

Collaborating*

Alessandro Bonatti and Johannes Hörner

February 15, 2009

Abstract

This paper examines how moral hazard in teams evolves. Agents are collectively engaged in an uncertain project, and their individual efforts are unobserved. Free-riding leads to postponement of effort. 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 ever, increases with the number of agents. We show how synergies affect these conclusions, and why deadlines are beneficial.

1 Introduction

Cooperation evolves over time. Lack of tangible results often breeds mistrust, and mistrust leads to lower levels of commitment. Agents grow suspicious that other team members do not pull their weight in the common enterprise, and scale back their own involvement in response.

Is this bleak scenario indeed the lot of every team project? How does the evolution of free-riding depend on the characteristics of the project and of the team? And what can be done to avert such a scenario? This paper develops a formal framework to address these questions.

A set of agents is engaged in a project. They share a common value from completing the project. All it takes to complete the project is one breakthrough, but making a breakthrough

*First draft. We would like to thank Dirk Bergemann, Martin Cripps, David Miller, Sven Rady, Roland Strausz and Xavier Vives for useful discussions, and Nicolas Klein for excellent comments.

requires costly effort. Some projects, however, are inherently doomed to failure, no matter how much effort is put into them. As for the other kind of projects, the instantaneous probability of a breakthrough increases in the combined effort of the agents at this moment. Making a breakthrough is the only way to find out that a project is good.

Unlike in the literature on experimentation, 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 less optimistic about the prospects of the project as time goes on. Agents, thus, have the usual incentive to free-ride on the effort of others.

Our focus is descriptive. There is a large and important literature on how free-riding affects the amount of effort that is provided in teams. Here instead, we focus on how moral hazard affects the allocation of effort over time, and what factors determine this allocation. This, in turn, will provide us with an understanding of the circumstances under which such alliances, or teams, are likely to form in the first place. This allows us to evaluate the efficiency of intrinsically dynamic mechanisms, such as deadlines.

Our main results are as follows. When efforts are perfect substitutes, free-riding comes in the form of a strategy of procrastination and postponement. First, the total effort exerted altogether decreases, as it is bounded away from the cooperative amount. But effort is not only scaled down, it is delayed. Agents keep exerting some effort as long the common belief exceeds the Marshallian threshold, but the times at which intermediate beliefs are reached move back, and individual effort decays sufficiently fast for this threshold never to be reached. Work dwindles over time, but the plug on the project is never pulled. Therefore, the more agents involved in the project, the later (stochastically) the project gets completed, if ever. Nevertheless, the individual payoff increases with the number of participants.

If agents are not identical, there is a unique equilibrium. The more capable agent behaves as if he were alone, and the other agents exert no effort whatsoever. (This is also one of the asymmetric equilibria in the symmetric game.) Of course, while this is an efficient outcome, it is also a very unbalanced one, and if joining the alliance involves any additional fixed set-up cost, the strongest agent will be unwilling to participate. Therefore, only alliances among similar agents should be expected to form.

How can the unfortunate outcome in the symmetric set-up be remedied? We show that, if agents have enough commitment power to fix themselves a deadline, it is indeed optimal to do so. Faced with a deadline, an agent's equilibrium effort is U-shaped. It is low and decreasing over time at the beginning, but once the deadline looms close enough, effort jumps back to its maximum. The downside of a deadline is that, in equilibrium, agents pace themselves so that, if the deadline is hit, the project is abandoned at a point at which the belief still exceeds the threshold at which they would stop if unconstrained. If agents could re-negotiate at this time, they would. Nevertheless, a deadline is desirable, as we show that the reduction in wasteful delay more than compensates the shortfall from this event. The optimal deadline is precisely the longest time such that agents have incentives to exert maximal effort throughout. We also examine the optimal dynamic wage scheme for a principal who owns the project's returns.

It is often argued that observability, or monitoring, may mitigate the free-riding issue. We show that this is not necessarily the case. In fact, in the unique Markovian equilibrium, effort is even lower under observability than non-observability. This is because individual efforts are strategic substitutes. The prospects of the team improve if it turns out that an agent has slacked off, since 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. There are, of course, other, non-Markovian equilibria, some of which are even efficient.

As Alchian and Demsetz already pointed out, synergies are likely explanations for team formations and alliances. That is, total output might not be separable in the agents' efforts. To account for this phenomenon, we extend our model into two directions. First, we consider the case in which agents possess different skills, so that it might be that case that both, none, or one or the other worker is able to succeed. Alternatively, 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 is a C.E.S. function of their individual efforts.

These two kinds of synergies lead to very different behaviors. In the first case, there is a unique equilibrium. The agent that is viewed as most likely to succeed starts by exerting effort by himself, at the efficient level. As time passes by and no breakthrough occurs, the difference in the likelihoods of success across agents level off. When both agents are equally likely to succeed,

they then both start exerting effort, but these effort levels are low because of free-riding. Unless a breakthrough occurs, both agents will keep on exerting effort forever, albeit at rapidly declining levels. In the second case, there are multiple equilibria, which are all inefficient. In all equilibria, both agents put in some effort. These equilibria illustrate the difference between the amount of effort, and the allocation of effort: the symmetric equilibrium is characterized by free-riding and significant delay, but it also the one in which the probability of a breakthrough is greatest.

