

MIT Open Access Articles

Collaborating

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Bonatti, Alessandro, and Johannes Hörner. 2011. "Collaborating." American Economic Review, 101(2): 632–63.

As Published: <http://dx.doi.org/10.1257/aer.101.2.632>

Publisher: American Economic Association

Persistent URL: <http://hdl.handle.net/1721.1/65916>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Collaborating

By ALESSANDRO BONATTI AND JOHANNES HÖRNER*

This paper examines moral hazard in teams over time. Agents are collectively engaged in a project whose duration and outcome are uncertain, and their individual efforts are unobserved. Free-riding leads not only to a reduction in effort, but also to procrastination. Collaboration among agents dwindles over time, but does not cease as long as the project has not succeeded. In addition, the delay until the project succeeds, if it ever does, increases with the number of agents. We show why deadlines, but not necessarily better monitoring, help to mitigate moral hazard. (JEL D81, D82, D83)

Teams and partnerships are playing an increasingly important role in economic activity, at the levels of individual agents and firms alike. In research, for instance, we observe an apparent shift from the individual-based model to a teamwork model, as documented by Stefan Wuchty, Benjamin F. Jones, and Brian Uzzi (2007). In business, firms are using collaborative methods of decision making more and more, so the responsibility for managing a project is shared by the participants (see Michael Hergert and Deigan Morris 1988). Yet despite the popularity of joint ventures, firms experience difficulties in running them.

The success of such ventures relies on mutual trust, such that each member of the team believes that every other member will pull his weight. However, it is often difficult to distinguish earnest, but unlucky, effort from free-riding because the duration and outcome of the project are uncertain. This difficulty has practical consequences in cases of persistent lack of success. In such cases, agents may grow suspicious that other team members are not pulling their weight in the common enterprise and scale back their own involvement in response, which will result in disruptions and, in some cases, the ultimate dissolution of the team. There is a large literature in management that documents the difficulties involved in running joint ventures and the risks of shirking that are associated with them (see, for instance, Anoop Madhok 1995; Yadong Luo 2002). Yaping Gong et al. (2007) show how cooperation in international joint ventures decreases as the number of partners increases, and relate their finding to the increased room for opportunistic behavior that arises as the venture becomes larger and more complex.

* Bonatti: MIT Sloan School of Management, 50 Memorial Drive, Cambridge MA 02142 (e-mail: bonatti@mit.edu); Hörner: Department of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06520 (e-mail: johannes.horner@yale.edu). We would like to thank Dirk Bergemann, Martin Cripps, Glenn Ellison, Meg Meyer, David Miller, Motty Perry, Sven Rady, Larry Samuelson, Roland Strausz, and Xavier Vives for useful discussions, and Nicolas Klein for excellent comments. We thank Glenn Ellison for providing us with the data that is used in the working paper. Finally, we thank three anonymous referees, whose suggestions have led to numerous improvements to the paper. In particular, their suggestions led to Sections I, VE, and a large part of IVA. Hörner gratefully acknowledges support from NSF grant 0920985.

This paper provides a model of how teams work toward their goals under conditions of uncertainty regarding the duration and outcome of the project, and examines possible remedies. Our model may be applied both intrafirm, for example, to research teams, and interfirm, for example, to R&D joint ventures and alliances. The benefits of collaborative research are widely recognized, and R&D joint ventures are encouraged under both US and EU competition law and funding programs.¹ Nevertheless, any firm that considers investing resources in such projects faces all the obstacles that are associated with contributing to a public good.

With very few exceptions, previous work on public goods deals with situations that are either static or involve complete information. This work has provided invaluable insights into the *quantitative* underprovision of the public good. In contrast, we account explicitly for uncertainty and learning, and therefore concentrate our attention on the *dynamics* of this provision.²

The key features of our model are as follows: (i) *Benefits are public, costs are private*: the value from completing the project is common to all agents. All it takes to complete the project is one breakthrough, but making a breakthrough requires costly effort. (ii) *Success is uncertain*: some projects will fail, no matter how much effort is put into them, due to the nature of the project, though this will not, of course, be known at the start.³ As for those projects that can succeed, the probability of making a breakthrough increases as the combined effort of the agents increases. Achieving a breakthrough is the only way to ascertain whether a project is one that must fail or can succeed. (iii) *Effort is hidden*: the choice of effort exerted by an agent is unobserved by the other agents. As long as there is no breakthrough, agents will not have any hard evidence as to whether the project can succeed; they simply become (weakly) more pessimistic about the prospects of the project as time goes on. This captures the idea that output is observable, but effort is not. We shall contrast our findings with the case in which effort is observable. At the end of the paper, we also discuss the intermediate case, in which the project involves several observable steps.

