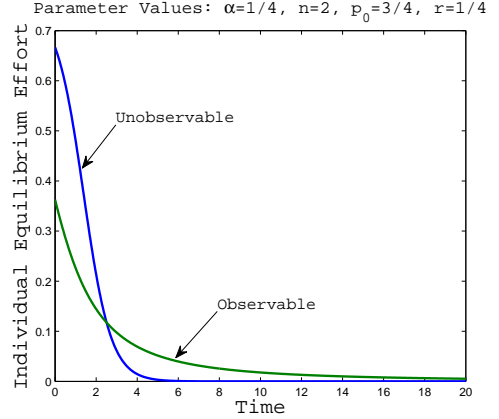
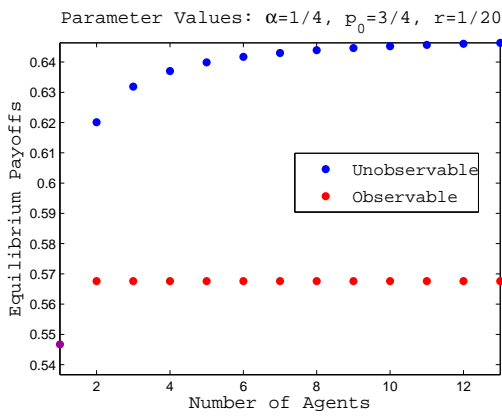


Figure 3: Welfare and effort in the observable vs. non-observable case

which the identity of the agent exerting effort switches. Such switching is impossible without observability. Suppose that agent i is expected to exert effort alone over the interval $[t - \varepsilon, t]$, for some $\varepsilon > 0$, while another agent j exerts effort alone during some time interval starting at time t . Under observability, if agent i procrastinates, this “freezes” the common belief and therefore postpones the time of switching until agent i makes up for the wasted time. So, the punishment for procrastination is automatic. Here instead, agent i ’s deviation is not observable; hence, agent j will start exerting effort at time t , as expected. Further, given that an agent working alone, such as j , must be exerting maximal effort (because of discounting), procrastination is profitable for i , because the induced probability that the postponed effort will not be exerted more than offsets the loss in value due to discounting. Therefore, such switching is impossible without observability, independently of the agents’ discount rate.

In the observable case, there exist other, non-Markovian symmetric equilibria. As mentioned above, appropriate concepts of equilibrium have been defined carefully elsewhere (see, for instance, Bergin and McLeod (1993)). It is not difficult to see how one can define formally a “grim-trigger” equilibrium, for low enough discount rates, in which all agents exert effort at a maximal rate until time T_1 at which $p = \alpha$, and if there is a unilateral deviation by agent i , all other agents stop exerting effort, leaving agent i with no choice but to exert effort at a maximal rate from this point on until the common belief reaches α . While this equilibrium is not first-best, it clearly does better than the Markovian equilibrium in the observable case, and than the

symmetric equilibrium in the unobservable case.⁹

Which equilibrium is more likely to emerge? As a simple application, we reconsider the data that was used by Ellison (2002). This data, as mentioned above, corroborates one of our main findings: as the number of coauthors increases, the time lag between submission and acceptance of the paper increases. As a rudimentary measure of observability, we consider, for each paper, the number of coauthors affiliated with the same department. The underlying assumption is that it is easier to monitor effort when coauthors are physically close; hence, we view papers written by several authors from the same department as being projects with observable levels of effort, and papers written by distant coauthors as projects with unobservable levels of effort. As Table 1 illustrates, the difference between these two kinds of papers is striking: an additional coauthor increases the time lag by a month on average; however, it actually increases it by one month and a half if the coauthors are geographically distant, while it reduces the time lag by two weeks when they are neighbors. Both results are statistically significant. On one hand, this reinforces the main conclusion of the unobservable case. On the other hand, this is suggestive evidence that coauthors in the same department do not follow the prescription of the Markovian equilibrium. The grim-trigger equilibrium, for instance, is consistent with the data. Delay is decreasing in the number of coauthors, because coauthors exert effort at a maximal rate. Therefore, increasing the number of authors increases the total amount of effort exerted at any instant.

It would be interesting to examine whether the quality of monitoring affects the amount of effort that is exerted, in addition to the delay. The acceptance rate could be used as a proxy variable. Unfortunately, the data is not suited for such an analysis, because it is restricted to papers that were accepted eventually.

4 Deadlines and Other Mechanisms

In the absence of any kind of commitment, the equilibrium outcome described above seems inevitable. Pleas and appeals to cooperate are given no heed and deadlines are disregarded. In this section, we examine what can be done to improve upon this outcome. Three mechanisms

⁹It is tempting to consider grim-trigger strategy profiles in which agents exert effort for beliefs below α . We ignore them here, because such strategy profiles cannot be limits of equilibria of discretized versions of the game.

Table 1: Submit-Accept Time Regressions

	(1)	(2)
	Submit-accept time	Submit-accept time
Total number of authors	28.970** [14.232]	46.430*** [15.511]
Number of coauthors in the same department		-61.156** [26.032]
Journal dummies	Yes	Yes
Journal trends	Yes	Yes
Field dummies	Yes	Yes
Observations	1393	1393

Note: The dependent variable is the length of time between submission of a paper to a journal and its acceptance in days. The sample is a subset of the set of papers published in the top five or six general-interest economics journals between 1990 and 1998. The total number of authors ranges from 1 to 3. All regressions include journal dummies, journal-specific linear time trends, dummies for 17 fields of economics, and all the control variables used by Ellison (2002) in Table 6. Robust standard errors are reported in brackets (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

are considered.

The first mechanism is a self-imposed time-limit. We assume that agents can commit to a deadline and that they choose the optimal deadline such that, if no breakthrough has occurred by this time, all agents stop exerting effort and the project is abandoned. Effectively, this is equivalent to considering the game with a finite horizon. We examine which is the optimal horizon.

A deadline is an extreme version of a symmetric reward scheme, in which the entire surplus is shared evenly if a breakthrough obtains before the deadline expires, but is completely destroyed afterwards. In the second mechanism, agents can commit to any symmetric reward scheme (as a function of time) they wish to, but the scheme has to be *ex ante* budget-balanced. Clearly, this must at least weakly improve on a deadline, and we show that, in fact, a deadline is not the optimal mechanism satisfying *ex ante* budget balance; it is the optimal mechanism satisfying *ex post* budget balance.

Finally, we contrast this “team” reward scheme with the one that a principal who reaps the benefits of the project would wish to implement. Here as well, because the identity of the actual agent responsible for the breakthrough is not observed, attention is restricted to symmetric schemes in which the principal promises the same reward to all agents, but this reward is allowed to depend on time.

In this section, we normalize the capacity λ to 1.

4.1 Deadlines

For some possibly infinite deadline $T \in \mathbb{R}_+ \cup \{\infty\}$, and some strategy profile $(u_1, \dots, u_n) : [0, T] \rightarrow [0, 1]^n$, agent i 's (expected) payoff over the horizon $[0, T]$ is now defined as

$$r \int_0^T (p_t(u_{i,t} + u_{-i,t}) - \alpha u_{i,t}) e^{-\int_0^t (p_s(u_{i,s} + u_{-i,s}) + r) ds} dt.$$

That is, if time T arrives and no breakthrough has occurred, the continuation payoff of the agents is nil. In the previous section, the case in which $T = \infty$ was studied exhaustively. The next lemma, which we prove in the appendix, describes the symmetric equilibrium for $T < \infty$.

Throughout this subsection, we maintain the restriction on the discount rate r that we imposed in Section 3.2.

Lemma 2 *Given $T < \infty$, there exists a unique symmetric equilibrium, characterized by $T_1 \in [0, T)$, in which the level of effort is given by*

$$u_{i,t} = u_{i,t}^* \text{ for } t < T_1, \text{ and } u_{i,t} = 1 \text{ for } t \in [T_1, T],$$

where u_i^* is as in Theorem 1. The time T_1 is non-decreasing in the parameter T and strictly increasing for T large enough. Moreover, the belief at time T strictly exceeds α .

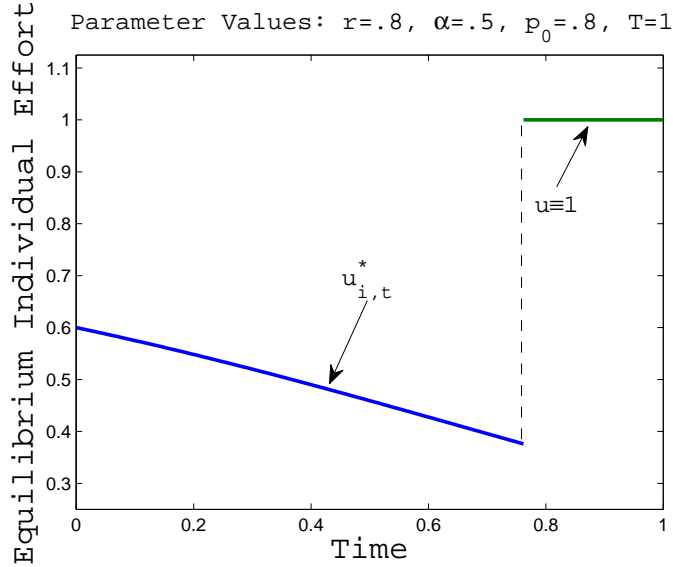


Figure 4: Optimal strategy given a deadline of $T = 1$ ($n = 2$)

See Figure 4. According to Lemma 2, effort is first decreasing over time, and over this time interval, it is equal to its value when the deadline is infinite. At that point, the deadline is far enough in the future not to affect the agents' incentives. However, at some point, the deadline looms large above the agents. Agents recognize that the deadline is near and exert maximal effort from then on. However, it is then too late to catch up with the overall effort exerted in the infinite-horizon case, and $p_T > \alpha$. By waiting until time T_1 , agents take a chance. It is not

difficult to see that the eventual belief p_T must exceed α : if the deadline were not binding, each agent would prefer to procrastinate at instant T_1 , given that all agents then exert maximal effort until the end.

The next theorem establishes that it is in the agents' best interest to fix such a deadline. That is, agents gain from restricting the set of strategies that they can choose from. Furthermore, the deadline is set precisely in order that agents will have strong incentives throughout.

Theorem 3 *The optimal deadline T is finite and is given by*

$$T = \frac{1}{n+r} \ln \frac{(n-\alpha)\bar{p}}{\alpha(n-\bar{p}) - r(\bar{p}-\alpha)}.$$

It is the longest time for which it is optimal for all agents to exert effort at a maximal rate throughout.

Note that the deadline is decreasing in n , because it is the product of two positive and decreasing functions of n . That is, tighter deadlines need to be set when teams are larger. This is a consequence of the stronger incentives to shirk in larger teams. Furthermore nT decreases in n as well. That is, the total amount of experimentation is lower in larger teams. However, it is easy to verify that the agent's payoff is increasing in the team size. Larger teams are bad in terms of overall efficiency, but good in terms of individual payoffs.

In the appendix, it is also shown that T is increasing in r : the more impatient the agents, the longer the optimal deadline. This should not come as a surprise, because it is easier to induce agents who have a greater level of impatience to work longer.

One might suspect that the extreme features of the equilibrium effort pattern in presence of a deadline is driven by the linearity in the cost function. Indeed, as was reported in the last section, the path of equilibrium effort appears to be continuous over time when the cost function is quadratic. Nevertheless, a deadline provides additional incentives to exert effort when time is running short, and the graph of the level of effort is U-shaped in a variety of circumstances.¹⁰

¹⁰Readers have pointed out to us that this effort pattern in the presence of a deadline reminds them of their own behavior as a single author. Sadly, it appears that such behavior is hopelessly suboptimal for $n = 1$. We refer to O'Donoghue and Rabin (1999) for a behavioral model with which they might identify.

4.2 The Optimal Budget-Balanced Mechanism

The setting of a deadline is a rather extreme way in which the team can affect incentives. We now consider a more general class of mechanisms. Agents can commit to a common wage that is an arbitrary function of time. In the case of a breakthrough, each agent collects this wage. The budget is required to be balanced on average, given that the project is worth $v = n$ to the team (because success is a public good that yields a benefit of 1 to each of the n agents, its value v equals n).

Therefore the agents choose $w : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which is measurable, so as to maximize

$$r \int_0^\infty np_t u_t (nw_t - n\alpha) e^{-\int_0^t (np_s u_s + r) ds} dt,$$

subject to there being an *ex ante* budget balance, namely,

$$r \int_0^\infty np_t u_t (v - nw_t) e^{-\int_0^t (np_s u_s + r) ds} dt = 0.$$

The function $u := u_i : \mathbb{R}_+ \rightarrow [0, 1]$, which is measurable, must maximize

$$r \int_0^\infty (p_t (u_{i,t} + u_{-i,t}) w_t - \alpha u_{i,t}) e^{-\int_0^t (p_s (u_{i,t} + u_{-i,t}) + r) ds} dt,$$

where $u_{-i} := (n-1)u$, because agents choose the level of effort that they will exert noncooperatively. We find that fixing a constant wage is not optimal, and neither is fixing a deadline, in the sense described above (i.e. fixing a constant wage up to some time-limit, after which the wage is set to zero). To state the next theorem, it is necessary to introduce some notation. Let

$$\delta := \frac{\bar{p}}{2n(1-\bar{p})} \left(\sqrt{(n-1)^2 + 4n \frac{1-\bar{p}}{\bar{p}} (v/\alpha - 1)} - (n-1) \right),$$

and

$$w(\delta) := \frac{\alpha}{pr(n-r)} \left(p(r + n(n-1)) - r^2 + (nr(1-p)\delta - (n-1)p(n-r)) \delta^{-r/n} \right).$$

Finally, let

$$w_t^* = w(\delta)e^{rt} + \alpha \frac{n-1+r}{r}(1-e^{rt}) - \alpha \frac{r(1-\bar{p})}{(n-r)\bar{p}}(e^{nt} - e^{rt}).$$

Theorem 4 *The optimal wage scheme w is given by*

$$w_t = w_t^* \text{ for } t \leq \hat{T}, \text{ and } w_t = 0 \text{ for } t > \hat{T},$$

for some \hat{T} , such that $u_t = 1$ until time \hat{T} , and $u_t = 0$ thereafter.¹¹

The final time \hat{T} cannot be solved in closed-form, and is the unique strictly positive solution to

$$\frac{r}{n+r} \left(\frac{v}{\alpha} - 1 \right) (e^{(n+r)\hat{T}} - 1) - (n-1)e^{n\hat{T}}(e^{r\hat{T}} - 1) - \frac{nr}{n-r} \frac{1-\bar{p}}{\bar{p}} e^{n\hat{T}}(e^{n\hat{T}} - e^{r\hat{T}}) = 0.$$

This theorem is proved in the appendix, where it is further shown that, at least for small enough discount rates, the belief at the final time \hat{T} is above the asymptotic belief in the baseline model. Thus, if agents are patient enough, the optimal mechanism helps to alleviate the free-rider problem, but does not bring the total effort back to its efficient level. As is the case with a deadline, it is optimal to have agents exert maximal effort until effort stops completely. For low discount rates, the wage is decreasing over time: frontloading payments allows to incentivize agents, because the shrinking wage counteracts the incentive to procrastinate. Frontloading cannot be achieved with a deadline; hence, maximal effort is sustained here over a longer horizon, as $\hat{T} > T$.