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 additive tasks, in which the payoffs is additively separable in the tasks, conjunctive tasks, in which both tasks must be completed, and 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. 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 (2006) and Murto and Välimäki (2008) consider the case in which the outcomes of each agent's action is unobservable, but their actions are. This is precisely the opposite of what is assumed in this paper, since 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.

This paper is also related to 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).

Dynamic versions of public-goods game, with observable contributions, are examined in Admati and Perry (1991), Compte and Jehiel (2004), Lockwood and Thomas (2002), and Marx and Matthews (2000). 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, or 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 evolution of project teams.

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 among 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 during that time would be exponentially distributed 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. 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 a 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 a little more precise, a strategy for agent i is a measurable function $u_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, with the interpretation that $u_{i,t}$ is the instantaneous effort exerted by agent i at time t , conditional on no breakthrough having occurred. Here and in what follows, we omit the specification of the strategy after histories following an agent's own deviations. Given a strategy profile $u := (u_1, \dots, u_n)$, it follows from Bayes' rule that the common belief assigned by the agent to the project being good, p , is given by the solution to the differential equation

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

with $p_0 = \bar{p}$. Since the project is good with probability p_t , and the 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 $p_t f(u_t) - c_i(u_{i,t})$. Since the probability that a breakthrough occurs before time t is given by $\int_0^t p_s f(u_s) ds$, it follows that the average (expected) payoff that agent i seeks to maximize is equal to

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

Since there is a positive probability that the game lasts forever, and since 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, and so our objective is to identify the Nash equilibria of this game. (We shall nevertheless describe the sequentially rational behavior off the equilibrium path 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:

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

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

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

subject to

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

Observe that the parameter α is the Marshallian threshold. That is, it is equal to the belief threshold at which a myopic agent would stop experimenting, since at this point the instantaneous marginal revenue from effort, p_t , equals 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

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

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

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

Not surprisingly, the resulting payoff is decreasing in the discount rate r and the normalized cost α , and increasing in the prior 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 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 , is the well-known free-riding effect in teams, eloquently described in Alchian and Demsetz (1972), and extensively studied thereafter. In a non-cooperative equilibrium, the amount of experimentation is too low.¹ Here instead, our focus is on how free-riding affects the timing of the provision of effort.

¹This 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.

3.2 The Non-Cooperative Solution

As mentioned, 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 otherwise. Further, we shall assume throughout this section that agents are sufficiently patient. More precisely, we assume that r satisfies

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

The case of agents that are sufficiently impatient for this assumption to fail is discussed at the end of this section.

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.²

What is the trade-off between exerting effort at some instant and the next? Consider some time t and assume that players 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

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

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

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

²Applying dynamic programming more generally is possible, but difficult in this setting, since 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. Of course, verifying optimality of some strategy profile does require evaluating the continuation payoff after such deviations.

The first-order effect of effort today and 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 pin down the effort level. To do so, we must consider the next term in the Taylor expansion, and use the fact that $p' = p - p(1 - p)(u_i + u_{-i})dt$.³ Then

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

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 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 to save 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,

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

where the first three second-order terms have analogous interpretations. The penultimate term corresponds to the saving in marginal cost from exerting effort later: since with probability $p(u_i + u_{-i})dt$, that next instant is not reached, the cost of effort is spared in the 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

³We must also not forget that $e^{-p(u_i + u_{-i})dt} = 1 - p(u_i + u_{-i})dt + \frac{p^2}{2}(u_i + u_{-i})^2(dt)^2$.

procrastinating does not strictly dominate exerting effort now, it must be that

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

But 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 $p \rightarrow \alpha$. So it must be that effort tends 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 effort level. By planning to work a bit more tomorrow instead if necessary, the saving is proportional to the combined effort levels, that determine whether it will indeed be necessary to produce that effort. As a result, agents procrastinate.

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

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

which is clearly impossible, as dt is arbitrarily small. Therefore, not only does effort go to zero as p tends to α , but it goes 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 equilibrium value of u_i must be. Namely, since agent i must be indifferent between exerting effort at different moments in time, it must be that

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

From this, it follows that the common belief tends to the Marshallian threshold α asymptotically, and that total effort, as a function of the belief, is actually decreasing in n . Of course, the belief p is itself a function of the path of effort. The following theorem, proved in appendix, provides the closed-form solution.⁴

⁴In the case $\bar{p} = 1$ that has been ruled out earlier, the game reduces essentially to the static game. Effort is constant and, because of free-riding, inefficiently low ($u_{i,t} = \frac{r(\alpha^{-1} - 1)}{n - 1}$).

Theorem 1 *There exists a unique symmetric equilibrium, in which the effort level 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.

Lemma 1

1. *Individual effort is decreasing in t . For a fixed time t , individual effort is decreasing in r and α , and increasing in p .*
2. *As a function of t , individual and total effort are first decreasing, and then increasing in n .*
3. *The payoff $V_i(\bar{p})$ is increasing in n and $\bar{p}$, and decreasing in α . It is independent of r .*
4. *The probability of a breakthrough occurring is independent of n . Conditional on a breakthrough occurring, its expected time increases in the number of agents.*

As discussed above, total 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. But since this means that the belief does not decrease as fast with more agents, and since effort increases in the belief, it must be that eventually, total 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 overall effort exerted over the infinite horizon is independent of n as well. Ultimately, then, larger teams must catch up in terms of effort.