Our main findings are the following:

- *Agents procrastinate*: as is to be expected, agents slack off, i.e., there is underprovision of effort overall. Agents exert too little effort, and exert it too late. In the hope that the effort of others will suffice, they work less than they should early on, postponing their effort to a later time. Nevertheless, due to growing pessimism, the effort expended dwindles over time, but the plug is never pulled on the project. Although the overall effort expended is independent of the size

¹ More specifically, R&D joint ventures enjoy a block exemption from the application of Article 81(3) of the EU Treaty, and receive similar protection in the United States under the National Cooperative Research and Production Act and the Standards Development Organization Advancement Act. Further, cooperation on transnational research projects was awarded 64 percent of the nonnuclear, €51 billion budget under the 2007–2013 7th Framework Programs of the European Union for R&D.

² The management literature also stresses the importance of learning in alliances. For example, Yves L. Doz and Gary Hamel (1998) claim that to sustain successful cooperation, partners typically need to learn in five key areas: the environment in which the alliance will operate, the tasks to be performed, the process of collaboration, the partners' skills, and their intended and emerging goals.

³ This is typically the case in pharmaceutical research. Doz and Hamel (1998) report that the pharmaceutical giant Merck assembled a complex network of research institutes, universities, and biotechnology companies to develop AIDS vaccines and cures in the early 1990s. Given the serendipitous nature of pharmaceutical innovations, there was no way of knowing which, if any, part of this web of R&D alliances would be productive.

of the team, the more agents are involved in the project, the later the project gets completed on average, if ever.

- *Deadlines are beneficial*: if agents have enough resolve to fix their own deadline, it is optimal to do so. This is the case despite the fact that agents pace themselves so that the project is still worthwhile when the deadline is hit and the project is abandoned. If agents could renegotiate at this time, they would. However, even though agents pace themselves too slowly, the deadline gives them an incentive to exert effort once it looms closely enough. The deadline is desirable, because the reduction in wasteful delay more than offsets the value that is forfeited if the deadline is reached. In this sense, delay is more costly than the underprovision of effort.
- *Better monitoring need not reduce delay*: when effort is observed, there are multiple equilibria. Depending on the equilibrium, the delay might be greater or smaller than under nonobservability. In the unique symmetric Markov equilibrium, the delay is greater. This is because individual efforts are strategic substitutes. The prospects of the team improve if it is found that an agent has slacked off, because this mitigates the growing pessimism in the team. Therefore, an observable reduction in current effort encourages later effort by other members, and this depresses equilibrium effort. Hence, better monitoring need not alleviate moral hazard. Nevertheless, there are also non-Markovian, grim-trigger equilibria for which the delay is smaller.

We finally discuss how our results can be generalized in several dimensions. We discuss more general mechanisms than deadlines. In particular, we consider the optimal dynamic (budget-balanced) compensation scheme, as well as the profit-maximizing wage scheme for a principal who owns the project's returns. One rationale that underlies the formation of partnerships is the possible synergies that may arise between agents. We consider two types of synergy: (i) agents might be identical, but their combined efforts might be more productive than the sum of their isolated efforts, and (ii) agents might have different skills, and depending on the problem at hand, one or the other might be more likely to succeed.⁴ We also examine how private information regarding the agent's productivity and learning-by-doing affect our results. Finally, we remark on the case in which the project involves completing multiple tasks.

This paper is related to several strands of literature. First, our model can be viewed as a model of experimentation. There is a growing literature in economics on experimentation in teams. For instance, Patrick Bolton and Christopher Harris (1999), Godfrey Keller, Sven Rady, and Martin Cripps (2005), and Nicolas Klein and Rady (2008) study a two-armed bandit problem in which different agents may choose different arms. While free-riding plays an important role in these studies as well, effort is always observable. Dinah Rosenberg, Eilan Solan, and Nicolas Vieille (2007), Hugo Hopenhayn and Francesco Squintani (2004), and Pauli Murto and

⁴This kind of synergy typically facilitates successful collaborations. Doz and Hamel (1998) report the case of the Japanese Very Large Scale Integration (VLSI) research cooperative. The VLSI project involved five semiconductor manufacturers: Fujitsu, Hitachi, Mitsubishi Electric, NEC, and Toshiba. Interteam cooperation was facilitated by the choice of overlapping research areas: several teams worked on the same types of topic, but in slightly differentiated ways, thereby providing for original contributions from each team.