This analysis supports the relevance of staggered prizes in the design of scientific competitions. Such degressive rewards are implemented, among others, by the X Prize foundation (in, for example, the design of the Google Lunar X Prize).

4.3 The Principal Agent Problem

Up to this point, attention has been on what the team could do to help itself. Now, we consider the problem from the perspective of a principal who designs the payment that the agent or agents receive (the analysis that follows holds for all $n \geq 1$). As mentioned, we restrict

¹¹Clearly, the choice of w_t for $t > T^*$ is to a large extent arbitrary.

attention to symmetric schemes. Given that it is clear that the principal cannot benefit from paying wages to unsuccessful agents, such a scheme can be summarized by a wage schedule w_t , with the interpretation that each agent receives w_t if the project gets completed at time t (as will be clear, agents do not gain by delaying the announcement of a breakthrough). The principal can commit to any wage path he would like to.

The project is worth v to the principal, who chooses $w : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which is measurable, so as to maximize

$$r \int_0^\infty np_t u_t (v - nw_t) e^{-\int_0^t (np_s u_s + r) ds} dt.$$

The function $u := u_i : \mathbb{R}_+ \rightarrow [0, 1]$, which is measurable, maximizes

$$r \int_0^\infty (p_t(u_{i,t} + u_{-i,t})w_t - \alpha u_{i,t}) e^{-\int_0^t (p_s(u_{i,t} + u_{-i,t}) + r) ds} dt.$$

Indeed, observe that a breakthrough arrives at rate $u_i + u_{-i} = nu$, and that the total wage bill is nw_t . Given that an agent can only be provided with a sufficient incentive to exert effort if the marginal benefit, pw , covers the marginal cost, α , and given that the principal's mark-up in case of a breakthrough is $v - nw$, there is no scope for a profitable scheme if $v \leq n\alpha/\bar{p}$. Conversely, whenever $v > n\alpha/\bar{p}$, the principal may always design some wage scheme that is profitable, because fixing a constant wage that is equal to the average of v/n and α/\bar{p} is feasible.

Theorem 5 *The optimal wage scheme w is given by*

$$w_t = w_t^* \text{ for } t \leq T^*, \text{ and } w_t = 0 \text{ for } t > T^*,$$

which is such that $u_t = 1$ until time $T^ := \frac{1}{n} \ln \delta < \hat{T}$, and $u_t = 0$ thereafter.*¹²

Observe that the principal's wage scheme follows the same dynamics as the wage scheme that is optimal for the team. For both problems, it is optimal to induce maximal effort as cheaply as possible, as long as any effort is worthwhile, but effort stops being worthwhile earlier for the monopolist.

¹²Here as well, the choice of w_t for $t > T^*$ is to a large extent arbitrary.

In the principal's problem, the wage dynamics must offset the free-riding incentives. The principal must concede some rents to the agents to induce them to work early. This is easier when agents are impatient (the wage decreases in r), and more difficult when they are more agents. In fact, the principal's payoff is decreasing in the number of agents: a larger team means more effort early, but the cost in terms of additional rents more than offsets this benefit. In fact, the maximal amount of total effort produced ($nT^* = \ln \delta$) is decreasing in the number of agents: the larger the team, the less the principal is willing to experiment. Again, this can be easily understood in terms of the familiar trade-off faced by a monopolist, who must trade off lower output against lower rents. Observe that this amount of effort is independent of the discount rate. While discounting affects the rents, it also affects the cost of providing this rent, and because the principal and the agents have the same discount rate, these effects cancel out.

As the discount rate r tends to zero, the wage approaches the affine function $\bar{w}(\delta) - \alpha t(n-1)$, where $\bar{w}(\delta) := \alpha \left(1 + \frac{1-\bar{p}}{\bar{p}} \delta + \frac{n-1}{n} \ln \delta \right)$. To provide incentives when agents are perfectly patient, the wage must decay at the rate of the marginal cost (or be constant when $n = 1$).

Figure 5 illustrates how the wage scheme varies with the parameters. Effort stops when $p_t w_t = \alpha$, or $w_t = p_t / \alpha$. Note that this time is independent of the discount rate, as is already clear from the formula for δ . This means that the total amount of effort is independent of the discount rate, and so is efficiency. On the other hand, the amount of effort varies with the size of the group. As is easy to verify, the value of δ is decreasing in n : the larger the group, the higher the aggregate effort. However, the principal's payoff need not increase in n . Smaller groups are typically better, because it is less costly to overcome the incentive to free-ride.

One might wonder whether total effort is lower under an output-restricting, but sophisticated principal, or under a mechanism as imperfect as a self-imposed deadline. It is shown in the appendix that, when $v = n$, so that the comparison is meaningful, and the discount rate is sufficiently low, $T^P \leq T$: a principal stops the project before the optimal deadline. Clearly, agents are then worse off.




Figure 5: Optimal wages

5 Synergies

Social psychology stresses the role of synergies as an important factor in team success, and Alchian and Demsetz (1972) emphasized their importance for moral hazard in teams. Among the many kinds of possible synergies, we focus here on two extreme versions. In the first case, agents are more effective when working together than working when separately. That is, the rate at which a success arrives if the project is good displays complementarities in their effort choices. In the second case, in which attention is restricted to two agents, agents have different skills, so that, depending on the type of project, one skill, the other, both or neither might be appropriate.

5.1 Complementarities

In this subsection, we assume that the instantaneous probability of a breakthrough, conditional on the project being good, is given by

$$f(u_1, \dots, u_n) = \left(\sum_i u_i^\rho \right)^{1/\rho}, \text{ where } \rho \in (0, 1),$$

and, as before, capacity λ is normalized to 1. The instantaneous rate of arrival of a success in case the project is good has the property of constant elasticity of substitution with respect to the agents' levels of effort.¹³ The assumption that ρ is positive guarantees that $\lim_{u_{-i} \rightarrow 0} f_i(u_i, u_{-i}) > 0$ for all $u_i > 0$, where $f_i = \partial f / \partial u_i$.¹⁴ The baseline model corresponds to the special case $\rho = 1$. Agents are assumed to be identical for now ($\alpha_i = \alpha$). We assume that the discount rate satisfies

$$(n-1)r^{-1} \geq \alpha^{-1} - \frac{n^{1-\frac{1}{\rho}}}{\bar{p}} > 0.$$

Due to the fact that $\lim_{u_{-i} \rightarrow 0} f_i(u_i, u_{-i}) = \infty$ for all $u_{-i} > 0$, it is no longer possible that, in equilibrium, some agents exert no effort whatsoever while some other agents exert effort. No matter what the other agents do, the returns from some sufficiently small amount of effort are arbitrarily large. As we shall see, this does not imply that the equilibrium is necessarily symmetric.

This specification captures the notion that by working together, agents are more productive than by themselves. Indeed, observe that, in the team problem in which agents work cooperatively, it is optimal to set $u_i = 1$ for all i , so that $f(u_i, \dots, u_n) = n^{\frac{1}{\rho}}$, up to the time at which the common belief drops to the level $n^{-\frac{1}{\rho}}\alpha$. That is, there is strictly more experimentation here than in i 's single-decision problem (i.e., when $u_{-i} = 0$). In this case, the instantaneous probability of success is equal to u_i , as in the baseline case, and effort is only exerted at a maximal rate up to the time at which the common belief reaches the threshold α .

As in our baseline case, it is more convenient to represent the equilibrium level of effort in terms of the common equilibrium belief at that time. However, recall that effort is not observable, and that effort is a function of time only, while the common belief is a function of time and effort, which can be derived from the equilibrium strategies.

¹³Of course, there are many other technologies with the property that levels of effort are imperfect substitutes (for example $f(u) = \delta \min \{u_i\} + (1-\delta) \max_{j \neq i} \{u_j\}$). The C.E.S. function provides the clearest comparison with the baseline model, as shown in Theorem 6.

¹⁴Otherwise, there always exist trivial equilibria in which no agent exerts any effort.

Theorem 6 *There exists a unique symmetric equilibrium, in which the level of effort exerted at time t is given by*

$$u_i(p_t) = u(p_t) := \frac{r}{n-1} \left(\alpha^{-1} - \frac{n^{1-\frac{1}{\rho}}}{p_t} \right),$$

given the equilibrium value of p_t . Effort is positive and strictly decreasing, tending to 0 as $t \rightarrow \infty$.

This result generalizes Theorem 1 in the natural way. With synergies as well, free-riding leads to delay and effort dwindles over time. The belief p_t converges to $n^{1-\frac{1}{\rho}}\alpha$, which corresponds, here as well, to the effort exerted in the team problem when the prize is divided by n . Observe that this belief is no longer equal to the threshold in the single-player decision problem. As mentioned above, in the single-player decision problem, effort is exerted up to the point at which the belief is equal to α . Given that $\alpha > n^{1-\frac{1}{\rho}}\alpha$, aggregate equilibrium effort is higher with synergies and tends to the efficient level as the parameter ρ tends to zero. The threshold at which effort asymptotically stops, $n^{1-\frac{1}{\rho}}\alpha$, is decreasing in n , unlike in the case without synergies in which it was constant. Provided that the belief is low enough, effort is greater with more agents, although not compared to the first-best level.

It is easy to derive the corresponding symmetric Markovian equilibrium in the observable case. The equilibrium effort is given by

$$u_i(p) = \frac{r}{(n-1)\alpha} \left(p - \frac{\alpha}{k_n} - \frac{\alpha}{k_n}(1-p) \ln \frac{p}{1-p} \frac{1 - \alpha/k_n}{\alpha/k_n} \right),$$

where $k_n := n^{\frac{1}{\rho}-1}$, which gives rise to the same asymptotic threshold as in the non-observable case. It is easy to verify, as in Theorem 2, that the effort exerted is less than in the non-observable case, for any given degree of belief.

To discuss the asymmetric equilibria, it is simpler to assume that $n = 2$, an assumption that is maintained throughout this subsection. As mentioned above, synergies ensure that there cannot be an equilibrium in which some agent does not exert any effort at all, while his partner does. However, they might exert different levels of effort. To describe the asymmetric equilibria, define the function

$$g(\sigma) := \frac{(\sigma^\rho + 1)^{2-\frac{1}{\rho}}}{\sigma^{2\rho-1} + 1}.$$

It is easy to verify that this function is strictly increasing for $\rho > 1/2$, so that its inverse function, g^{-1} , is well-defined and increasing as well. Without loss of generality, assume that $u_{1,0} \geq u_{2,0}$, i.e. agent 1 exerts at least as much effort as agent 2 at the initial instant.

Theorem 7 *For $\rho < 1/2$, there exists no asymmetric equilibrium. For $\rho \geq 1/2$, there exists a continuum of asymmetric equilibria. Each asymmetric equilibrium is uniquely identified by the value of $u_{1,0}/u_{2,0}$, which is in $(1, g^{-1}(\bar{p}/\alpha)]$ if $\rho > 1/2$, and is unrestricted for $\rho = 1/2$.*

The symmetric equilibrium corresponds to the special case in which $u_{1,0}/u_{2,0} = 1$. In the proof in the appendix, we further show that the aggregate effort is strictly lower in any asymmetric equilibrium than in the symmetric equilibrium for $\rho > 1/2$, and equal to it for $\rho = 1/2$. Agents stop exerting effort at the same finite time if $\rho = 1/2$, but never stop if $\rho > 1/2$. The roles of agents are never reversed: because agent 1 exerts more effort than agent 2 at the initial time, he keeps on exerting more effort throughout. Figure 6 displays the range of values of initial ratios $\sigma_0 := u_{1,0}/u_{2,0}$ for which an asymmetric equilibrium exists, for a given prior $p_0 = \bar{p}$ and a level of complementarity ρ .




Figure 6: Admissible initial values for asymmetric equilibria

As in the baseline model, there is a trade-off between efficiency and fairness: it is always best for the team's aggregate payoff if the team members' choices about how much effort to exert are, to a certain extent, asymmetrical when synergies are not too strong. However, as one would expect, simulations show that the optimal degree of asymmetry decreases as the strength of the synergies increases.

What if agents have different costs? Suppose that $\alpha_1 < \alpha_2$, so that agent 1 is more efficient. There exists a continuum of equilibria, indexed by the initial ratio of effort levels, $\sigma_0 = u_{1,0}/u_{2,0}$, which uniquely defines the equilibrium path. (Depending on \bar{p} , there might be a bound on the admissible values of σ_0 . The appendix contains a more formal discussion.) The equilibrium trajectories of effort can be decomposed into two time intervals. Initially, the more efficient agent works more than his counterpart. Afterwards, the opposite is true. Depending on the initial value of σ_0 , one or the other interval might be empty. This ordering of roles can be understood in terms of the agents' relative incentives to procrastinate. The more inefficient agent stands to lose less by procrastinating (the project is worth less to him, and so discounting his value is less costly), and has more to gain (effort being costlier for him, he appreciates more the potential cost saving from procrastination). Therefore, if agents must "take turns," he must be the second agent.

The total amount of experimentation, as measured by the limit value of the belief p , is maximized when the more efficient agent works more throughout; more precisely, he must work sufficiently more for the ratio σ to reach asymptotically $(\alpha_2/\alpha_1)^{1/(1-\rho)}$, as the common belief p tends to its limiting value $\alpha_2((\alpha_2/\alpha_1)^{\rho/(1-\rho)} + 1)^{1-1/\rho}$. This means that, not only are there asymmetric equilibria that tend to the symmetric equilibrium as $\alpha_2 \rightarrow \alpha_1$, but this is in particular the case for the equilibrium that maximizes the amount of experimentation. Moreover, this statement does not depend on the value parameter ρ . Therefore, the discontinuity between equilibria of the baseline model with symmetric and asymmetric players is not robust to the introduction of (an arbitrarily small level of) complementarities. The two panels of Figure 7 illustrate how the respective efforts might vary with the initial value σ_0 .





Figure 7: Equilibria when $\alpha_1 \neq \alpha_2$.

5.2 Different Skills

In this subsection, the baseline model is generalized as follows. Instead of the project's being simply good or bad, it may be any of four possible types. (1) It is good (type 0), and both agents are equally able to achieve a breakthrough, as before. In this case, we maintain the assumption that the arrival rate of a breakthrough has instantaneous probability $u_{1,t} + u_{2,t}$. (2 and 3) It is of type $i = 1, 2$, in which case only agent i 's effort might lead to a breakthrough. That is, the instantaneous probability of a breakthrough is now $u_{i,t}$, independently of $u_{-i,t}$. (4) It is bad (type 3), and efforts are then wasted, because breakthroughs are impossible.