The last conclusion is an immediate consequence of this temporal pattern. Slightly surprisingly also, the payoff is independent of the discount rate, which, as it turns out, is simply a scaling factor. For instance, the expected cost of delay until the random time τ of a breakthrough, conditional on a breakthrough occurring eventually, $1 - \mathbb{E}[e^{-r\tau} | \tau < \infty]$, is independent of the discount rate. While larger teams are slower to succeed, their members are better off, as the payoff is increasing in n . This result is obvious for $n = 1$ relative to $n = 2$, since an agent may always choose to behave as if he were by himself, securing the payoff from a single-agent




Figure 1: Individual and collective effort




Figure 2: Payoffs and cost of delay

team. The result is far less obvious for larger teams. Figures 1 and 2 provide a quantitative illustration of these results.

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 what it would have been on the equilibrium path. This means that agent i 's private belief is higher than the common belief of the other agents. As agent i would precisely be indifferent between exerting effort or not if he shared the common belief, his optimism leads him to optimally 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 higher, 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, if ever.

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 in the cooperative solution. That is, he would exert maximal effort up to time T_1 at which his belief attains α . Faced with such behavior, the best-reply of another 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 appendix, we prove that this is precisely the set of all equilibria for $n = 2$ (as long as $\lambda/r \geq \alpha^{-1} - p^{-1}$). More generally, with n agents, an equilibrium is indexed by a collection of nested subsets of agents, $\{i\} \subset \{i, j\} \subset \{i, j, k\} \subset \dots \subset \{1, \dots, n\}$, and (not necessarily distinct) times $t_1 \leq t_2 \leq \dots \leq t_{n-1}$, with $t_1 \leq T_1$, $t_k \in \mathbb{R}_+ \cup \{\infty\}$ for $k \geq 2$ and $t_{n-1} = \infty$, 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 $nu_i < \lambda$), 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. According to the maximin rule, however, the ranking is reversed, as 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 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 on patience was necessary to make sure that the agent's individual effort characterized in Theorem 1 was indeed 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 has then the feature that all agents exert maximal effort up to the time t at which, given the resulting belief, the effort level $u_i(p)$ given above is exactly 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 level $u_i(p)$, and the qualitative features of the equilibrium are as before.

So far, agents were assumed 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 u_{-i} is the smaller, the more patient agent 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 appendix. The intuition is straightforward. Since 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.⁵ 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 since agent i will be starting to exert maximum effort in an instant, the probability that he might spare this effort altogether is high enough for him to defer. This implies that $t = 0$ and so, that agent i is indeed 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.

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

In this subsection, the previous findings are contrasted 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 effort choice at time t 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 known among agents, even after histories off the equilibrium path.

A strategy for agent i , then, is a map $u_i : [0, 1] \rightarrow [0, 1]$ from possible beliefs p into an effort choice $u_i(p)$, such that (i) u_i is left-continuous; (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

⁵More precisely, it cannot be that $\lim_t p_t > \alpha_i$, since, 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 .

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, the continuation payoff given p , $V_i(p)$, must satisfy the optimality equation given by, for all p , and $dt > 0$,

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

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. Since 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

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

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

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

It is standard to verify that the resulting u is the unique equilibrium strategy profile provided that $\bar{p}$ is such that $u \leq 1$ for all $p < \bar{p}$. In particular, this is satisfied when, as assumed in the unobservable case, $\lambda/r \geq \alpha^{-1} - p^{-1}$, which we maintain henceforth. In the model without observability, recall that, in terms of the belief p , the effort is given by $u(p) = \frac{r}{n-1}(\alpha^{-1} - p^{-1})$.

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

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

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

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 $(1-\alpha)^{-1} > 0$. So 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, the equilibrium level of effort is lower when previous choices are observable, and so is the welfare. While this may sound a little surprising, it is an immediate consequence of the fact that effort choices are strategic substitutes. Since effort is increasing in the common belief, and 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 higher level of effort by other agents. That is, to some extent, the other agents take up the slack. This depresses the incentives to provide effort, and leads to lower equilibrium levels. This, of course, cannot happen when effort is unobservable, since an agent cannot provoke the other agents into doing the effort for him. Figure 3 illustrates this relationship. As can be seen from the right panel, a lower effort level for every value of the belief p does not imply a lower effort level for every time t : since total effort over the infinite horizon is the same in both models, effort levels 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 directly seen from the formula for effort, as 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 giving 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. But p is not observed: it is the common belief about the unobserved past total efforts, including i 's effort, and therefore, it does depend 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.




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

There also other, asymmetric Markovian equilibria in the observable case similar to those described in Keller, Rady and Cripps (2005), involving possibly many switching points between the agents. 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 on some time interval starting at time t . Under observability, if agent i procrastinates, this ‘freezes’ the common belief, and therefore this 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, and so agent j will start exerting effort at time t , as expected. Further, since an agent working alone, such as j , must necessarily be exerting maximal effort (because of discounting), procrastination is profitable for i , as the induced probability that the postponed effort won't have to be carried out more than offsets the loss in value due to discounting. Therefore, such switching is impossible without observability, independently of the agents' discount rate.

Finally, there exist other, non-Markovian equilibria. As mentioned, appropriate concepts of equilibrium have been defined carefully elsewhere (see, for instance, Bergin and McLeod (1993)).