Juuso Välimäki (2008) consider the case in which the outcome of each agent's action is unobservable, while their actions are observable. This is precisely the opposite of what is assumed in this paper; here, actions are not observed, but outcomes are. Dirk Bergemann and Ulrich Hege (2005) study a principal-agent relationship with an information structure similar to the one considered here. All these models provide valuable insights into how much total experimentation is socially desirable, and how much can be expected in equilibrium. As will become clear, these questions admit trivial answers in our model, which is therefore not well suited to address them.

Robin Mason and Välimäki (2008) consider a dynamic moral hazard problem in which effort by a single agent is unobservable. Although there is no learning, the optimal wage declines over time, to provide incentives for effort. Their model shares a number of features with ours. In particular, the strategic substitutability between current and later efforts plays an important role in both models, so that, in both cases, deadlines have beneficial effects. See also Flavio Toxvaerd (2007) on deadlines, and Tracy R. Lewis and Marco Ottaviani (2008) on similar effects in the optimal provision of incentives in sequential search.

Second, our model ties into the literature on free-riding in groups, starting with Mancur Olson (1965) and Armen A. Alchian and Harold Demsetz (1972), and further studied in Bengt Holmström (1982), Patrick Legros and Steven A. Matthews (1993), and Eyal Winter (2004). In a sequential setting, Roland Strausz (1999) describes an optimal sharing rule. Our model ties in with this literature in that it covers the process by which free-riding occurs over time in teams that are working on a project whose duration and outcome is uncertain; i.e., it is a dynamic version of moral hazard in teams with uncertain output. A static version was introduced by Steven R. Williams and Roy Radner (1995) and also studied by Ching-To Ma, John Moore, and Stephen Turnbull (1988). The inefficiency of equilibria of repeated partnership games with imperfect monitoring was first demonstrated by Radner, Roger Myerson, and Eric Maskin (1986).

Third, our paper is related to the literature on dynamic contributions to public goods. Games with observable contributions are examined in Anat R. Admati and Motty Perry (1991), Olivier Compte and Philippe Jehiel (2004), Chaim Fershtman and Shmuel Nitzan (1991), Ben Lockwood and Jonathan P. Thomas (2002), and Leslie M. Marx and Matthews (2000). Fershtman and Nitzan (1991) compare open- and closed-loop equilibria in a setup with complete information and find that observability exacerbates free-riding. In Parimal Kanti Bag and Santanu Roy (2008), Christopher Bliss and Barry Nalebuff (1984), and Mark Gradstein (1992), agents have independently drawn and privately known values for the public good. This type of private information is briefly discussed in the conclusion. Applications to partnerships include Jonathan Levin and Steven Tadelis (2005) and Barton H. Hamilton, Jack A. Nickerson, and Hideo Owan (2003). Also related is the literature in management on alliances, including, for instance, Doz (1996), Ranjay Gulati (1995), and Gulati and Harbir Singh (1998).

There is a vast literature on free-riding, also known as social loafing, in social psychology. See, for instance, Bibb Latané, Kipling Williams, and Stephen Harkins (1979), or Steven J. Karau and Williams (1993). Levi (2007) provides a survey of group dynamics and team theory. The stage theory, developed by Bruce W. Tuckman

and Mary Ann C. Jensen (1977) and the theory by Joseph E. McGrath (1991) are two of the better-known theories regarding the development of project teams: the patterning of change and continuity in team structure and behavior over time.

I. A Simple Example

Consider the following two-period game. Agent $i = 1, 2$ may exert effort in two periods, $t = 1, 2$, in order to achieve a breakthrough. Whether a breakthrough is possible depends on the quality of the project. If the project is good, the probability of a breakthrough in period t (assuming that there was no breakthrough before) is given by the sum of the effort levels $u_{i,t}$ that the two agents choose in that period. However, the project might be bad, in which case a breakthrough is impossible. Agents share a common prior belief $\bar{p} < 1$ that the project is good.

The project ends if a breakthrough occurs. A breakthrough is worth a payoff of one to both agents, independently of who is actually responsible for this breakthrough. Effort, on the other hand, entails a private cost given by $c(u_{i,t})$ in each period. Payoffs from the second period are discounted at a common factor $\delta \leq 1$.