The initial belief is now given by a vector $\bar{p} = (\bar{p}^0, \bar{p}^1, \bar{p}^2)$, where \bar{p}^k is the initial belief that the project is of type k . More generally, we write p_t^k for the equilibrium belief at time t that the project is of type k , so that $p_t^3 = 1 - p_t^0 - p_t^1 - p_t^2$, and we assume that $\bar{p}^3 > 0$.¹⁵ Let $u_0 := u_1 + u_2, u_3 := 0$. Using Bayes' rule, it is readily verified that, for $k = 0, \dots, 3$,

$$\dot{p}^k / p^k = \sum_{j=0}^3 p^j u_j - u_k.$$

¹⁵The case $\bar{p}^3 = 0$ can be studied independently. In particular, when skills are perfectly negatively correlated ($\bar{p}^0 = \bar{p}^3 = 0, \bar{p}^1 = \bar{p}^2 = 1/2$), as in Klein and Rady (2008), belief and effort remain constant, with a level of effort equal to $r(\alpha^{-1} - (1/2)^{-1})$, which is the same effort as in the baseline model for a belief $\bar{p} = 1/2$.

Agent i seeks to maximize

$$V_i(\bar{p}) = r \int_0^\infty \left(\sum_{j=0}^3 p_t^j u_{j,t} - \alpha_i u_{i,t} \right) e^{-\int_0^t (\sum_{j=0}^3 p_s^j u_{j,s} + r) ds} dt.$$

We assume that agents are symmetric, i.e. α_i is independent of i . We focus first on the case in which $\bar{p}^1 = \bar{p}^2$.

Observe that it is no longer an equilibrium for only one agent to exert effort throughout the project. Indeed, if agent $i = 1, 2$ works by himself, p^i will decrease while p^j will increase, where j is the index of the other agent. That is, the absence of a breakthrough leads agent i to become more pessimistic about his chances of making a breakthrough than agent j . Therefore, at the point at which it does not pay for i to continue exerting effort on his own, it would still be profitable for j to exert effort.

It cannot be, either, that one agent works after another agent stops working, because the agent that remains active would exert maximal effort as soon he was working by himself, say at time t . Yet the other agent would be unwilling to exert any effort at time $t - dt$, for small enough $dt > 0$. So when one agent works, both must work, and if they are patient enough, the unique solution for equilibrium effort is symmetric and interior. The restriction on the discount rate and on the parameters must be changed to

$$r^{-1} \geq \alpha^{-1} - \frac{1}{1 - \bar{p}^3} \left(1 + \left(1 + \frac{\bar{p}^0(1 - \bar{p}^3)}{\bar{p}^1 \bar{p}^2} \right)^{-1/2} \right) > 0,$$

which is assumed here. If the second inequality fails, effort is identically zero in every equilibrium.

As stated, the problem is multidimensional. However, it turns out that it can be solved explicitly and that, on the equilibrium path, effort only depends on p^3 . To state the result, we define

$$C := \frac{\bar{p}^0 \bar{p}^3}{\bar{p}^1 \bar{p}^2}, \text{ and } \tilde{p}^3 := \frac{1 - 2C - 2\alpha(1 - C) + \sqrt{1 - 4\alpha(1 - \alpha)(1 - C)}}{2(1 - C)}.$$

Theorem 8 *Assume $\bar{p}^1 = \bar{p}^2 > 0$. There exists a unique equilibrium, which is symmetric. In this equilibrium, at time t , given the equilibrium value of p_t^3 , agents exert a level of effort equal*

to

$$u_i(p_t^3) = u(p_t^3) := \frac{r}{\alpha} \left(1 - \frac{\alpha}{1 - p_t^3} \left(1 + (1 + C(1 - p_t^3)/p_t^3)^{-1/2} \right) \right).$$

Effort is positive for all $t \geq 0$, and $\lim_t p_t^3 = \tilde{p}^3$.

The exact relationship between time t and belief p_t^3 is given in the appendix. Given this result, it would appear that having different skills does not fundamentally change the incentives of agents to free-ride. As in the baseline model, the equilibrium reflects dilatory behavior. The project suffers a protracted delay and effort dwindles over time. Indeed, it follows from the theorem that $\lim_{t \rightarrow \infty} u_{i,t} = 0$. As the common value $\bar{p}^1 = \bar{p}^2$ tends to zero, the effort approaches the symmetric equilibrium from Theorem 1. Therefore, introducing agent-specific skills restores uniqueness and singles out the symmetric equilibrium in the limit.

Somewhat surprisingly, the limiting threshold \tilde{p}^3 only depends on the prior belief \bar{p} via a one-dimensional statistic, C , in which it is increasing. One extreme case is obtained by taking $\bar{p}^1 = \bar{p}^2$ to zero. We are then back to the baseline model, in which the belief $1 - p^3$ tends to α . The other extreme case is obtained by taking \bar{p}^0 to zero. Skills are then entirely independent, yet free-riding persists, because effort is not maximal, and $1 - p^3$ tends to a higher threshold, 2α , which reflects the lower probability that a particular agent's skill is the appropriate one.

If one agent is more likely to solve the problem than the other, in the sense that $\bar{p}^1 > \bar{p}^2$, say, the equilibrium is unique as well. As we show in the appendix, along the equilibrium path, the more "optimistic" agent (agent 1, here), starts by exerting maximal effort by himself, up to the point at which $p^1 = p^2$ (recall that p^1 will decrease and p^2 will increase), provided that the resulting p^3 is still lower than \tilde{p}^3 (otherwise, agent 1 stops at some point, and neither agent works thereafter). From that point on, both agents work symmetrically, as a function of p^3 , as described in Theorem 8.

As we show in the appendix, the belief \tilde{p}^3 has a natural interpretation. In the team problem in which agents behave cooperatively, if a breakthrough is worth $1/2$ to each agent, rather than 1 , the optimal strategy profile calls for both agents to exert maximal effort up to the time t at which $p_t^3 = \tilde{p}^3$. This means that, as in the baseline model, the effect of free-riding can be decomposed into two components: (i) it affects the total exertion of effort in the usual way (\tilde{p}^3 falls short of the cooperative threshold); and (ii) it also affects the timing of this exertion, as

described above. Figure 8 below illustrates the pattern of effort and belief over time.




Figure 8: Effort and beliefs with synergies

6 Multitask Projects

So far, we have presented a project as consisting of a single breakthrough. This is a gross simplification. Most projects involve several steps, or tasks, each of which must be completed for the project to be a success. In this section, we adapt some of our earlier findings to the design of optimal collaborations in the case of multiple tasks. For simplicity, we restrict attention to the case of two tasks. In addition, given that the information structure is no longer trivial, we require strategies to be sequentially rational. At any time, agents now observe whether a task has been already (successfully) completed, which task it was, and at what time it was completed. Neither the level nor the allocation of effort is observed. We also assume that if both agents work on one task simultaneously and the task gets completed at that instant, the specific agent that is responsible for the success cannot be identified.¹⁶

¹⁶This is not to say that the other case is uninteresting or intractable, but for the sake of concision, one or the other modeling choice had to be made.

The literature on project design has emphasized the relevance of the type of task for the effect of social loafing. Particular attention is devoted to the distinction between: conjunctive tasks, all of which must be completed for the project to be a success; disjunctive tasks, only of which must be completed to guarantee the success of the team (and completing any further task provides no further benefit); and additive tasks, for which the value of the project is additive in the completion of the tasks (if only one project is completed, it is worth a payoff normalized to 1 to each agent. If both are completed, each agent obtains a payoff of 2.)

Another dimension along which projects vary relates to the timing of tasks. For some projects, it is possible to work on two tasks simultaneously. For others, however, it is imperative to complete a specific task before tackling the second one.

As before, we assume that each task is either good or bad. The type of task is statistically independent across tasks and the initial probability that any given task is good is still denoted by \bar{p} . Effort is additively separable across tasks. Agents are assumed to be identical, with marginal cost α , and the total capacity for effort normalized to 1. We describe here some equilibria. It is easy to verify that these strategy profiles are indeed equilibria. Proofs are omitted and available upon request.

6.1 The Team Problem

To understand the agents' incentives better, it is useful to start with the case in which there is only one agent (or, equivalently, in which agents act cooperatively). The following result holds for all discount rates $r \geq 0$. In all cases, effort is exerted at the maximal rate until some point at which effort stops altogether. Our focus is on the allocation of effort across tasks, so we omit a discussion of these stopping times. The next proposition describes the optimal sequencing.

Proposition 1 (The Team solution) *If $n = 1$, it is optimal to:*

1. *Work sequentially on tasks if they are conjunctive. If the prior probability \bar{p} of each task being good is identical across tasks, it is better to start with the task for which the parameter α is higher. If the cost parameter α is the same, it is better to start with the task for which \bar{p} is lower.*

2. *Split work equally across tasks if they are disjunctive and identical. If they are not identical, effort must be devoted exclusively to the task for which the difference $p - \alpha$ is larger, until these differences are equalized, after which effort must be split so as to maintain these differences equal.*
3. *Tackle the tasks in any order, if tasks are additive and $r = 0$. If $r > 0$, effort with additive tasks should be allocated as in the disjunctive case.*

The first conclusion might first sound a little surprising. However, recall that all conjunctive tasks must be completed for the project to be successful. If one task is likely to be impossible to complete, either because agents are quite pessimistic about it (low \bar{p}) or because it is demanding (high α), then it makes sense to avoid wasting effort on the “easier” task by postponing tackling it until it has been determined whether or not the more difficult task can be completed.¹⁷

If tasks are disjunctive, on the other hand, it makes sense to devote the effort to whichever task yields the higher immediate return, that is, the task for which the spread $p - \alpha$ is larger. This also holds for the additive case, but only because agents are impatient. Otherwise, since tasks are independent, the order becomes irrelevant.

6.2 Sequential Conjunctive Tasks

We first consider the case in which task 1 needs to be completed in order to start task 2. Agents each receive a payoff of 1 if both tasks are completed, and nothing otherwise.

Observe first that, in the single-agent problem with only one task, which is worth v , the agent would exert maximal effort up to the point at which his belief p would satisfy $pv = \alpha$, or $p = \alpha/v$. That is, increasing the prize has the same effect as decreasing the marginal cost, and if agents had different values for the project, the unique equilibrium would involve the agent with the higher valuation exerting all the effort.

Next, observe that, in the asymmetric equilibrium of the single-task project in which one agent exerts effort at a maximal rate and the other agent does not exert any effort at all, the idle agent has quite obviously a strictly higher payoff.

¹⁷Obviously, the threshold for the first task in the sequential problem is not the threshold α for that task considered on its own.

It follows that, if agent i performs the last task all by himself, his continuation payoff v_i , at the time at which the first task is completed, is strictly lower than the payoff of the other agent, v_j . From the point of view of performing the first task, the second task can be summarized by the continuation payoffs (v_1, v_2) . Therefore, if agent i performs the last task by himself, independently of the time at which the first task is successfully completed, it must be that agent $j \neq i$ is the only one exerting effort on the first task. With two tasks, there is no longer a trade-off between efficiency and fairness and there exists a unique equilibrium such that the last task is performed by one agent only.

This reasoning can be extended to multiple tasks. With two tasks left, the agent who performs the last task has a slightly higher continuation payoff. This is because he will only exert effort if (and after) the other agent is successful. Therefore, he must be the one working by himself on the first of the three tasks. This reasoning can obviously be extended to any number of tasks: there exists a unique equilibrium such that the last task is performed by one agent only; in this equilibrium, agents alternate in executing tasks, as long as they are successful.

6.3 Conjunctive Tasks

Now consider the case in which the two tasks are conjunctive (*i.e.*, the payoff is awarded only upon completion of both). However, there are no restrictions on the timing of players' efforts. We focus on the case in which players are symmetric and sufficiently patient, and discuss the following two equilibria.¹⁸

Proposition 2 (Conjunctive Tasks) *If agents are sufficiently patient, the following are pure-strategy equilibria of the game with two tasks.*

1. *Agents work sequentially, each on one task. Agent 2 begins to work only if, and after, agent 1 has completed the first task successfully .*
2. *Agents work simultaneously, each on one task. Agents exert maximal effort until some time T , at which they stop working. Upon completing a task, an agent stops working, while the other one keeps on exerting maximal effort up to some time.*

¹⁸The specification of the appropriate bound on r is omitted here.

In the equilibrium with sequential efforts, beliefs eventually reach the efficient thresholds. In fact, the first agent works until his beliefs offset the payoff from having the second agent complete the remaining one. After and if the first agent has completed the task, the second agent works until the beliefs reach α , because his value from a success is equal to 1. Furthermore, both agents exert maximal effort, which makes this equilibrium the most efficient noncooperative solution.

The equilibrium with simultaneous work and specialization is supported by the threat that, if an agent is successful after time T , he is required to work alone on the remaining task, up to the appropriate thresholds. As agents work, they become increasingly pessimistic about their partner's chances of completing the other task. This reduces the value of completing their own task, so that the belief threshold at which they would stop actually increases over time. This threshold is reached when α/p equals the expected value from having the other agent work alone on the remaining task, starting from a belief p . However, if a task gets completed, the remaining agent works until the usual threshold $p = \alpha$. Therefore, this equilibrium is less efficient than the sequential one, purely in terms of total effort (because both individual thresholds increase over time). The continuation strategy of stopping work after a success is efficient both from an *ex post* perspective (because having only agent work on one task is efficient), and from an *ex ante* perspective. It is of critical importance that each agent knows that he is alone working on his task. Otherwise, he would be tempted to wait for the other agent to complete his own task. Continuation play prescribes the strongest possible punishment for deviating. In fact, if the first success is obtained after the time at which both agents were supposed to stop, the first agent to succeed must also complete the remaining task. This specification is admittedly extreme, although the equilibrium outcome can also be supported by weaker ones (under stronger restrictions on the parameters).

Finally, there are also several equilibria with simultaneous and non maximal effort levels. For instance, agents use the symmetric (baseline) equilibrium strategy on task 1, given the continuation payoff (which affects the limit threshold) and then, if successful, on task 2.

6.4 Additive Tasks

Now consider the case in which tasks are additive and payoffs are given by the total number of successes.

Proposition 3 (Additive Tasks) *If players are sufficiently patient, the following are pure-strategy equilibria of the game with two tasks.*

1. *Agents work sequentially, each on one task. Agent 2 begins to work only if, and after, agent 1 has completed his task.*
2. *Agents work simultaneously, each on one task. Agents exert effort at the maximal rate until the single-task threshold $p = \alpha$ is reached. Upon obtaining a success, an agent stops working, while the other one completes his task.*

In the equilibrium in which agents work simultaneously, both agents work until they reach the single-task threshold, because the value of each success is independent of their beliefs about other tasks. This also deters an agent from delaying his efforts to until after his partner has completed his task. Unlike in the case of conjunctive tasks, this equilibrium is more efficient than the sequential one. In fact, both equilibria involve both agents working until the single-project threshold, but the sequential one has a longer expected completion time.