It is not hard to see how one can define formally a ‘grim-trigger’ equilibrium, for a low enough discount rate, in which all agents exert effort at 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 maximal effort from this point until the common belief reaches α . While this equilibrium is not first-best, it does clearly better than the Markovian equilibrium in the observable case, and than the symmetric equilibrium in the unobservable case.⁶

4 Deadlines and Optimal Incentive Schemes

In the absence of any kind of commitment, the equilibrium outcome described above seems inevitable. Pleas and appeals to cooperate are given no heed to, and deadlines are disregarded. In this section instead, we consider limited forms of commitment.

The first kind of commitment is a self-imposed time-limit. We assume that agents 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, and we examine which is the optimal horizon.

In the second case we consider, a principal owns the project and decides what the reward for the successful project completion should be. Because the identity of the actual agent responsible for the breakthrough is not observable, we restrict attention 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 one.

4.1 The Optimal Deadline

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

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

⁶It is tempting to consider grim-trigger strategy profiles in which agents exert effort for beliefs below α . We ignore them here, since such strategy profiles cannot be limits of equilibria of discretized versions of the game.

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 $T = \infty$ has been exhaustively studied. The next lemma, proved in appendix, describes the symmetric equilibrium for $T < \infty$.

Lemma 2 *Given $T < \infty$, there exists a unique symmetric equilibrium, characterized by $T_1 \in [0, T)$, in which the effort level 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 given by Theorem 1. The time T_1 is non-decreasing in 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 not to affect the agents' incentives. At some point, however, the deadline looms large above the agents. Agents pick it up and exert maximal effort from then on. But it is then too late to catch up with the overall effort exerted in the infinite-horizon case, as $p_T > \alpha$. By waiting

until time T_1 , agents take a chance. That p_T must exceed α is not hard to see: if the deadline was not binding, each agent would prefer to procrastinate at instant T_1 , given that all agents exert then maximal effort until the end.

The next theorem establishes that it is indeed 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 precisely set so that agents 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 maximal effort throughout.

Note that the deadline is decreasing in n , as it is the product of two positive and decreasing functions of n . That is, larger teams call for tighter deadlines. 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 check 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.

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 reported in the last section, the path of equilibrium effort appears to be continuous over time when the cost is quadratic. Nevertheless, a deadline provides additional incentives to exert effort when it is close enough, and effort is U-shaped in a variety of circumstances.⁷

4.2 The Optimal Incentive Scheme

In the last subsection, attention was focused on what the team could do to help itself. Here instead, 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

⁷Readers have pointed out to us that this effort pattern in presence of a deadline reminds them of their own behavior as a single author. Sadly, it appears that such behavior is hopelessly suboptimal for $n = 1$.

restrict attention to symmetric schemes. Since 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 $w : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, measurable, so as to maximize

$$r \int_0^\infty n p_t u_t (v - n w_t) e^{-\int_0^t (n p_s u_s + r) ds} dt,$$

given that $u := u_i : \mathbb{R}_+ \rightarrow [0, 1]$, measurable, maximizes

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

where $u_{-i} := (n - 1)u$. Indeed, observe that a breakthrough arrives at rate $u_i + u_{-i} = nu$, and that the total wage bill is nw_t . Since an agent can only be incentivized to exert effort if the marginal benefit, pw , covers the marginal cost, α , and since 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, as fixing a constant wage equal to the average of v/n and $\alpha/\bar{p}$ is one of the feasible schemes.

However, 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

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

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

$$w_t = w(\delta)e^{rt} + \alpha \frac{n-1+r}{r}(1-e^{rt}) - \alpha \frac{r(1-\bar{p})}{(n-r)\bar{p}}(e^{nt} - e^{rt}) \text{ for } t \leq T^*, \text{ and } w_t = 0 \text{ for } t > T^*,$$

which is such that $u_t = 1$ until time $T^ = \frac{1}{n} \ln \delta$, and $u_t = 0$ thereafter.*⁸

That is, it is optimal for the principal to induce maximal effort, as long as any effort is worthwhile. This requires that the wage dynamics offset the free-riding incentives. The principal must be giving some rents to the agents to induce them to work early. This is easier when agents are impatient (the wage decreases in r), and harder when there 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 compensates 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 $n = 1$).

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). Figure 5 illustrates how the wage scheme varies with the parameters. Effort stops at the time at which $p_t v_t = \alpha$, or $v_t = p_t / \alpha$. Note that this time is independent of the discount rate, as 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, this amount of effort varies with the size of the group. As is easy to check, δ is decreasing in n : there is more aggregate effort exerted in larger groups. However, the principal's

⁸Clearly, the choice of w_t for $t > T^*$ is to a large extent arbitrary.

payoff need not increase in n . Smaller groups are typically better, since it is less costly to overcome the incentive to free-ride.

While the previous scheme is optimal from the principal's point of view, it is inefficient. From the perspective of total welfare, it is best to have agents exert maximal effort all the way to the point at which $p_{T^*}v = \alpha$. This time is given by $T^* = \frac{1}{n} \ln \delta^*$, where

$$\delta^* = \frac{\bar{p}}{1 - \bar{p}}(v/\alpha - 1).$$

The corresponding wage dynamics are unchanged, provided $w(\delta^*)$ is substituted for $w(\delta)$.