Agents do not observe their partner's effort choice. All they observe is whether a breakthrough occurs. Therefore, if there is no breakthrough at the end of the first period, agents update their belief about the quality of the project based only on their own effort choice, and their expectation about the other agent's effort choice. Thus, if an agent chooses an effort level $u_{i,t}$ in each period, and expects his opponent to exert effort $\hat{u}_{-i,t}$, his expected payoff is given by

$$(1) \quad \underbrace{\bar{p} \cdot (u_{i,1} + \hat{u}_{-i,1}) - c(u_{i,1})}_{\text{First period payoff}} + \delta(1 - \bar{p} \cdot (u_{i,1} + \hat{u}_{-i,1})) \underbrace{[\rho(u_{i,1}, \hat{u}_{-i,1}) \cdot (u_{i,2} + \hat{u}_{-i,2}) - c(u_{i,2})]}_{\text{Second period payoff}},$$

where $\rho(u_{i,1}, \hat{u}_{-i,1})$ is his posterior belief that the project is good. To understand (1), note that the probability of a breakthrough in the first period is the product of the prior belief assigned to the project being good (\bar{p}), and the sum of effort levels exerted ($u_{i,1} + \hat{u}_{-i,1}$). The payoff of such a breakthrough is one. The cost of effort in the first period, $c(u_{i,1})$, is paid in any event. If a breakthrough does not occur, agent i updates his belief to $\rho(u_{i,1}, \hat{u}_{-i,1})$, and the structure of the payoff in the second period is as in the first period.

By Bayes's rule, the posterior belief of agent i is given by

$$(2) \quad \rho(u_{i,1}, \hat{u}_{-i,1}) = \frac{\bar{p} \cdot (1 - u_{i,1} - \hat{u}_{-i,1})}{1 - \bar{p} \cdot (u_{i,1} + \hat{u}_{-i,1})} \leq \bar{p}.$$

Note that this is based on agent i 's expectation of agent $-i$'s effort choice. That is, agents' beliefs are private, and they coincide only on the equilibrium path. If, for

instance, agent i decides to exert more effort than he is expected to by agent $-i$, yet no breakthrough occurs, agent i will become more pessimistic than agent $-i$, unbeknownst to him. Off path, beliefs are no longer common knowledge.

In a perfect Bayesian equilibrium, agent i 's effort levels $(u_{i,1}, u_{i,2})$ are optimal given $(\hat{u}_{-i,1}, \hat{u}_{-i,2})$, and expectations are correct: $\hat{u}_{-i,t} = u_{-i,t}$. Letting V_i denote the agent's payoff, it must be that

$$\frac{\partial V_i}{\partial u_{i,1}} = \bar{p} - c'(u_{i,1}) - \delta \bar{p} \cdot (u_{i,2} + \hat{u}_{-i,2} - c(u_{i,2})) = 0,$$

and

$$\frac{\partial V_i}{\partial u_{i,2}} \propto \rho(u_{i,1}, \hat{u}_{-i,1}) - c'(u_{i,2}) = 0,$$

so that, in particular, $c'(u_{i,2}) < 1$. It follows that (i) the two agents' first-period effort choices are neither strategic complements nor substitutes, but (ii) an agent's effort choices across periods are strategic substitutes, as are (iii) an agent's current effort choice and the other agent's future effort choices.

It is evident from (ii) that the option to delay reduces effort in the first period. Comparing the one- and two-period models is equivalent to comparing the first-period effort choice for $u_{i,2} = \hat{u}_{i,2} = 0$ on the one hand, and a higher value on the other. This is what we refer to as *procrastination*: some of the work that would otherwise be carried out by some date gets postponed when agents get further opportunities to work afterward.⁵ In our example, imposing a deadline of one period heightens incentives in the initial period.

Further, it is also clear from (iii) that observability of the first period's action will lead to a decline in effort provision. With observability, a small decrease in the first-period effort level increases the other agent's effort tomorrow. Therefore, relative to the case in which effort choices are unobservable, each agent has an incentive to lower his first-period effort level in order to induce his partner to work harder in the second period, when his choice is observable.

As we shall see, these findings carry through with longer horizons: deadlines are desirable, while observability, or better monitoring, is not. However, this two-period model is ill-suited to describe the dynamics of effort over time when there is no last period. To address this and related issues, it is best to consider a baseline model in which the horizon is infinite. This model is described next.

II. The Setup

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

⁵To procrastinate is to "delay or postpone action," as defined by the *Oxford English Dictionary*.