There also exists a symmetric equilibrium without maximal effort. Agents can work on the two tasks both sequentially or simultaneously. In either case, agents adopt the equilibrium strategies described in our baseline model with one task. The equilibrium with sequential efforts requires a minimal level of patience to ensure that agents actually want to wait for one task to be completed (or abandoned) before starting to work on the other.

6.5 Disjunctive Tasks

Now consider the case in which projects are disjunctive, which means that success in any one project ends the game with a unit payoff for both agents.

Proposition 4 (Disjunctive Tasks) *If agents are sufficiently patient, the following are equilibria of the game with two tasks.*

1. *Agents work simultaneously, each on one task. Each agent exerts a lower amount of effort than in the case of a common single-task project.*
2. *Agents work simultaneously and divide their efforts equally across tasks.*

As in the team problem, agents maintain the spread $p - \alpha$ constant across tasks. In the equilibrium with agents working each on one task, the incentives to procrastinate are stronger than in our baseline case. By shirking today, an agent “freezes” his beliefs about his task. Exerting effort tomorrow will therefore be relatively more productive. Analogously, in the equilibrium in which efforts are divided across tasks, the total amount of effort that is exerted on each task is lower than in the equilibrium with division of labor. Indeed, suppose that the total amounts of effort exerted on each task were equal to the case of division of labor. Holding fixed the effort devoted to one task, each agent would be indifferent between working and not working at all on the other task, provided his partner does not collaborate on it. Since his partner is now exerting positive effort on both tasks, each player then has an incentive to free-ride and reduce his effort, relative to the equilibrium with division of labor.

7 Concluding Remarks

We have shown that moral hazard distorts not only the amount of effort, but also its timing. Agents work too little, too late. Downsizing the team might help, provided that agents’ skills are not too complementary. On the other hand, increasing transparency might aggravate the delay. Setting an appropriate deadline is beneficial, in as much as the reduction in delay more than offsets the further reduction in effort.

The model that we have considered is quite stylized, partly for reasons of simplicity, partly for tractability. We discuss here how relaxing two of the assumptions affects the main results.

Convex Costs: Throughout the analysis, we have maintained the assumption that the cost is linear in the level of effort. While this affords tractability, it is natural to ask whether the

findings are robust to this assumption. This is especially relevant given that, with linear cost, agents are actually indifferent between any level of effort at a symmetric equilibrium, so that such an equilibrium has the flavor of a mixed-strategy equilibrium, for which comparative statics are sometimes counterintuitive.

While it is no longer possible to obtain closed-form formulas for the solution of the Euler-Lagrange equations that characterize the interior solution in the case of nonlinear cost, we present here a few numerical illustrations for the case of cost functions that are power functions, i.e. $c(u_i) = c \cdot u_i^\gamma$, $\gamma > 1, c > 0$. We focus here on the case of symmetric agents, and the instantaneous probability of a breakthrough is still given by the sum of the efforts. That is, the model is otherwise identical to the baseline case.




Figure 9: Cost of delay and payoffs with convex cost

The first remark is that convex costs are similar to synergies, in the sense that, with more agents, it is possible to achieve the same total level of effort at a lower cost (because dividing the same total effort across more agents lowers the overall cost, when the cost is convex). This should favor larger teams, and we might expect that the amount of effort exerted by an agent does not decrease as quickly, as we increase the number of agents, relative to the baseline model. In turn, this might lead to the time within which a breakthrough is expected to occur being reduced (conditional on a breakthrough occurring in finite time) for larger teams, while the

impact on payoff is ambiguous.¹⁹ Indeed, this is precisely what we find, provided the convexity is sufficiently pronounced. See Figure 9 for an illustration of welfare and the cost of delay, and see the left panel of Figure 10 below, which shows that the effort path becomes flatter as the convexity becomes stronger (the cost functions have been normalized so that the value of the single-decision problem remains constant).




Figure 10: Effort with convex cost with or without a deadline

The other finding that seems to depend significantly on the linear cost structure is the discontinuity in equilibrium levels of effort when there is a deadline. Indeed, with convex costs, one suspects that the equilibrium effort should be a continuous function of time. This is indeed the case, as shown in the right panel of Figure 10. Furthermore, for a deadline that is far enough in the future, the graph for effort is approximately U-shaped.

Incomplete Information: In many applications, uncertainty pertains not only to the quality of the project, but also to the productivity of agents. That is, the value of the parameter α of each agent might be unknown. We may model this by assuming that $\alpha \in \{\alpha_L, \alpha_H\}$ is drawn independently across players, with $\alpha_L < \alpha_H$, and some probability q_0 that each agent is of the low type α_L , that is, that his productivity is high.

Solving for an equilibrium is difficult, because agents are updating on both the quality of the project and the productivity of their opponent, and this updating depends on the effort

¹⁹Such conditioning is meaningful, since the overall effort over time is still independent of n .

choices, which are private information. While we have not attempted to establish uniqueness, the following constitutes a symmetric Bayes Nash equilibrium path (for those values of the parameters for which the solution is not trivial). The low-type (high-productivity) agent starts by exerting effort by himself. As time passes, he quickly becomes more pessimistic, while the high-type agent, who does not work during that time, does not update his belief downward as fast (the lack of success is not as surprising to him because he does not work). That is, the private beliefs of the two agent's types about the quality of the project diverge. At some time, the relative optimism of the high-cost agent more than offsets his cost disadvantage and he starts exerting effort. Simultaneously, the low-cost agent stops working once and for all. The high-cost agent's effort then dwindles over time and his belief converges asymptotically to α_H , his Marshallian threshold. Because of his initial effort level, the private belief of the low-cost agent remains forever below the high-cost agent's and converges asymptotically to a level below his own Marshallian threshold, namely α_L . The effort trajectories in such an equilibrium are shown in Figure 11 below.²⁰

An interesting open question is what happens when effort is observable. Since the level of effort conveys information about the agent's type, this might give rise to ratcheting, as low-cost agents might want to hide their private information. This might further depress the exertion of effort in the observable case, but a careful analysis is left for future research.

References

- [1] Admati, A.R. and M. Perry 1991. "Joint Projects without Commitment," *Review of Economic Studies*, **58**, 259–276.
- [2] Alchian, A. and H. Demsetz, 1972. "Production, information costs, and economic organization," *American Economic Review*, **62**, 777–795.

²⁰Note that effort can be increasing over time. This is due to the fact that, because the probability that agent $-i$ assigns to agent i being of the low type decreases over time, the effort agent $-i$ needs to exert to keep player i indifferent between exerting effort or not might actually increase.




Figure 11: Effort under incomplete information

- [3] Bag, P.K. and S. Ray, 2008. “Repeated Charitable Contributions Under Incomplete Information,” *Economic Journal*, **118**, 60–91.
- [4] Bergemann, D. and U. Hege, 2005. “The Financing of Innovation: Learning and Stopping,” *RAND Journal of Economics*, **36**, 719–752.
- [5] Bergin, J. and W.B. McLeod, 1993. “Continuous Time Repeated Games,” *International Economic Review*, **34**, 21–37.
- [6] Bliss, C. and B. Nalebuff, 1984. “Dragon-Slaying and Ballroom Dancing,” *Journal of Public Economics*, **25**, 1–12.
- [7] Bolton P. and C. Harris, 1999. “Strategic Experimentation,” *Econometrica*, **67**, 349–374.
- [8] Compte, O. and P. Jehiel, 2004. “Gradualism in Bargaining and Contribution Games,” *Review of Economic Studies*, **71**(4), 975–1000.
- [9] Dasgupta, P., 1988. “Patents, Priority and Imitation or, The Economics of Races and Waiting Games,” *Economic Journal*, **98**, 66–80.
- [10] Doz, Y.L., 1996. “The Evolution of Cooperation in Strategic Alliances: Initial Conditions or Learning Processes?,” *Strategic Management Journal*, **17**(7), 619–652.

- [11] Ellison, G., 2002. "The Slowdown of the Economics Publishing Process," *Journal of Political Economy*, **105(5)**, 947–993.
- [12] Fershtman, C. and S. Nitzan, 1991. "Dynamic voluntary provision of public goods," *European Economic Review*, **35**, 1057–1067.
- [13] Gradstein, M., 1992. "Time Dynamics and Incomplete Information in the Private Provision of Public Goods," *Journal of Political Economy*, **100**, 581–597.
- [14] Gulati, R., 1995. "Social Structure and Alliance Formation Patterns: A Longitudinal Analysis," *Administrative Science Quarterly*, **40**, 85–112.
- [15] Gulati, R. and H. Singh, 1998. "The Architecture of Cooperation: Managing Coordination Costs and Appropriation Concerns in Strategic Alliances," *Administrative Science Quarterly*, **43(4)**, 781–814.
- [16] Hamilton, B.H., J.A. Nickerson and H. Owan, 2003. "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation," *Journal of Political Economy*, **111(3)**, 465–497.
- [17] Holmström, B., 1982. "Moral Hazard in Teams," *The Bell Journal of Economics*, **13(2)**, 324–340.
- [18] Hopenhayn, H. and F. Squintani, 2008. "Preemption Games with Private Information," working paper, University College London.
- [19] Kamihigashi, T., 2001. "Necessity of Transversality Conditions for Infinite Horizon Problems," *Econometrica*, **69(4)**, 995–1012.
- [20] Karau, S. and K.D. Williams, 1993. "Social Loafing: A Meta-analytic Review and Theoretical Integration," *Journal of Personality and Social Psychology*, **65**, 681–706.
- [21] Keller, G., S. Rady and M. Cripps, 2005. "Strategic Experimentation with Exponential Bandits," *Econometrica*, **73**, 39–68.

- [22] Klein, N. and S. Rady, 2008. “Negatively Correlated Bandits,” Discussion Paper No. 243, SFB/TR 15, Munich.
- [23] Latané, B., K. Williams and S. Harkins, 1979. “Many hands make light the work: The causes and consequences of social loafing,” *Journal of Personality and Social Psychology*, **37(6)**, 822–832.
- [24] Legros, P. and S. Matthews, 1993. “Efficient and Nearly-Efficient Partnerships,” *Review of Economic Studies*, **60(3)**, 599–611.
- [25] Levi D., 2007. *Group Dynamics for Teams*. Sage Publications.
- [26] Levin, J. and S. Tadelis, 2005. “Profit Sharing and the Role of Professional Partnerships,” *The Quarterly Journal of Economics*, **120**, 131–171.
- [27] Lockwood, B. and J.P. Thomas, 2002. “Gradualism and Irreversibility,” *Review of Economic Studies*, **69(2)**, 339–356.
- [28] Ma, C.-t. A., J. Moore, and S. Turnbull, 1988. “Stopping agents from Cheating,” *Journal of Economic Theory*, **46(2)**, 355–372.
- [29] Marx, L. and S. Matthews, 2000. “Dynamic Voluntary Contribution to a Public Project,” *Review of Economic Studies*, **67(2)**, 327–358.
- [30] McGrath, 1991. “Time, interaction, and performance (TIP): A theory of groups,” *Small Group Behavior*, **4(3)**, 147–174.
- [31] Murto, P. and J. Välimäki, 2008. “Learning and Information Aggregation in an Exit Game,” working paper, Helsinki Center of Economic Research.
- [32] O’Donoghue, T. and M. Rabin, 1999. “Incentives for procrastinators,” *Quarterly Journal of Economics* **104**, 769–816.
- [33] Olson, M., 1965. *The Logic of Collective Action*. Harvard University Press.

- [34] Presman, E.L., 1990. “Poisson Version of the Two-Armed Bandit Problem with Discounting,” *Theory of Probability and its Applications*, **35**, 307–317.
- [35] Presman, E.L. and I.N. Sonin, 1990. *Sequential Control with Incomplete Information: The Bayesian Approach to Multi-Armed Bandit Problems*. Academic Press, London.
- [36] Rosenberg, D., E. Solan and N. Vieille, 2007. “Social Learning in One-Arm Bandit Problems,” *Econometrica*, **75(6)**, 1591–1611.
- [37] Seierstad, A. and K. Sydsaeter, 1987. *Optimal Control Theory with Economic Applications*. North-Holland.
- [38] Strausz, R., 1999. “Efficiency in Sequential Partnerships,” *Journal of Economic Theory*, **85**, 140–156.
- [39] Tuckman B.W., and M.C. Jensen, 1977. “Stages of small-group development revisited,” *Group and Organization Studies*, **2(4)**, 419–427.
- [40] Williams S. and R. Radner, 1988. “Efficiency in partnerships when the joint output is uncertain,” in *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability: Essays in Honor of Stanley Reiter*. J. Ledyard, Ed., Kluwer Academic, Boston.
- [41] Winter, E., 2004. “Incentives and Discrimination,” *American Economic Review*, **94(3)**, 764–773.

Appendix

A Proofs for Section 3

Proof of Theorem 1: (Preliminaries.) Observe first that, since $\dot{p}_t = -p_t(1 - p_t) \sum_i u_{i,t}$, we can further rewrite the objective in (1) as

$$r \int_0^\infty \left(-\frac{\dot{p}_t}{1 - p_t} + \alpha \left(\frac{\dot{p}_t}{p_t(1 - p_t)} + u_{-i,t} \right) \right) \frac{1 - \bar{p}}{1 - p_t} e^{-rt} dt,$$

where $u_{-i,t} := \sum_{j \neq i} u_{j,t}$. Applying integration by parts to the objective, we obtain, ignoring constant terms,

$$\int_0^\infty \left(r\alpha \ln \frac{p_t}{1 - p_t} + \frac{r(\alpha - 1) + \alpha u_{-i,t}}{1 - p_t} \right) e^{-rt} dt.$$

Making the further change of variable $x_t = \ln((1 - p_t)/p_t)$, and defining $\beta := 1/\alpha - 1$, agent i maximizes

$$\int_0^\infty (-x_t + e^{-x_t}(u_{-i,t}/r - \beta)) e^{-rt} dt, \text{ such that } \dot{x}_t = u_{i,t} + u_{-i,t},$$

over functions $u_{i,t}$ in $[0, \lambda]$, given the function $u_{-i,t}$. The Hamiltonian for this problem is given by

$$H(u_{i,t}, x_t, \gamma_{i,t}) = (-x_t + e^{-x_t}(u_{-i,t}/r - \beta)) e^{-rt} + \gamma_{i,t}(u_{i,t} + u_{-i,t}).$$

We now drop the subscript i . It is easy to see that no agent exerts effort if $p_t < \alpha$ (consider the original objective function: if $p_t < \alpha$, then choosing $u_{i,t} = 0$ is clearly optimal). We therefore assume that $\bar{p} > \alpha$, which is equivalent to $x_0 < \ln \beta$, where $x_0 = \ln((1 - \bar{p})/\bar{p})$. The assumption on the discount rate is equivalent to $1 + e^{-x_0}(\lambda/r - \beta) > 0$.