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 in moral hazard in teams. Among the many kinds of possible synergies, we focus here on two extreme versions. In the first case, agents are identical, but working together is more effective than working separately. That is, the arrival rate displays complementarities in the agents' effort choices. In the second case, agents

have different skills, so that, depending on the type of project, one or the other skill might be appropriate.

As before, capacity λ is normalized to one.

5.1 Complementarities

In this subsection, we go back to the case in which the project is either good or bad, but assume that the instantaneous probability of a breakthrough, conditional on the project being good, is given by

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

That is, this function has the property of constant elasticity of substitution. 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 $f_i = \partial f / \partial u_i$.⁹ The baseline model corresponds to the special case $\rho = 1$. Agents are assumed identical for now ($\alpha_i = \alpha$). We assume that the discount rate satisfies

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

Because $\lim_{u_{-i} \rightarrow 0} f_i(u_i, u_{-i}) = \infty$ for all $u_i > 0$, it is no longer possible that, in equilibrium, some agents exert no effort whatsoever while some other agents exert positive 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 belief drops to $n^{-\frac{1}{\rho}}\alpha$. 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 maximal effort is only exerted up to the time at which the belief reaches α .

⁹Otherwise, there always exist trivial equilibria in which no agent exerts any effort.

In this case as well, it is more convenient to represent the equilibrium effort level in terms of the common equilibrium belief at that time. Recall however that effort is not observable, and that effort is a function of time only, but that the common belief is a function of time and effort, which can be derived from the equilibrium strategies.

Theorem 5 *There exists a unique symmetric equilibrium, in which effort 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 provision 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-decision problem. As mentioned, in the single-decision problem, effort is exerted up to the point at which the belief is equal to α . Since $\alpha > n^{1-\frac{1}{\rho}}\alpha$, aggregate equilibrium effort is higher with synergies, and indeed, 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 , so that, unlike in the case without synergies, effort is higher with more agents, provided the belief is low enough.

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, but as is easy to verify as in Theorem 2, specifies effort levels that are lower than in the non-observable case, for any given level of belief.

To discuss the asymmetric equilibria, it is simpler to assume that $n = 2$, an assumption which is maintained throughout this subsection. As mentioned, synergies ensure that there cannot be

an equilibrium in which some agent does not exert any effort at all, while his partner does. But 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 check 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 6 *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 appendix, we further show that 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: since agent 1 exerts more effort than agent 2 at the initial time, he keeps on exerting more effort throughout. Figure 6 displays how the the range of values of initial ratios $\sigma_0 := u_{1,0}/u_{2,0}$ that are consistent, for a given prior p_0 and a level of complementarity ρ , with an asymmetric equilibrium.

As in the baseline model, there is a trade-off between efficiency and fairness: it is always best for the team's aggregate payoff to have some asymmetry in their effort choices when synergies are not too strong, but, as simulations show, it appears that, as expected, the optimal degree of asymmetry decreases with the strength of the synergies.

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 prevails. Depending on the initial




Figure 6: Admissible initial values for asymmetric equilibria

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 less to lose by procrastinating (the project is worth less to him, and so discounting its value is less costly), and more to gain (effort being costlier for him, he appreciates more the potential cost saving from procrastination). Therefore, if agents must 'take turn,' 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 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. The two panels of Figure 7 illustrate how the respective efforts might vary with the initial value of σ .





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

5.2 Different Skills

The baseline model is generalized as follows. Instead of being simply good or bad, the project is of one of four possible types. Either 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 $u_{1,t} + u_{2,t}$. Or the project 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 $u_{-i,t}$. Finally, the project might be bad (type 3), and efforts are then wasted, as 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 $\bar{p}^3 > 0$.¹⁰ Let $u_0 := u_1 + u_2, u_3 := 0$. Using Bayes' rule, it is readily verified that, for $k = 0, \dots, 3$,

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

¹⁰The 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 an effort level equal to $r(\alpha^{-1} - (1/2)^{-1})$, which is the same effort as in the baseline model for a belief $p = 1/2$.

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 (α_i is independent of i), and we first focus 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 all the way. Indeed, if agent $i = 1, 2$ works by himself, p^i will decrease faster than p^j , where j is the index of the other agent. In fact, p^j increases. 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 any longer, 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 would be by himself, say at time t . But then the other agent would be unwilling to exert any effort at time $t - dt$, for $dt > 0$ small enough. So when 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 multi-dimensional. However, it turns out that it can be explicitly solved, 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 7 *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 effort equal to

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

The exact relationship between time t and belief p_t^3 is given in the proof. Effort is positive for all $t \geq 0$, and $\lim_t p_t^3 = \tilde{p}^3$.

Given this result, it appears as if 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 protracts tediously, and effort wastes away over time. Indeed, it follows from the theorem that $\lim_{t \rightarrow \infty} u_{i,t} = 0$. As $\bar{p}^1 = \bar{p}^2$ 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 $\bar{p}$ 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, as effort is not maximal, and $1 - p^3$ tends to a higher threshold, 2α , reflecting 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 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 faster than p^2), 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 5.

As shown in 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. It affects the total provision of effort in the usual way ($\tilde{p}^3$ falls short of

the cooperative threshold). And it also affects the timing of this provision, as described above. Figure 11 below illustrates the pattern of effort and belief over time.




Figure 8: Effort and beliefs with synergies

6 Multitask Projects

So far, a project has consisted of a single breakthrough. This is a gross simplification. Most projects involve several steps, or tasks, that 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. Also, given that the information structure is no longer trivial, we require strategies to be sequentially rational. At any point time, agents now observe whether a task has been already (successfully) completed, which task it was, and at what time it was completed. As before, the level of effort is not observed, and the allocation of effort is not observed either. If both agents work on one task simultaneously, and the task gets completed at that instant, we also assume that the specific

agent that is responsible for the success cannot be identified.¹¹

The literature on project design has emphasized the relevance of the type of task to the impact of social loafing. Conjunctive tasks are tasks that must all be completed for the project to be a success. Disjunctive tasks are tasks for which completing either one of them guarantees the success of the team, and completing any further task provides no further benefit. Finally, additive tasks are such that 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 one to each agent. If both are completed, it yields a payoff of two to each agent.

Another dimension along which projects vary relates to the timing of tasks. For some projects, it is possible to work on both 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 tasks is statistically independent across tasks, and the initial probability that any given task is good is still denoted $\bar{p}$. Effort is additively separable across tasks, and agents are assumed identical, with marginal cost α , and total capacity for effort normalized to one. The arguments showing that the ones we discuss are indeed equilibria are relatively straightforward. They are therefore omitted and available upon request.

6.1 The Team Problem

To better understand the agents' incentives, 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 maximal rate until some point at which effort stops altogether. Our focus is on the allocation of effort across tasks, and we thus omit a discussion of these stopping times. The next proposition describes the optimal sequencing.

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

(1) work sequentially on tasks if they are conjunctive. If both tasks have the same parameter $\bar{p}$, it is better to start with the task for which the parameter α is higher. If the parameter α is the

¹¹This 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.

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 exclusively devoted 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. Recall, however, that conjunctive tasks must all 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 this task until it is determined whether the harder task can be completed or not.¹²

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 is also the case in 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 one 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 providing all the effort.

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

¹²Obviously, the threshold for the first task in the sequential problem is not the threshold α of 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 payoff 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 producing 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 with the feature that the last task is performed by one agent only.

This reasoning can be extended to multiple tasks. With two tasks left, the agent performing 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 with the feature that the last task is performed by one agent only; in this equilibrium, agents alternate executing tasks, as long as they are successful.

6.3 Conjunctive Tasks

Now consider the case in which the two tasks are conjunctive, or complements (*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.¹³

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 successfully completed the first task.*
- (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.*

¹³The specification of the appropriate bound on r is omitted here.

In the equilibrium with sequential efforts, beliefs reach the efficient thresholds. In fact, the first agent works until her beliefs offset the expected payoff from having the second agent completing the remaining one. Once and if the first agent has completed the task, the second agent works until the beliefs reach α , since her value from a success is equal to one. Furthermore, both agents exert maximal effort, making this equilibrium the most efficient noncooperative solution.

The equilibrium with simultaneous work and specialization is supported by the threat that, in case an agent obtains a success after time T , she 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 (as both individual thresholds shift up through time). The continuation strategy of stopping to work after a success is efficient both from an ex post perspective (as having only agent work on one task is efficient), and from an ex ante perspective. It is key 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. Finally, note that 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, then 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 in mixed strategies. For instance, agents play the symmetric (baseline) mixed strategy equilibrium 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 projects 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 maximal effort until the single-task threshold $p = \alpha$ is reached. Upon obtaining a success, an agent stops working, while the other one completes her task.*

In the equilibrium in which agents work simultaneously, both agents work until they reach the single-task threshold, as the value of each success is independent of their beliefs about other tasks. This also deters an agent from delaying her efforts to after her partner has completed her task. Furthermore, 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 entails a longer time-to-completion.

There also exists a symmetric equilibrium in mixed strategies. In this case, agents can work on the two tasks both sequentially or simultaneously. In either case, agents adopt the mixed strategies described in our baseline equilibrium with one task. The equilibrium with sequential efforts requires a minimal level of patience to ensure that agents actually wish 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, meaning 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 mixed-strategy equilibria of the game with two tasks:

- (1) Agents work simultaneously, each on one task. Each agent exerts an amount of effort lower 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 in which agents work 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 devoted to each project 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 her effort, relatively to the equilibrium with division of labor.

7 Concluding Remarks

Convex Costs: Throughout the analysis, we have maintained the assumption that the cost is linear in the effort. While this affords tractability, it is natural to ask whether the findings are robust to this assumption. This is especially relevant, since, with linear cost, agents are actually indifferent between both effort levels 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 characterizing 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.

$c(u_i) = u_i^\gamma$, $\gamma > 1$. 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: Effort and delay 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 (since 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 individual effort level does not go as quickly down, as we increase the number of agents, relative to the baseline model. In turn, this might lead to a lower expected time of a breakthrough (conditional on a breakthrough occurring in finite time) for larger teams.¹⁴ 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).

The other finding that seems to depend significantly on the linear cost structure is the equilibrium characterization with a deadline. Indeed, with convex costs, one suspects that the equi-

¹⁴Such conditioning is meaningful, since the overall effort over time is still independent of n .




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

librium effort should be a continuous function of time, and this is indeed the case, as shown in the right panel of Figure 10.