(Necessary Conditions.) By Pontryagin's principle, there must exist a continuous function γ_i such that, for each i ,

1. $u_{i,t}$ maximizes $\gamma_{i,t}(u_{i,t} + u_{-i,t})$.
2. $\dot{\gamma}_{i,t} = r\gamma_t + 1 + e^{-x_t}(u_{-i,t}/r - \beta)$.
3. If x^* is the optimal trajectory, $\lim_{t \rightarrow \infty} \gamma_{i,t}(x_t^* - x_t) \leq 0$ for all feasible trajectories x_t .

The last condition is the transversality condition, which follows here from Kamihigashi (2001). Since there is a co-state variable γ_i for each player, we are led to consider a phase diagram in \mathbb{R}^{n+1} , with dimensions $\gamma_1, \dots, \gamma_n$, and x .

(Candidate Equilibrium.) We first show that the candidate equilibrium strategy $u_{i,t}^*$ and the corresponding beliefs function x_t^* satisfy the necessary conditions. Consider a strategy generating a trajectory that starts at $(\gamma_1, \dots, \gamma_n, x_0) = (0, \dots, 0, x_0)$, and has $u_{i,t} = u_{i,t}^* := r(\beta - e^{x_t})/(n-1)$. This implies that $\gamma_{i,t} = 0$ along the trajectory. Observe that $u_{i,t}^* > 0$ as long as $x_t < \ln \beta$, and is decreasing in t , with limit 0 as $t \rightarrow \infty$. Indeed, the solution is $x_t^* = \ln \beta - \ln(1 + (\beta e^{-x_0} - 1)e^{-(n/(n-1))r\beta t})$, and $u_{i,t}^* = r\beta/((\beta e^{-x_0} - 1)^{-1}e^{(n/(n-1))r\beta t} + 1)$, which corresponds to the expression in the text. Indeed, this trajectory has $x_t \rightarrow \ln \beta$, and $\gamma_{i,t}^* = 0$, for all t .

(Uniqueness.) We now use the trajectory $(\gamma_{1,t}^*, \dots, \gamma_{n,t}^*, x_t^*)$ as a reference to eliminate other trajectories, by virtue of the transversality condition. We shall divide all possible paths into several subsets:

1. Consider paths that start with $\gamma_j \geq 0$ for all j , with strict inequality $\gamma_i > 0$ for some i . Since $\gamma_i > 0$, $u_i = \lambda$, and so $\dot{\gamma}_j > 0$ for all j . So we might as well consider the case $\gamma_j > 0$ for all j . Then for all j , we have $u_j > 0$ and γ_j strictly increasing. It follows that $\gamma_1, \dots, \gamma_n$, and x all diverge to $+\infty$. Given the reference path along which x converges, such paths violate the transversality condition.
2. Consider paths that start with $\gamma_i \leq 0$ for all i , with strict inequality $\gamma_i < 0$ for all but one agent j . We then have $u_{-j} = 0$. Since $\bar{p} > \alpha$ implies that $r\gamma_j + 1 - \beta e^{-x_0} < 0$, it follows that $\dot{\gamma}_j < 0$, and we might as well assume that $\gamma_i < 0$ for all i . So we have $u_i = 0$ for all i , and x remains constant, and all γ_i diverge to $-\infty$. Since x_0 is less than $\ln \beta$, the limit of our reference trajectory, this again violates the transversality condition. The same argument rules out any path that enters this subset of the state space, provided it does so for $x_t < \ln \beta$. However, we do not rule out the case of $\gamma = (\gamma_1, \dots, \gamma_n) \leq 0$ with two or more $j : \gamma_j = 0$ and $u_{-j} > 0$.
3. Consider paths that start with some $\gamma_i < 0$ for all agents $i \neq j$, and with $\gamma_j > 0$. Assume further that $r\gamma_j + 1 - \beta e^{-x_0} \geq 0$. Because $u_j > 0$, $\dot{x}_t > 0$ and so we might as well assume

that $\dot{\gamma}_j \geq r\gamma_j + 1 - \beta e^{-x_0} > 0$. It then follows that $u_j > 0$ forever, and so γ_j diverges to $+\infty$, as does x . This again violates the transversality condition. If there is more than one $j : \gamma_j > 0$, then $u_{-j} \geq \lambda$, and $1 + e^{-x_0}(\lambda/r - \beta) > 0$ implies that a fortiori $\dot{\gamma}_j > 0$. The same argument then applies.

4. Consider paths that start with some $\gamma_i < 0$ for all $i \neq j$, and with $\gamma_j > 0$. Assume further that $r\gamma_j + 1 - \beta e^{-x_0} < 0$. Since $u_j > 0$ as long as $\gamma_j > 0$, the trajectory must eventually leave this subset of parameters, either because $\gamma_j \leq 0$, and then we are back to case 2, or because $\dot{\gamma}_j \geq 0$, and then we are back to case 3. If the trajectory enters one of the previous subsets, it is not admissible for the reasons above. Therefore, the only case left is if this trajectory hits $(\gamma_1, \dots, \gamma_n, x) \leq (0, \dots, 0, \ln \beta)$ with at least two j such that $\gamma_j = 0$. This is case 5. Notice that if there were more than one $j : \gamma_j > 0$, then $u_{-j} \geq \lambda$, and $1 + e^{-x_0}(\lambda/r - \beta) > 0$ would imply $\dot{\gamma}_j > 0$ even if $r\gamma_j + 1 - \beta e^{-x_0} < 0$, bringing us back to case 3.
5. Consider paths that start with $(\gamma_1, \dots, \gamma_n, x) \leq (0, \dots, 0, \ln \beta)$ with at least two j such that $\gamma_j = 0$. Let $A_t = \{j : \gamma_{j,t} = 0\}$. Then there is a unique solution $u_{j,t}^*(|A|) := r(\beta - e^{x_t})/(|A| - 1)$, such that $\dot{\gamma}_j = 0$ for all $j \in A$ (as long as $x_t < \ln \beta$, since $u_{j,t}^* = 0$ when $x_t = \ln \beta$). Along this trajectory, $x_t \rightarrow \ln \beta$. Furthermore, the effort levels must switch to $u_{j,t}^*(|A_t| + 1)$ for all $j \in A \cup \{i\}$ whenever $\gamma_i = 0$ for $i \notin A_t$. Similarly if two or more $i \notin A_t$ hit $\gamma_i = 0$ at the same time. We show this by ruling out all other cases. Any policy with $u_{-j} < u_{-j}^*(|A|)$ for all j implies $\dot{\gamma}_j < 0$, leading to case 2. Any policy with $u_{-j} > u_{-j}^*(|A|)$ leads to case 1. Finally, any policy different from $u_j^*(|A|)$ can lead to two or more $\gamma_j > 0$ (case 3), or to a single $\gamma_j > 0$ (cases 3 and 4). This leaves us with the only possible scenario, $u_j = u_j^*(|A|)$ for all $j \in A$, and this is precisely the candidate trajectory examined earlier.

We have thus eliminated all but one family of paths. These paths start with at most one i exerting $u = \lambda$, then switching (before the beliefs have reached $\ln \beta$) to two or more agents (including i) who play the reference strategy $u_{i,t}^*(|A_t|)$, as if only agents $i \in A_t$ were present in the team. At any point in time before the beliefs have reached the critical level, more agents may be added to A (but not subtracted). In that case, the policy switches to the appropriate strategy

$u_j^*(|A|)$. That is, all candidate equilibria have several phases. In the first phase, one player experiments alone. In the subsequent phases, all (active) players experiment equally, adding new players at any point in time. Of course, there are extreme cases in which some phase is non-existent. Therefore, the only symmetric equilibrium is one in which the $|A_0| = n$, that is, all players exert effort $u_j^*(n)$ from the start.

(Sufficiency.) We are left with proving that these candidate equilibria are indeed equilibria. While the optimization programme described above is not necessarily concave in x , observe that, defining $q_t := p_t/(1 - p_t)$, it is equivalent to

$$\max_{u_i} \int_0^\infty (\ln q_t + q_t \left(\frac{u_{-i,t}}{r} - \beta \right)) e^{-rt} dt \text{ s.t. } \dot{q}_t = -q_t(u_{i,t} + u_{-i,t}).$$

so that the maximized Hamiltonian is concave in q , and sufficiency then follows from the Arrow sufficiency theorem (see Seierstad and Sydsaeter (1987), Thm. 3.17). Therefore, all these paths are equilibria. \square

Proof of Lemma 1: (1.) From expression (2), it is clear that individual effort is decreasing in t , and that for a fixed t , $u_{i,t}^*$ is increasing in r and \bar{p} , and decreasing in α .

(2.) Similarly, individual and total efforts are decreasing in n for low values of t , and increasing in n when t is higher.

(3.) Substituting the expression for x_t^* , equilibrium payoffs can be written as

$$V(n) = r \int_0^\infty \left(\ln \left(1 + (\beta e^{-x_0} - 1) e^{-\frac{n}{n-1} r \beta t} \right) - \ln \beta - 1 \right) e^{-rt} dt.$$

Let $k := (\beta e^{-x_0} - 1)^{-1}$. The change of variable $y = e^{-nt}$ allows us to write payoff as

$$V(n) = -1 - \ln \beta + \frac{n-1}{n\beta} \int_0^1 \ln(1 + y/k) y^{\frac{n-1}{n\beta} - 1} dy,$$

which is independent of r .

(4.) Given the equilibrium strategies, the probability of a success occurring is given by

$$\int_0^\infty f(s) ds = \int_0^\infty \frac{1}{1 + ke^s} e^{-\frac{s}{1+ke^s}} ds,$$

where $s = nr\beta t / (n - 1)$. It is therefore independent of n . Let $\tau \in \mathbb{R}_+ \cup \{\infty\}$ denote the random time at which a breakthrough arrives. The conditional expected time of a breakthrough is then given by

$$\mathbb{E}[\tau | \tau < \infty] = \frac{n-1}{nr\beta} \int_0^\infty s f(s) ds \Big/ \int_0^\infty f(s) ds ,$$

which is increasing in n . Also note that $\mathbb{E}[r\tau | \tau < \infty]$ is independent of r . \square

Two-player, asymmetric case: Assume that players are asymmetric, in the sense that $\alpha_1 < \alpha_2$, which implies $\ln \beta_1 > \ln \beta_2$. The nontriviality condition becomes now that $p > \alpha_1$, while we maintain the patience assumption $1 + e^{-x_0}(\lambda/r - \beta_i) > 0$.

(Necessary Conditions.) There must exist a continuous function γ_i such that, for each i ,

1. $u_{i,t}$ maximizes $\gamma_{i,t}(u_{i,t} + u_{-i,t})$.
2. $\dot{\gamma}_{i,t} = r\gamma_t + 1 + e^{-x_t}(u_{-i,t}/r - \beta_i)$.
3. If x^* is the optimal trajectory, $\lim_{t \rightarrow \infty} \gamma_{i,t}(x_t^* - x_t) \leq 0$ for all feasible trajectories x_t .

(Candidate Equilibrium.) We consider a phase diagram in \mathbb{R}^3 , with dimensions γ_1 , γ_2 , and x . Consider the trajectory that starts from some (γ_1, γ_2, x) , with $\gamma_1 > 0$, $\gamma_2 < 0$ and $\dot{\gamma}_1 = r\gamma_1 + 1 - \beta_1 e^{-x} < 0$ (i.e. it has $u_{1,t} = \lambda$ and $u_{2,t} = 0$ to begin with) such that it reaches $(\bar{\gamma}_1, \bar{\gamma}_2, \bar{x})$, with $\bar{\gamma}_1 = 0$, $\bar{\gamma}_2 < 0$, $\dot{\gamma}_1 = \dot{\gamma}_2 = 0$, and $\bar{x} = \ln \beta_1$. At this point $(\bar{\gamma}_1, \bar{\gamma}_2, \bar{x})$, $u_{i,t} = 0$ for all t , and the trajectory stops.

(Uniqueness.) To prove that this is the unique equilibrium outcome, we divide this space into several subsets.

1. Consider any path that starts with $\gamma_i > 0, \gamma_j \geq 0$. This case is analogous to case (1) in the proof of Theorem 1.
2. Consider paths that start with $\gamma_i < 0, \gamma_j \leq 0$. There are several subcases, depending on the initial condition.
 - (a) $r\gamma_k + 1 - \beta_k e^{-x} > 0$ for either $k = i$ or j . Then γ_k diverges to $+\infty$, and so must x . This violates the transversality condition.

- (b) $r\gamma_2 + 1 - \beta_2 e^{-x} = 0$. Given that $x_0 < \ln \beta_1$, it follows that $\dot{\gamma}_1 < 0$ (unless possibly $\gamma_2 = 0$ and $u_2 > 0$, in which case, however, $\dot{x} > 0$ and so the trajectory immediately enters the previous subcase). So γ_1 diverges to $-\infty$, and x remains constant at a level strictly below \bar{x} . Again, this violates the transversality condition.
- (c) $r\gamma_2 + 1 - \beta_2 e^{-x} < 0$. As in the previous case, it follows that $\dot{\gamma}_1 < 0$ (with the same caveat as before), so γ_1 diverges to $-\infty$, and so does γ_2 ; x remains constant, and again, the transversality condition is violated.
3. Consider paths that start with $\gamma_i < 0, \gamma_j > 0$, and $r\gamma_j + 1 - \beta_j e^{-x_0} \geq 0$. This case is analogous to (3.) in the proof of Theorem 1.
4. Consider paths that start with $\gamma_i < 0, \gamma_j > 0, r\gamma_j + 1 - \beta_j e^{-x_0} < 0$. There are two subcases:
- (a) $i = 1$. Because $u_2 > 0$ as long as $\gamma_2 > 0$, so that $\dot{x} > 0$, the trajectory must eventually leave this subset of parameters. Note that it must do so for a value of x no larger than $\ln \beta_2$. The only possibility that has not already been ruled out previously is if this trajectory hits $(\gamma_1, \gamma_2, x_t) = (0, 0, x_t)$, for some $x_t \leq \ln \beta_2$. This is ruled out in case 5.
- (b) $i = 2$. Since $u_1 > 0, \dot{x} > 0$; hence, here as well, we must eventually leave this region. The cases not covered so far are if the trajectory hits $(\gamma_1, \gamma_2, x_t) = (0, 0, x)$ for some $x < \ln \beta_1$ (ruled out in case 5), or if it hits $(\gamma_1, \gamma_2, x_t) = (0, \gamma_2, \ln \beta_1)$ for some $\gamma_2 \leq 0$. If $\gamma_2 > \bar{\gamma}_2$, then as in case 3, γ_2 must diverge to $+\infty$, and so must x , violating the transversality condition. If instead $\gamma_2 < \bar{\gamma}_2$, then $u_1 = 0$ identically thereafter, in which case $\gamma_{2,t} \rightarrow -\infty$, and player 2 never exerts effort. This outcome is identical to the one in our reference trajectory. In fact, if player 2 exerts any effort, x must increase at some point, from which point on γ_1 will diverge to $+\infty$, and so will x , violating the transversality condition.
5. Consider paths that start from $(\gamma_1, \gamma_2, x_t) = (0, 0, x)$, for some $x \leq \ln \beta_1$. There are two subcases:

- (a) $x \in (\beta_2, \beta_1)$. Then $\dot{\gamma}_2 > 0$; if also $\dot{\gamma}_1 \geq 0$, we are back to the first case; if instead, $\dot{\gamma}_1 < 0$, we are in the third case. Both cases have already been ruled out.
- (b) $x \leq \ln \beta_2$. Then there is a unique solution $u_{2,t}^*$, given in Theorem 1, such that $\dot{\gamma}_i = 0$. Observe that, unlike in the symmetric case, $u_2^* > 0$ for all $x \leq \ln \beta_2$. For $u_j > u_j^*$, $\dot{\gamma}_i > 0$. There are four cases to consider. Either $\dot{\gamma}_i > 0, \dot{\gamma}_j \geq 0$. Then the region covered in case 1 is entered, and such a path cannot satisfy the necessary conditions. Or $\dot{\gamma}_i \leq 0, \dot{\gamma}_j < 0$, but then the region covered in case 2 is entered, and again this path can be ruled out. Or $\dot{\gamma}_i > 0, \dot{\gamma}_j < 0$, for some $i = 1, 2$, but this would lead to one of the two regions covered in case 4, and given the dynamics there, such a region cannot be entered for a positive interval of time starting from $(0, 0, x)$. Or, finally, $u_i = u_i^*$ for both $i = 1, 2$, but since $u_2^* > 0, \dot{x} > 0$, and so, as in the case 3, γ_2 must diverge to $+\infty$, and so must x , violating the transversality condition.