Incomplete Information: In many applications, uncertainty does not only pertain 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 $\alpha \in \{\alpha_L, \alpha_H\}$ it 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, since agents are both updating on the quality of the project, and the productivity of their opponent, and this updating depends on the effort 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 by, he becomes quickly more pessimistic, while the high-type agent, who does not work during that time, does not update his belief downward as fast (since he knows he does not work, the lack of success is not as surprising to him). 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 wilts 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 illustrated below.




Figure 11: Effort under incomplete information

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.
- [3] Bergemann, D. and U. Hege, 2005. "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, **36**, 719-52.
- [4] Bergin, J. and W. B. McLeod, 1993. "Continuous Time Repeated Games," *International Economic Review*, **34**, 21-37.

- [5] Bolton P. and C. Harris, 1999. "Strategic Experimentation," *Econometrica*, **67**, 349-374.
- [6] Compte, O. and P. Jehiel, 2004. "Gradualism in Bargaining and Contribution Games," *Review of Economic Studies*, **71(4)**, 975-1000.
- [7] Doz, Y. L., 1996. "The Evolution of Cooperation in Strategic Alliances: Initial Conditions or Learning Processes?," *Strategic Management Journal*, **17(7)**, 619-52.
- [8] Gulati, R., 1995. "Social Structure and Alliance Formation Patterns: A Longitudinal Analysis," *Administrative Science Quarterly*, **40**, 85-112.
- [9] 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.
- [10] 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.
- [11] Holmström, B., 1982. "Moral Hazard in Teams," *The Bell Journal of Economics*, **13(2)**, 324-340.
- [12] Hopenhayn, H. and F. Squintani, 2006. "Preemption Games with Private Information," working paper, University College London.
- [13] Kamihigashi, T., 2001. "Necessity of Transversality Conditions for Infinite Horizon Problems," *Econometrica*, **69(4)**, 995-1012.
- [14] 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.
- [15] Keller, G., S. Rady and M. Cripps, 2005. "Strategic Experimentation with Exponential Bandits," *Econometrica*, **73**, 39-68.
- [16] Klein, N. and S. Rady, 2008. "Negatively Correlated Bandits," Discussion Paper No. 243, SFB/TR 15.

- [17] 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.
- [18] Legros, P., S. Matthews, 1993. "Efficient and Nearly-Efficient Partnerships," *Review of Economic Studies*, **60(3)**, 599-611.
- [19] Levi D., 2007. *Group Dynamics for Teams*. Sage Publications.
- [20] Levin, J. and S. Tadelis, 2005. "Profit Sharing and the Role of Professional Partnerships," *The Quarterly Journal of Economics*, **120**, 131-171.
- [21] Lockwood, B. and J. P. Thomas, 2002. "Gradualism and Irreversibility," *Review of Economic Studies*, **69(2)**, 339-356.
- [22] Ma, C.-t. A. 1988. "Unique implementation of incentive contracts with many agents," *Review of Economic Studies*, **55**, 555-572.
- [23] Marx, L. and S. Matthews, 2000. "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies*, **67(2)**, 327-358.
- [24] McGrath, 1991. "Time, interaction, and performance (TIP): A theory of groups," *Small Group Behavior*, **4(3)**, 147-174.
- [25] Murto, P. and J. Välimäki, 2008. "Learning and Information Aggregation in an Exit Game," working paper, Helsinki Center of Economic Research.
- [26] Olson, M., 1965. *The Logic of Collective Action*. Harvard University Press.
- [27] Presman, E.L. and I.N. Sonin, 1990. *Sequential Control with Incomplete Information: The Bayesian Approach to Multi-Armed Bandit Problems*. Academic Press, London.
- [28] Presman, E.L., 1990. "Poisson Version of the Two-Armed Bandit Problem with Discounting," *Theory of Probability and its Applications*, **35**, 307-317.

- [29] Rosenberg, D., E. Solan and N. Vieille, 2007. “Social Learning in One-Arm Bandit Problems,” *Econometrica*, **75(6)**, 1591-1611.
- [30] Seierstad, A. and K. Sydsaeter, 1987. *Optimal Control Theory with Economic Applications*. North-Holland.
- [31] Strausz, R., 1999. “Efficiency in Sequential Partnerships,” *Journal of Economic Theory*, **85**, 140-156.
- [32] Tuckman B. W., and M. C. Jensen, 1977. “Stages of small-group development revisited,” *Group and Organization Studies*, **2(4)**, 419-427.
- [33] S. 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.
- [34] 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

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

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 $p > \alpha$, which is equivalent to $x_0 < \ln \beta$, where $x_0 = \ln((1 - \bar{p})/\bar{p})$. We also assume that players are sufficiently patient. In particular, we assume that $1 + e^{-x_0}(\lambda/r - \beta) > 0$.

(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 $p > \alpha$ 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 in 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 $u_{j,t}^*(|A| + 1)$ 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 economy. 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 $q_t := p_t/(1 - p_t)$, is equivalent to

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

so that the maximized Hamiltonian is concave in q , and sufficiency then follows from the Arrow sufficiency theorem (see Seierstad and Sydsaeter, 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 decreasing in r and α , and increasing in $\bar{p}$.