(Sufficiency.) We have thus eliminated all but one outcome: the one in which the strongest player experiments alone as long as he finds it profitable to do so. Sufficient conditions did not rely on symmetry, hence this path is an equilibrium. \square

B Proofs for Section 4

Throughout, let $x := \ln \frac{1-p}{p}$, so that in particular $x_0 = \ln \frac{1-\bar{p}}{\bar{p}}$, and $\beta = \alpha^{-1} - 1$.

B.1 Proofs for Section 4.1.

Proof of Lemma 2: Suppose that all agents $-i$ exert effort at level $u = 1$ over the entire interval $[0, T]$, where T is the deadline. It is easy to see that agent i has the weakest incentives to exert high effort himself at time 0. Let us compute the payoff of agent i from exerting no effort on time $[0, \delta]$, and effort on the interval $[\delta, T]$, where $\delta \in [0, T]$. It is equal to

$$V(\bar{p}, T) := (1 - \bar{p}) \left(\int_0^\delta \frac{(n-1)p_t}{1-p_t} e^{-rt} dt + \int_\delta^T \frac{np_t - \alpha}{1-p_t} e^{-rt} dt \right),$$

where $p_t = \frac{\bar{p}}{\bar{p} + (1-\bar{p})e^{(n-1)t}}$ for $t \in [0, \delta]$, and $\frac{p_\delta}{p_\delta + (1-p_\delta)e^{n(t-\delta)}}$ for $t \in [\delta, T]$. The term in brackets is decreasing in δ , and setting the derivative with respect to δ equal to 0 at $\delta = 0$ gives

$$T \leq T(\bar{p}) := \frac{1}{n+r} \ln \frac{(n-\alpha)\bar{p}}{\alpha(n-\bar{p}) - r(\bar{p}-\alpha)}.$$

(the argument of the logarithm is larger than 1 for $n \geq 2$.) It is immediate to verify that $T(\bar{p}) < T_n$, which is the time it takes for beliefs to reach α : stopping occurs before beliefs have gone down to the Marshallian threshold. This establishes sufficiency. Necessity follows from Pontryagin's principle, as in Theorem 1, and is accordingly omitted.

If the horizon exceeds $T(\bar{p})$, in a symmetric equilibrium, agents cannot choose effort at level $u = 1$ throughout, and they cannot both choose $u = 0$ either. Arguments similar to those of Theorem 1 establish then that effort must be interior, and as in baseline case, i.e. $u_t = u_t^*$, up to the time T_1 at which $T - T_1 = T(p_{T_1})$. \square

Proof of Theorem 3: Let $V_i(\bar{p}) := V_i(\bar{p}, T(\bar{p}))$. If the deadline is such that effort switches to 1 at time T_1 , the payoff of agent i is then

$$V_i(T_1) := (1-\bar{p}) \left(\int_0^{T_1} \frac{(np_t - \alpha)}{1-p_t} u_{i,t}^* e^{-rt} dt + e^{-rT_1} \frac{V(p_{T_1})}{1-p_{T_1}} \right),$$

where p_t solves $\dot{p}_t = -p_t(1-p_t)nu_{i,t}^*$, $p_0 = \bar{p}$. Taking derivatives with respect to T_1 , and considering the derivative at $T_1 = 0$ gives

$$\left. \frac{dV_i(T_1)}{dT_1} \right|_{T_1=0} = \frac{\alpha((n-1)\bar{p} + r) - \bar{p}r}{(n-\alpha)(n-1)\bar{p}^2} \left((n-\alpha)\bar{p} - \alpha(n-\bar{p}) \left(\frac{(n-\alpha)\bar{p}}{\alpha(n-\bar{p}) - r(\bar{p}-\alpha)} \right)^{\frac{n}{n+r}} \right).$$

The derivative with respect to r of the term in parenthesis has a derivative equal to (up to a positive multiplicative constant)

$$\ln \left(\frac{(n-\alpha)\bar{p}}{\alpha(n-\bar{p}) - r(\bar{p}-\alpha)} \right) - \left(\frac{(n-\alpha)\bar{p}}{\alpha(n-\bar{p}) - r(\bar{p}-\alpha)} - 1 \right) \leq 0, \quad (2)$$

so $\left. \frac{dV_i(T_1)}{dT_1} \right|_{T_1=0}$ is decreasing in r . Since $\left. \frac{dV_i(T_1)}{dT_1} \right|_{T_1=0, r=0} = 0$, it follows that $\left. \frac{dV_i(T_1)}{dT_1} \right|_{T_1=0} \leq 0$.

Because it is optimal to have agents choose level of effort $u = 1$ as long as possible, the optimal value is $T_1 = 0$: agents should be given a deadline for which it is optimal to exert at a maximal rate immediately.

Finally, differentiating $T(\bar{p})$ with respect to r , one can show that $-\partial T / \partial r$ is exactly equal to expression (2), and so the optimal deadline is increasing in the discount rate.

B.2 Proofs for Section 4.2.

The Agent's Problem: Given a wage function w_t , each agent maximizes

$$\int_0^\infty \frac{\sum_j u_j p_t w_t - \alpha u_i}{1 - p_t} e^{-rt} dt.$$

Integrating by parts and ignoring constant terms (assuming, as will be verified, that w is differentiable) gives

$$\frac{w_0}{1 - \bar{p}} + \int_0^\infty \left(\frac{r(\alpha - w) + \alpha u_{-i} + \dot{w}}{1 - p} + r\alpha \ln \frac{p}{1 - p} \right) e^{-rt} dt.$$

In terms of the function x , agent i maximizes

$$w_0 (1 + e^{-x_0}) + \int_0^\infty (-r\alpha x + (r(\alpha - w) + \alpha u_{-i} + \dot{w})e^{-x}) e^{-rt} dt,$$

over functions u_i such that $\dot{x} = \sum_j u_j$. Applying Pontryagin's principle gives

$$\dot{w} - rw = -\alpha (re^x + r + u_{-i}),$$

which generalizes the earlier formula for $w = 1$. Note that this formula will hold even if $u = 1$ over some interval, as the principal cannot gain from giving agents strict rather than weak incentives to exert maximal effort.

Proof of Theorem 4 (Agents' Optimal Wage): Agents induce symmetric levels of effort $u_{i,t} = u_t$ in order to maximize

$$\int_0^{\infty} nu_t (e^{-x}v - \alpha (1 + e^{-x})) e^{-rt} dt,$$

subject to

$$\begin{aligned} \int_0^{\infty} nu_t (v - nw_t) e^{-x} e^{-rt} dt &= 0, \\ rw - \alpha (re^x + r + (n-1)u_t) &= \dot{w}, \\ nu_t &= \dot{x}. \end{aligned}$$

Integrating by parts both the objective and the first constraint gives the following program

$$\begin{aligned} \max_{u_t} \int_0^{\infty} - (e^{-x} (v - \alpha) + \alpha x) e^{-rt} dt, \\ \text{s.t. } \dot{x} = nu_t, \dot{w} = rw - \alpha (re^x + r + (n-1)u_t), \\ e^{-x_0} (v - nw_0 + \alpha (n-1)) + n\alpha - r(v - \alpha) \int_0^{\infty} e^{-x} e^{-rt} dt = 0. \end{aligned}$$

Associate the co-state μ to the constraint $\dot{w} = rw - \alpha (re^x + r + (n-1)u)$, σ to the isoperimetric constraint $k_0 = \int_0^{\infty} e^{-x} e^{-rt} dt$ and γ to $\dot{x} = nu$. Given the constraint $u \leq 1$, with associated Lagrangian coefficient λ , the Hamiltonian is

$$H = - (e^{-x} (v - \alpha) + \alpha x) e^{-rt} + \sigma e^{-x} e^{-rt} + \gamma nu + \mu (rw - \alpha (re^x + r + (n-1)u)) + \lambda (1 - u),$$

with $\lambda \geq 0$ and $\lambda(1 - u) = 0$. So $\mu = \mu_0 e^{-rt}$, and

$$\gamma n - \alpha \mu (n-1) - \lambda = 0.$$

Observe that, if $\lambda = 0$ over some interval, then $\gamma = \alpha \frac{n-1}{n} \mu_0 e^{-rt}$ and therefore $\dot{\gamma} = -r\alpha \frac{n-1}{n} \mu_0 e^{-rt}$.

But we would then have

$$\begin{aligned} -\dot{\gamma} &= -\left(-e^{-x}(v - \alpha) + \alpha\right)e^{-rt} - \sigma e^{-x}e^{-rt} - \alpha\mu r e^x, \\ r\alpha\frac{n-1}{n}\mu_0 &= e^{-x}(v - \alpha - \sigma) - \alpha - \alpha\mu_0 r e^x, \end{aligned}$$

from which it follows that x is constant over that interval, a contradiction. So $\lambda > 0$ and $u = 1$. Hence, $x = x_0 + nt$ (or more precisely, on any subinterval on which $u = 1$), and we obtain the following ordinary differential equation (hereafter, ODE) for the wage:

$$\dot{w} = rw - \alpha\left(re^{x_0+nt} + r + (n-1)\right).$$

It remains to determine the initial value of w , or equivalently, the time at which all effort stops. Once this time T is set, we know that the wage must satisfy the terminal condition $w_T = \alpha/p_T$. Therefore, the planner's problem reduces to finding the positive root T of the "budget" function $\int_0^T n(v - nw_t)e^{-x}e^{-rt}dt$, where w_t is given by the previous ODE. This gives, upon simplification, that T is the root of

$$\frac{r}{n+r}\left(\frac{v}{\alpha} - 1\right)(e^{(n+r)T} - 1) - (n-1)e^{nT}(e^{rT} - 1) - \frac{nr}{n-r}e^{x_0+nT}(e^{nT} - e^{rT}) = 0.$$

Setting $\delta := e^{nT}$, the left-hand side admits a single positive critical point, which gives precisely the solution to the monopolist's problem (to be described below). This means that either the left-hand side is positive when $T \rightarrow \infty$, in which case it admits no strictly positive root, or it is strictly negative, in which it admits exactly one root, and that root exceeds the optimal T for the monopolist. Clearly, if $n > r$, the left-hand side is negative for large enough r , and we assume so thereafter. To compare the belief at the time of stopping with the asymptotic belief in the baseline model, $\alpha n/v$ (the solution to $pv/n = \alpha$), we plug in the value of v/α corresponding to this asymptotic threshold (i.e., $v/\alpha = ne^{x_0+nT}$), and argue that the left-hand side is negative for r small enough (recall that the result was only claimed for r low enough). Indeed, define $X = nT$, and consider a series expansion around $r = 0$; the left-hand side equals, to the second

order,

$$-(n-1)(1+(X-1)e^X)r + o(r),$$

which is always negative, since $1+(X-1)e^X \geq 0$. Thus, the root is such that the belief has not yet reached the asymptotic belief of the baseline model. In fact, the limit of the left-hand side as $r \rightarrow 0$ does not vanish for $v/\alpha = ne^{x_0+nT}$, implying that the difference in those beliefs is strict, even in the limit.

Proof of Theorem 5 (Principal's Optimal Wage): Let us now turn to the problem of the profit-maximizing principal. The project is worth v to him. So the value of the project is (proportional to)

$$\int_0^\infty \frac{up_t(v-nw_t)}{1-p_t} e^{-rt} dt,$$

subject to $\dot{x} = nu$ and $\dot{w} = rw - \alpha(re^x + r + (n-1)u)$ ($u := u_i$). Integrating by parts and ignoring irrelevant terms gives

$$(v-nw_0)e^{-x_0} - \int_0^\infty e^{-x}(rv - n(rw - \dot{w}))e^{-rt} dt,$$

or

$$(v-nw_0)e^{-x_0} - \int_0^\infty e^{-x}(rv - n\alpha(re^x + r + (n-1)u))e^{-rt} dt.$$

(Observe that there is a term in the integrand that is independent of x, u, w and can be ignored.) Associate the co-state μ to the constraint $\dot{w} = rw - \alpha(re^x + r + (n-1)u)$, and γ to $\dot{x} = nu$. Given the constraint $u \leq 1$, with associated Lagrangian coefficient λ , the Hamiltonian is

$$-(e^{-x}(rv - n\alpha(r + (n-1)u))e^{-rt}) + \gamma nu + \mu(rw - \alpha(re^x + r + (n-1)u)) + \lambda(1-u),$$

with $\lambda \geq 0$ and $\lambda(1-u) = 0$. So $\mu = \mu_0 e^{-rt}$, and

$$(n\gamma - \lambda)e^{rt} = (n-1)\alpha\mu_0 - n(n-1)\alpha e^{-x}.$$

Therefore

$$n\dot{\gamma}e^{rt} - \dot{\lambda}e^{rt} + r((n-1)\alpha\mu_0 - n(n-1)\alpha e^{-x}) = n^2(n-1)u\alpha e^{-x}.$$

We also have

$$-e^{-x}(rv - n\alpha(r + (n-1)u))e^{-rt} + \mu_0 e^{rt}\alpha r e^x = \dot{\gamma},$$

which implies that

$$ne^{-x}(\alpha - v) + \alpha\mu_0(n-1 + ne^x) = \frac{\dot{\lambda}}{r}e^{rt}.$$

Observe that, if $\dot{\lambda} = 0$ over some interval, then

$$\mu_0 e^x(n-1 + ne^x) = n\frac{v-\alpha}{\alpha},$$

from which it follows that x is constant over that interval, a contradiction. So λ is not constant, and therefore, $u = 1$. Hence, $x = x_0 + nt$ (or more precisely, on any subinterval on which $u = 1$), which gives

$$ne^{-(x_0+nt)}(\alpha - v) + \alpha\mu_0(n-1 + ne^{x_0+nt}) = \frac{\dot{\lambda}}{r}e^{rt},$$

which can be integrated for λ and γ then follows. It remains to determine the initial value of w , or equivalently, the time at which all effort stops. Integrating the differential equation

$$\dot{w} - rw = -\alpha \left(r\frac{1-\bar{p}}{\bar{p}}e^{nt} + r + n - 1 \right)$$

gives w_t , as a function of the as yet undetermined initial value $w(\delta)$, as given in Theorem 4. Let T denote the time at which all effort stops. Given that $p_T w_T = \alpha$, and since $p_T = \frac{\bar{p}}{\bar{p} + (1-\bar{p})e^{nT}}$ we can solve for $w_0 = w(\delta)$ as a function of $\delta = e^{nT}$. We may now explicitly solve for the principal's payoff, as a function of δ , and directly verify that it is concave in δ , with a maximum achieved at the value of δ given in Section 4.2. The details are omitted.

To show that the profit-maximizing time horizon T^* is shorter than the welfare-maximizing deadline T , let $v = n$ and $r = 0$. The difference between the two horizons may be written as

$$e^{nT^*} - e^{nT} = \frac{\bar{p}}{2n(1-\bar{p})} \left(\sqrt{(n-1)^2 + 4n\frac{1-\bar{p}}{\bar{p}}(n/\alpha - 1)} - (n-1) \right) - \frac{(n-\alpha)\bar{p}}{\alpha(n-\bar{p})}.$$

Simplifying yields

$$e^{2nT^*} - e^{2nT} \alpha - (n - \bar{p})^{-2} \bar{p}^{-1} \alpha^{-2} (\bar{p} - \alpha) (1 - \bar{p}) n^2 < 0,$$

since $\alpha < \bar{p}$. \square

C Proofs for Section 5

Proof of Theorem 6: We now consider the case with complementarities

$$\dot{p}_t = -p_t(1 - p_t) f(u_{1,t}, \dots, u_{n,t}).$$

Writing in terms of the log-likelihood ratio, we obtain $x_t = \ln(1 - p_t)/p_t$, and $\dot{x}_t = f(u_{1,t}, \dots, u_{n,t})$.

The objective is

$$r \int_0^\infty (p_t f(u_{1,t}, \dots, u_{n,t}) - \alpha u_{i,t}) e^{-rt - \int_0^t p_s f(u_{1,s}, \dots, u_{n,s}) ds} dt = r \left(1 - (1 - \bar{p}) \int_0^\infty \frac{r + \alpha u_{i,t}}{1 - p_t} e^{-rt} dt \right),$$

where the equality follows from integration by parts. So we are minimizing

$$\int_0^\infty (r + \alpha u_{i,t}) (1 + e^{-x_t}) e^{-rt} dt \text{ such that } \dot{x}_t = f(u_{1,t}, \dots, u_{n,t}).$$

Pontryagin's principle gives (assuming an interior solution)

$$\alpha (1 + e^{-x}) e^{-rt} + \lambda f_i(u_1, \dots, u_n) = 0, \text{ and } \dot{\lambda} = (r + \alpha u_i) e^{-x-rt},$$

where f_i is the derivative of f with respect to u_i . Assuming a symmetric solution $(u_{1,t}, \dots, u_{n,t}) = u_t \in \mathbb{R}^n$, and dropping time subscripts, we have

$$\alpha (1 + e^{-x}) e^{-rt} = -\lambda f_i(u), \dot{\lambda} = -(r + \alpha u) e^{-x-rt}, \dot{x} = f(u).$$

So, differentiating the first equation, and defining $v(x) = u(t)$, so that $v'(x) \dot{x} = \dot{u}$,

$$1 - \left(\frac{r}{\alpha} + v\right) \frac{f_i(v)}{f(v)} = \frac{1}{p} \left(\frac{\sum_j f_{ij}(v)}{f_i(v)} v' - \frac{r}{f(v)} \right). \quad (3)$$

This ODE characterizes the interior solution (if any) of the problem. For $f(u_1, \dots, u_n) = (\sum_i u_i^\rho)^{\frac{1}{\rho}}$, and $u_1 = \dots = u_n$, $\sum_j f_{ij} = 0$, and so the equation can be solved for $u = v$, i.e., defining $k_n := n^{1-1/\rho}$,

$$u(p) = \frac{r}{n-1} \left(\frac{1}{\alpha} - \frac{k_n}{p} \right).$$

Again, it is simple to verify that, if $\alpha^{-1} \geq k_n/\bar{p}$, agents do not exert effort at a maximal rate on any interval of time in equilibrium. Sufficiency follows as in the baseline model. While it is not possible to solve for the function p_t , its inverse $t(p)$ can be computed, and it can be verified that $\lim_{p \rightarrow \alpha k_n} t(p) = \infty$, while $t(p) < \infty$ for $p > \alpha k_n$. \square

Proof of Theorem 7:

The case $\rho \geq 1/2$: Writing equation (3) for each player separately, we obtain

$$\frac{f(u_1, u_2)}{f_i(u_1, u_2)} - \left(\frac{r}{\alpha} + u_i\right) = \frac{1}{p} \left(-\frac{r}{f_i(u_1, u_2)} - \frac{f_{ii}(u_1, u_2)\dot{u}_i + f_{ij}(u_1, u_2)\dot{u}_j}{f_i(u_1, u_2)^2} \right).$$

Inserting the C.E.S. function, and defining the ratio $\sigma_i(p) = u_i(p)/u_j(p)$, We obtain

$$\left(1 - \frac{r}{\alpha u_i} \sigma_i^\rho\right) (u_1^\rho + u_2^\rho)^{1/\rho} = \frac{1}{p} \left((1 - \rho) \left(\frac{\dot{u}_i}{u_i} - \frac{\dot{u}_j}{u_j} \right) - r(1 + \sigma_i^\rho) \right). \quad (4)$$

By adding equations (4) for each player, we obtain the relationship that needs to hold between u_1 and u_2 . Formally,

$$u_2 = \frac{r}{2\alpha\sigma_1^\rho} \left(\sigma_1^{2\rho-1} + 1 - \frac{\alpha}{p} (\sigma_1^\rho + 1)^{2-\frac{1}{\rho}} \right). \quad (5)$$

Similarly, subtract equations (4), substitute the expression for u_2 in (5), and use $u_1 = \sigma_1 u_2(\sigma_1)$.

We then obtain the following ordinary differential equation for σ_1 :

$$\frac{(1-\rho)\sigma'_1}{\sigma_1} = \frac{1}{1-p} \left(1 - \frac{2 \left(1 - \frac{\alpha}{p} (\sigma_1^\rho + 1)^{1-1/\rho} \right)}{\sigma_1^{2\rho-1} + 1 - \frac{\alpha}{p} (\sigma_1^\rho + 1)^{2-1/\rho}} \right), \quad (6)$$

where $\sigma_1(\bar{p}) := u_{1,0}/u_{2,0}$. Notice that σ_1 will not be a function of r (hence patience will only influence levels of effort, not the allocation).

(σ bounded.) We now drop the subscript 1 and refer to σ_1 as σ . We first show that σ stays bounded for all p . Notice that when $u_2 = 0$, that is when $\alpha/p = (\sigma^{2\rho-1} + 1)/(\sigma^\rho + 1)^{2-1/\rho}$, σ' goes to $+\infty$ whenever $\sigma > 1$. However, observe that for $\rho \geq 1/2$,

$$\frac{\sigma^\rho}{\left(1 - \frac{\alpha}{p}\right)^\rho} \geq \frac{\sigma^\rho(1 - \sigma^{\rho-1})}{\sigma^{2\rho-1} + 1 - \frac{\alpha}{p}(\sigma^\rho + 1)^{2-1/\rho}}.$$

To see this, consider the function $x \mapsto \sigma^\rho(\sigma^{2\rho-1} + 1 - x(\sigma^\rho + 1)^{2-1/\rho}) - \sigma^\rho(1 - \sigma^{\rho-1})(1-x)^\rho$, which has a minimum on $[0, 1]$ that is positive whenever $\rho \geq 1/2$. Given this, and since the right-hand side of (6) is bounded by $-2/((1-p)(1-\alpha/p)^\rho)$, the solution to (6) must lie below the solution of the differential equation

$$\frac{(1-\rho)\sigma'}{\sigma} = -2 \frac{1}{1-p} \frac{1}{\left(1 - \frac{\alpha}{p}\right)^\rho},$$

with the same initial condition at \bar{p} . This differential equation is easy to integrate, and since $\rho < 1$, its solution is finite at $p = \alpha$. Furthermore, either σ is finite, or if it diverges, then $\alpha/p = 1$ when $\sigma \rightarrow \infty$, but then the previous argument applies. Note however that $\alpha/p \rightarrow 1$ requires

$$\frac{\sigma^{2\rho-1} + 1}{(\sigma^\rho + 1)^{2-1/\rho}} \rightarrow 1$$

and so $\rho > 1/2$.

(Experimentation in infinite time.) We know that $\sigma' \rightarrow \infty$ as $u_2 \rightarrow 0$. Now differentiate the identity $p(t(p)) = p$, and obtain the following ODE for the function $t(p)$:

$$t'(p) = -\frac{1}{p(1-p)u_2(\sigma_1^\rho + 1)^{1/\rho}}.$$

Then compare $t(q)$ with $-\ln(p - \bar{p})$. For p close to $\bar{p} := \alpha(\sigma^\rho + 1)^{2-1/\rho} / (\sigma^{2\rho-1} + 1)$, we would like to have $t'(p) < -1/(p - \bar{p})$, so that $t(p)$ (which is decreasing) is steeper than $-\ln(p - \bar{p})$.

We then require

$$p(1-p)u_2(\sigma_1^\rho + 1)^{1/\rho} < (p - \bar{p}).$$

But this is the case, since σ stays bounded when $\rho > 1/2$, and r is small enough. We have:

$$\begin{aligned} & p(1-p)u_2(\sigma_1^\rho + 1)^{1/\rho} - (p - \bar{p}) \\ = & p(1-p)\frac{r}{2\alpha\sigma_1^\rho} \left(\sigma_1^{2\rho-1} + 1 - \frac{\alpha}{p}(\sigma_1^\rho + 1)^{2-\frac{1}{\rho}} \right) (\sigma_1^\rho + 1)^{1/\rho} - \left(p - \frac{\alpha(\sigma^\rho + 1)^{2-1/\rho}}{(\sigma^{2\rho-1} + 1)} \right) \\ = & \left(p(\sigma_1^{2\rho-1} + 1) - \alpha(\sigma_1^\rho + 1)^{2-\frac{1}{\rho}} \right) \left(r \frac{(1-p)(\sigma_1^\rho + 1)^{1/\rho}(\sigma^{2\rho-1} + 1)}{2\alpha\sigma_1^\rho} - 1 \right) < 0. \end{aligned}$$

The case $\rho = 1/2$: Consider equation (6) again: let $y := \sigma^{1-\rho}$, so that, when $\rho = 1/2$, we have:

$$\frac{(1-y)\alpha}{(y+1)(2p-\alpha)(1-p)} = \frac{y'}{y}.$$

Let $y_0 > 1$ and define the function:

$$g(p) := \frac{(y_0 - 1)^2}{2y_0} \left(\frac{\bar{p} - \frac{1}{2}\alpha}{1 - \bar{p}} \frac{1-p}{p - \frac{1}{2}\alpha} \right)^{\frac{\alpha}{2-\alpha}}.$$

(Equilibrium.) The exact solution is then given by

$$\begin{aligned} y(p) &= 1 + g(p) + \sqrt{g(p)^2 + 2g(p)} \\ \Rightarrow \sigma(p) &= \left(1 + g(p) + \sqrt{g(p)^2 + 2g(p)} \right)^2, \end{aligned}$$

and hence $g(p) \rightarrow \infty$ and $\sigma \rightarrow \infty$ as $p \rightarrow \alpha$. We can derive the equilibrium levels of effort from equation (5), and obtain that:

$$f(u_1(p), u_2(p)) = 2r \left(\frac{p - \alpha/2}{p\alpha} \right) (g(p) + 2).$$

(Experimentation in finite time.) Finally, by setting $K := \frac{(y_0-1)^2}{y_0} \left(\frac{\bar{p}-\frac{1}{2}\alpha}{1-\bar{p}} \right)^{\frac{\alpha}{2-\alpha}}$, we can show that the solution to $p' = -p(1-p)f(u_1(p), u_2(p))$ lies below the solution to

$$p' = -(1-\bar{p})(p-\alpha/2) \left(\left(\frac{1-\bar{p}}{p-\frac{1}{2}\alpha} \right)^{\frac{\alpha}{2-\alpha}} + \frac{2}{K} \right) \frac{2rK}{\alpha},$$

which converges to $\alpha/2$ in finite time. So experimentation stops in finite time when $\rho = 1/2$.