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

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

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 conditions 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 exists 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 $u_1 > 0, \dot{x} > 0$, and 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 the case 3, γ_2 must diverge to $+\infty$, and so must x , violating the transversality condition. If instead $\gamma_2 < \bar{\gamma}_2$, then either $u_1 = 0$ 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 so 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 $\dot{\gamma}_i > 0$, but $\dot{\gamma}_j < 0$, for some $i = 1, 2$, but this would point 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 she 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 $u = 1$ 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 $t \in [\delta, T]$. It is easy to check that the term in brackets is decreasing in δ , and setting the derivative with respect to δ equal to 0 at $\delta = 0$ gives

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

(As is immediate to check, 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 $u = 1$ 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)

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

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 $\left. \frac{dV_i(T_1)}{dT_1} \right|_{T_1=0} \leq 0$. Since clearly it is optimal to have agents choose $u = 1$ as long as possible, it follows that the optimal value is $T_1 = 0$: agents should be given a deadline for which it is optimal to exert maximal effort immediately. $\square$

B.2 Proofs for Section 4.2.

Proof of Theorem 4: The 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 part 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.

Let us now turn to the problem of the 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. $\square$

C Proofs for Section 5

Proof of Theorem 5: 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 $v'(x) \dot{x} = \dot{u}$,

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

This O.D.E. characterizes the interior solution (if any) of the problem. For $f(u_1, \dots, u_n) = (\Sigma_i u_i^\rho)^{\frac{1}{\rho}}$, then $\Sigma_j f_{ij} = 0$, 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 maximal effort 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 , it is simple to solve for its inverse $t(p)$, and to verify that $\lim_{p \rightarrow \alpha k_n} t(p) = \infty$, while $t(p) < \infty$ for $p > \alpha k_n$. $\square$

Proof of Theorem 6:

The case $\rho \geq 1/2$: Writing equation (1) 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:

$$\left(1 - \frac{r}{\alpha u_i} \sigma_i^\rho \right) (u_1^\rho + u_2^\rho)^{1/\rho} = \frac{1}{p} ((1 - \rho) \left(\frac{\dot{u}_i}{u_i} - \frac{\dot{u}_j}{u_j} \right) - r(1 + \sigma_i^\rho)). \quad (2)$$

By adding equations (2) for each player, we obtain the relationship that needs to hold between u_1 and u_2 . Formally,

$$u_2 = \frac{r}{2\alpha\sigma_1^\rho} \left(\sigma_1^{2\rho-1} + 1 - \frac{\alpha}{p} (\sigma_1^\rho + 1)^{2-\frac{1}{\rho}} \right). \quad (3)$$

Similarly, subtract equations (2), substitute the expression for u_2 in (3), and use $u_1 = \sigma_1 u_2(\sigma_1)$.

We then obtain the following ordinary differential equation for σ_1 :

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

where $\sigma_1(p_0) = u_{1,0}/u_{2,0}$. Notice that σ_1 will not be a function of r (hence patience will only influence effort levels, not 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 $\rho \geq 1/2$,

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

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 (ODE) is bounded by $-2 / ((1-p)(1 - \alpha/p)^\rho)$, the solution to (ODE) must lie below the solution of the differential equation

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

with the same initial condition at p_0 . 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 $t(p)$ 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 (ODE) 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:

$$g(p) := \frac{(y_0 - 1)^2}{2y_0} \left(\frac{p_0 - \frac{1}{2}\alpha}{1 - p_0} \frac{1-p}{p - \frac{1}{2}\alpha} \right)^{\frac{\alpha}{2-\alpha}}.$$

(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 effort levels from

equation (3), 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 letting $K = \frac{(y_0-1)^2}{y_0} \left(\frac{p_0 - \frac{1}{2}\alpha}{1-p_0} \right)^{\frac{\alpha}{2-\alpha}}$, we can show that the solution to $p' = -p(1-p)f(u_1(p), u_2(p))$ lies below the solution to

$$p' = -(1-p_0)(p - \alpha/2) \left(\left(\frac{1-p_0}{p - \frac{1}{2}\alpha} \right)^{\frac{\alpha}{2-\alpha}} + \frac{2}{K} \right) \frac{2rK}{\alpha},$$

which converges to $\alpha/2$ in finite time. And 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 $p(\sigma)$,

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

It is immediate to check 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 (3) and (ODE) 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 7: Applying integration by parts and ignoring constant terms, the payoff that agent i maximizes equals

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

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

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

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

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

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 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 check that, for low enough discounting (more precisely, whenever $u(q_0)$, as given above, is less than 1), agents cannot exert maximal effort 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 maximal effort 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 interior effort, 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 her individual threshold p_1^* . If the asymmetry in $\bar{p}_i$ is sufficiently small, player 2 would then start working from there, since we would have $p_2 > p_1^*$. But then, anticipating that player 2 will start working at full speed, player 1 wants to deviate and shirk in the last instants before her 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 :

$$u_i = \frac{r}{\alpha} \frac{1 - q + e^{-x_i} + \alpha q}{1 - q + e^{-x_j}}. \quad (4)$$

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

$$\dot{q} = (1 - q) (u_1 + u_2) + u_1 e^{-x_2} + u_2 e^{-x_1}, \quad (5)$$

and obtain as a general solution:

$$q = 1 - k e^{-x_1-x_2} + e^{-x_1} + e^{-x_2} \quad (6)$$

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 (4) in (5), 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 (6), 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 working at full capacity, and she 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 maximal effort. 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)

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

Since 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. $\square$