The case $\rho < 1/2$: Recall that

$$\left(1 - \frac{2 \left(1 - \frac{\alpha}{p} (\sigma^\rho + 1)^{1-1/\rho} \right)}{\sigma^{2\rho-1} + 1 - \frac{\alpha}{p} (\sigma^\rho + 1)^{2-1/\rho}} \right) \frac{1}{1-p} = \frac{(1-\rho)\sigma'}{\sigma},$$

which is equivalent to, considering the inverse function $p(\sigma)$,

$$p'(\sigma) = \frac{(1-\rho)(1-p(\sigma))}{\sigma} \left(1 - \frac{2(1+\sigma^\rho)}{\frac{p}{\alpha}(1+\sigma^\rho)^{1/\rho}(\sigma^{2\rho-1}+1) - (\sigma^{2\rho}-1)} \right).$$

It is immediate to verify that $p'(\sigma)|_{p/\alpha=g(\sigma)} = 0$, while $g'(\sigma) < 0$ for $\rho < 1/2$. This implies that $\sigma(p)$ cannot converge to a finite value, but that it must diverge to infinity. Also, the second term in brackets converges to zero unless $p \rightarrow 0$, since

$$\lim_{\sigma \rightarrow \infty} \frac{1+\sigma^\rho}{\sigma^{2\rho}-1} = 0, \text{ and } \lim_{\sigma \rightarrow \infty} \frac{(1+\sigma^\rho)^{1/\rho}(\sigma^{2\rho-1}+1)}{\sigma^{2\rho}-1} = +\infty.$$

So if p does not converge to 0, $p'(\sigma)$ is eventually positive, which is impossible. So p must converge to zero. But exerting effort for beliefs arbitrarily close to zero yields strictly negative profits, as even the team would not exert effort for sufficiently low beliefs. So such an equilibrium cannot exist. \square

Remarks on the case $\alpha_1 \neq \alpha_2$: Back to the case, $\rho > 1/2$, we know analyze the equilibria with

asymmetric players. Equations (5) and (6) may now be written as

$$\begin{aligned} \frac{1}{1-p} \left(1 - \frac{2 \left(1 - \frac{\alpha_2}{p} (\sigma_1^\rho + 1)^{1-1/\rho} \right)}{\frac{\alpha_2}{\alpha_1} \sigma_1^{2\rho-1} + 1 - \frac{\alpha_2}{p} (\sigma_1^\rho + 1)^{2-1/\rho}} \right) &= (1-\rho) \frac{\sigma_1'}{\sigma_1}, \\ \frac{r}{2\alpha_2 \sigma_1^\rho} \left(\frac{\alpha_2}{\alpha_1} \sigma_1^{2\rho-1} + 1 - \frac{\alpha_2}{p} (\sigma_1^\rho + 1)^{2-\frac{1}{\rho}} \right) &= u_2. \end{aligned}$$

It is easy to show that when $u_2 > 0$ and $\sigma_1 = 1$, then $\sigma_1'(p)$ and $(\alpha_2 - \alpha_1)$ have the same sign. Let $\alpha_2 > \alpha_1$ so player 1 is the more efficient agent. Since $\dot{p}_t < 0$, the equilibrium can never move from a scenario with $\sigma_1 < 1$ to one with $\sigma_1 > 1$. Therefore, if agents take turns, the more efficient agent must exert higher effort first (and for σ_0 close enough to one, this happens indeed). Both agents stop working when

$$p = \frac{\alpha_2 (\sigma_1^\rho + 1)^{2-\frac{1}{\rho}}}{\frac{\alpha_2}{\alpha_1} \sigma_1^{2\rho-1} + 1}.$$

This expression is minimized at $\sigma^* = (\alpha_2/\alpha_1)^{\frac{1}{1-\rho}} > 1$, so the equilibrium with the highest experimentation level is one in which players never take turns. The corresponding threshold for beliefs is

$$p^* = \alpha_2 \left((\alpha_2/\alpha_1)^{\rho/(1-\rho)} + 1 \right)^{1-1/\rho}.$$

This threshold will be reached in infinite time, since the corresponding terminal value for σ_1 is finite (see the proof of Theorem 6). Finally, notice that the value of p^* is increasing in α_2 and equal to $\alpha 2^{1-\frac{1}{\rho}}$ if $\alpha_1 = \alpha_2 = \alpha$.

Proof of Theorem 8: Applying integration by parts and ignoring constant terms, the payoff that agent i maximizes is equal to

$$- \int_0^\infty (r + \alpha^i u_t^i) \frac{e^{-rt}}{p_t^3} dt \text{ such that } \dot{p}^k/p^k = \sum_{j=0}^3 p^j u_j - u_k, \quad k = 0, \dots, 3.$$

Defining x_i such that $p_i = p_3 e^{-x_i}$, for $i = 1, 2$, note that $\dot{x}_i = u_i$. Defining also $q := p_3^{-1}$, the

problem reduces to maximizing

$$- \int_0^{\infty} (r + \alpha u_i) q e^{-rt} dt \text{ such that } \dot{q} = (1 - q)(u_1 + u_2) + u_1 e^{-x_2} + u_2 e^{-x_1}, \dot{x}_i = u_i, i = 1, 2.$$

Let μ_i denote the co-state variable associated with x_i . Pontryagin's principle gives

$$\alpha q e^{-rt} = \gamma (1 - q + e^{-x_j}) + \mu_i, \dot{\gamma} = (r + \alpha u_i) e^{-rt} + \gamma (u_1 + u_2), \dot{\mu}_i = \gamma u_j e^{-x_i},$$

or, equivalently, if we let $\sigma = \gamma e^{-x_1 - x_2}$,

$$\alpha q e^{-rt} = \sigma e^{x_1 + x_2} (1 - q + e^{-x_j}) + \mu_i, \dot{\sigma} = (r + \alpha u_i) e^{-rt - x_1 - x_2}, \dot{\mu}_i = \sigma u_j e^{x_j}.$$

Since we are focusing on a symmetric solution, this means

$$\alpha q e^{-rt} = \sigma e^{2x} (1 - q + e^{-x}) + \mu, \dot{\sigma} = (r + \alpha u) e^{-rt - 2x}, \dot{\mu} = \sigma u e^x, \dot{q} = 2u (1 - q + e^{-x}), \dot{x} = u.$$

Differentiate $\alpha q e^{-rt} = \sigma e^{2x} (1 - q + e^{-x}) + \mu$ and substitute for $\dot{\sigma}$ and $\dot{\mu}$ to get

$$\alpha (\dot{q} - rq) e^{-rt} - (r + \alpha u) e^{-rt} (1 - q + e^{-x}) = 2u \sigma e^{2x} (1 - q + e^{-x}) - 2u (1 - q + e^{-x}) \sigma e^{2x} = 0,$$

so that

$$q = \frac{\alpha u - r}{\alpha u - (1 - \alpha) r} (1 + e^{-x}).$$

So we are left with the system

$$q = \frac{\alpha u - r}{\alpha u - (1 - \alpha) r} (1 + e^{-x}), \dot{q} = 2u (1 - q + e^{-x}), \dot{x} = u.$$

We make the following change of variable. Let $v(q) = x(t)$, so $v'(q) \dot{q} = \dot{x}$, or $2v'(q) (1 - q + e^{-v(q)}) = 1$, whose positive solution is

$$v(q) = \ln \left(\frac{1 + \sqrt{1 + C(q-1)}}{q-1} \right).$$

Since we have $v(q(0)) = x(0) = \ln(p_3(0)/p_1(0))$, we can solve for C to get $C = \frac{\bar{p}^0 \bar{p}^3}{\bar{p}^1 \bar{p}^2}$. It follows that

$$u(q) = \frac{r}{\alpha} \left(1 - \frac{\alpha q}{q-1} (1 + (1 + C(q-1))^{-1/2}) \right),$$

and, since $t'(q) = v'(q)/u$, we also get that $t(q)$, the time at which the (inverse) belief is q , is given by $t(q) = F(q) - F(q_0)$, where $q_0 = 1/\bar{p}^3$, and

$$F(q) = \frac{\alpha}{2r(\alpha-1)} \left(\ln(C(1-q(1-\alpha)) - 1 + \sqrt{1+C(q-1)}) - \frac{2 \arctan \left(\frac{2(\alpha-1)\sqrt{1+C(q-1)}-1}{\sqrt{4\alpha(1-\alpha)(1-C)-1}} \right)}{\sqrt{4\alpha(1-\alpha)(1-C)+1}} \right).$$

Observe that, as $q \rightarrow \infty$, $u = r(1-\alpha)/\alpha > 0$, while it is clearly negative for $q \downarrow 1$. So experimentation stops at some belief, although we may only reach this belief asymptotically. Let $s := \sqrt{1+C(q-1)}$. Solving for the (larger) root of $u(q) = 0$ gives

$$q^* := \frac{2(1-C)}{1-2C-2\alpha(1-C) + \sqrt{1-4\alpha(1-\alpha)(1-C)}},$$

and \tilde{p}^3 , as defined in the text, is the reciprocal of q^* . It follows from the explicit solution for F that $\lim_{q \rightarrow q^*} t(q) = \infty$, i.e. experimentation never stops. This characterizes the unique candidate for an interior, symmetric solution, and it is easy to verify that, for low enough discounting (more precisely, whenever $u(q_0)$, as given above, is less than 1), agents cannot exert effort at a maximal rate over some interval of time in equilibrium. Sufficiency follows from the linearity and concavity properties of the objective, as in the baseline model. \square

The case in which $\bar{p}^1 \neq \bar{p}^2$: (Equilibrium with asymmetric players.) Suppose that $\bar{p}_1 > \bar{p}_2$. This means $x_1(0) < x_2(0)$. In the unique equilibrium, player 1 exerts effort at the maximal rate until $x_1(t) = x_2(t)$. From that point on, both agents work symmetrically. Clearly, the second phase does not take place if p_3 reaches \tilde{p}_3 before $x_1 = x_2$. It is immediate to see that this is an equilibrium. While player 1 works, player 2 prefers to wait. When exerting effort at a level that is interior, the two players play the mixed strategy equilibrium described in Theorem 5. Finally, player 1 has incentives to work alone until the time at which $x_1 = x_2$, since player 2 is

not exerting effort, and player 1 expects to work even after that time.

(Uniqueness.) For the uniqueness part, we repeatedly use the following claim: there cannot be a last person working alone. While a detailed analysis is omitted, this result is intuitive: if there exists a last agent i working alone, then in the last instants t in which player j is required to work, he has an incentive to deviate and shirk. Since player i will start working at $t + dt$ anyway, the gains in saved effort exceed the (vanishing) loss due to delayed arrival of a success. This result rules out cases in which players work sequentially. Suppose that player 1 worked until his individual threshold p_1^* . If the asymmetry in \bar{p}_i is sufficiently small, player 2 would then start working from there, because we would have $p_2 > p_1^*$. But then, anticipating that player 2 will start working at the maximal rate, player 1 wants to deviate and shirk in the last instants before his beliefs reach the threshold. A qualitatively identical scenario arises if player 2 works alone until the individual threshold p_2^* . We now analyze the candidate interior equilibrium, and rule it out based on the same claim. Consider again the system of equations:

$$\begin{aligned}\alpha q e^{-rt} &= \sigma e^{x_1+x_2} (1 - q + e^{-x_j}) + \mu_i, \quad \dot{\sigma} = (r + \alpha u_i) e^{-rt-x_1-x_2}, \quad \dot{\mu}_i = \sigma u_j e^{x_j}, \\ \dot{q} &= (1 - q) (u_1 + u_2) + u_1 e^{-x_2} + u_2 e^{-x_1}, \quad \dot{x}_i = u_i.\end{aligned}$$

As before, differentiate the first equation, and plug in the formulas for $\dot{\sigma}, \dot{\mu}_i, \dot{q}$. Upon simplification, we obtain the following relation between u_j and x_i, x_j, q :

$$u_i = \frac{r}{\alpha} \frac{1 - q + e^{-x_i} + \alpha q}{1 - q + e^{-x_j}}. \quad (6)$$

Given that $\dot{x}_i = u_i$, we can also solve the equation

$$\dot{q} = (1 - q) (u_1 + u_2) + u_1 e^{-x_2} + u_2 e^{-x_1}, \quad (7)$$

and obtain as a general solution

$$q = 1 - k e^{-x_1-x_2} + e^{-x_1} + e^{-x_2}, \quad (8)$$

for some constant k . As before, we define $v_i(q) := x_i(t)$, so that

$$v'_i(q) \dot{q} = \dot{x}_i = u_i.$$

Substituting (6) in (7), we obtain the following first-order differential equation, separately for each agent's effort:

$$(2(1-q) + e^{-v_1} + e^{-v_2} + 2\alpha q) v'_i = \frac{1-q + e^{-v_i} + \alpha q}{1-q + e^{-v_j}}.$$

Finally, using (8), we obtain

$$\left(1 - q + e^{-v_i} + e^{-v_i} \frac{(1-q)k+1}{ke^{-v_i}-1} + 2\alpha q\right) v'_i = \frac{1-q + e^{-v_i} + \alpha q}{e^{-v_i} \frac{(1-q)k+1}{ke^{-v_i}-1}},$$

which is the desired ODE characterizing $v_i(q)$. Observe that the ODEs for $v_1(q)$ and $v_2(q)$ differ only because of the initial condition. In particular, $v_i(q(0)) = x_i(0)$, implying that $v_1(q(0)) < v_2(q(0))$. Since the paths of the solutions to the two ODEs cannot cross, $v_1(q)$ will reach zero for a level \tilde{q} for which $v_2(\tilde{q}) > 0$. The fact that the weaker player works harder is clearly necessary in order to have an interior (i.e. mixed strategy) equilibrium. When player 1 stops working, however, player 2 should continue exerting effort at the maximal rate, and he will be the last player working, which we ruled out. This rules out all but the equilibrium we described earlier.

Interpretation of \tilde{p}^3 in terms of the team problem: We prove here that the threshold \tilde{p}^3 is also the threshold at which the team would stop, if the value was $1/2$ per agent. In the team problem, agents would choose to allocate efforts equally (if $p^1 = p^2$), and they would exert effort at the maximal rate. That is, they would choose a time T to maximize

$$-\int_0^T (r+2\alpha)q_t e^{-rt} dt - \int_T^\infty r q_T e^{-rt} dt = -\int_0^T (r+2\alpha)q_t e^{-rt} dt - q_T e^{-rT}.$$

The optimal time then satisfies (taking first-order conditions with respect to T)

$$-(r + 2\alpha)q_T e^{-rT} + r q_T e^{-rT} - \dot{q}_T e^{-rT} = 0, \text{ or } \frac{\dot{q}_T}{q_T} = -2\alpha.$$

Given that q , as defined in the proof of Theorem 5, satisfies $\dot{q} = (1 - q)(u_1 + u_2)$, this means that

$$q_T = \frac{1 + \frac{\bar{p}^1}{\bar{p}^3} e^{-T}}{1 - \alpha}.$$

The solution to $\dot{q} = 2 \left(1 - q + \frac{\bar{p}^1}{\bar{p}^3} e^{-t}\right)$ is $1 - q_t = k e^{-2t} - 2 \frac{\bar{p}^1}{\bar{p}^3} e^{-t}$, with $k := -\frac{\bar{p}^0}{\bar{p}^3}$. This gives

$$-\alpha + (1 - \alpha) \frac{\bar{p}^0 \bar{p}^3}{(\bar{p}^1)^2} z^2 + (1 - 2\alpha) z = 0,$$

for $z := \frac{\bar{p}^1}{\bar{p}^3} e^{-T}$. That is, $z = \frac{-(1-2\alpha) + \sqrt{(1-2\alpha)^2 + 4\alpha(1-\alpha) \frac{\bar{p}^0 \bar{p}^3}{(\bar{p}^1)^2}}}{2(1-\alpha) \frac{\bar{p}^0 \bar{p}^3}{(\bar{p}^1)^2}}$, and therefore

$$q_T = \frac{2(1 - C)}{1 - 2C - 2\alpha(1 - C) + \sqrt{1 - 4\alpha(1 - \alpha)(1 - C)}} = 1/\tilde{p}^3.$$

Therefore, effort stops when the belief reaches the same threshold in both problems. \square