Agents continuously choose at which level to exert effort over the infinite horizon \mathbb{R}_+ . Effort is costly, and the instantaneous cost to agent $i = 1, \dots, n$ of exerting effort $u_i \in \mathbb{R}_+$ is $c_i(u_i)$, for some function $c_i(\cdot)$ that is differentiable and strictly increasing. In most of the paper, we assume that $c_i(u_i) = c_i \cdot u_i$, for some constant $c_i > 0$, and that the choice is restricted to the unit interval, i.e., $u_i \in [0, 1]$. The effort choice is, and remains, unobserved.

Effort is necessary for a breakthrough to occur. More precisely, a breakthrough occurs with instantaneous probability equal to $f(u_1, \dots, u_n)$, if the project is good, and to zero if the project is bad. That is, if agents were to exert a constant effort u_i over some interval of time, then the delay until they found out that the project is successful would be distributed exponentially over that time interval with parameter $f(u_1, \dots, u_n)$. The function f is differentiable and strictly increasing in each of its arguments. In the baseline model, we assume that f is additively separable and linear in effort choices, so that $f(u_1, \dots, u_n) = \sum_{i=1, \dots, n} \lambda_i u_i$, for some $\lambda_i > 0$, $i = 1, \dots, n$.

The game ends if a breakthrough occurs. Let $\tau \in \mathbb{R}_+ \cup \{+\infty\}$ denote the random time at which the breakthrough occurs ($\tau = +\infty$ if it never does). We interpret such a breakthrough as the successful completion of the project. A successful project is worth a net present value of one to each of the agents.⁶ As long as no breakthrough occurs, agents reap no benefits from the project. Agents are impatient, and discount future benefits and costs at a common discount rate r .

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

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

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

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

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

with $p_0 = \bar{p}$.⁷ Given that the probability that the project is good at time t is p_t , and that 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

⁶We discuss this assumption further in Section IIIB.

⁷To see this, note that, given p_t , the belief at time $t + dt$ is

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

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

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

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

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

III. The Benchmark Model

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

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

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

$$(3) \quad V_i(\bar{p}) := r \int_0^\infty \left(p_t \sum_i u_{i,t} - \alpha u_{i,t} \right) e^{-\int_0^t (p_s \sum_i u_{i,s} + r) ds} dt,$$

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: it is equal to the belief at which a myopic agent would stop working, because at this point the instantaneous marginal revenue from effort, p_t , equals the marginal cost, α .

A. The Team Problem

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

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

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

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

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

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

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

Observe that the instantaneous marginal benefit from effort to an agent is equal to p_t , which decreases over time, while the marginal cost is constant and equal to α . Therefore, it will not be possible to provide incentives for selfish agents to exert effort beyond the Marshallian threshold. The wedge between this threshold and the efficient one, α/n , captures the well-known free-riding effect in teams, which is described eloquently by Alchian and Demsetz (1972), and has since been studied extensively. In a noncooperative equilibrium, the amount of effort is too low.⁸ Here, instead, our focus is on how free-riding affects when effort is exerted.

B. The Noncooperative Solution

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

$$(4) \quad \frac{\lambda}{r} \geq \alpha^{-1} - \bar{p}^{-1} > 0.$$

This assumption ensures that the upper bound on the effort level does not affect the analysis.⁹ 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.

What is the trade-off between exerting effort at some instant and exerting it at the next? Fix some date t , and assume that players have followed the equilibrium strategies up to that date. Fix also some small $dt > 0$, and consider the gain or loss from

⁸There is neither an "encouragement effect" in our setup, unlike in some papers on experimentation (see, for instance, Bolton and Harris 1999), nor any effect of patience on the threshold. This is because a breakthrough yields a unique lump sum to all agents, rather than conditionally independent sequences of lump sum payoffs.

⁹See Bonatti and Hörner (2009) for an equilibrium analysis when (4) is not satisfied.

shifting some small effort ε from the time interval $[t, t + dt]$ (“today”) to the time interval $[t + dt, t + 2dt]$ (“tomorrow”). 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 . The payoff $V_{i,t}$ must satisfy the recursion

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

Because we are interested in the trade-off between effort today and tomorrow, we apply the same expansion to $V_{i,t+dt}$, to obtain

$$\begin{aligned} (5) \quad V_{i,t} &= (p(u_i + u_{-i}) - \alpha u_i) dt \\ &+ (1 - rdt)(1 - p(u_i + u_{-i}) dt) \\ &\times [(p'(u'_i + u'_{-i}) - \alpha u'_i) dt + (1 - rdt)(1 - p'(u'_i + u'_{-i}) dt) V_{i,t+2dt}], \end{aligned}$$

where $p' = p - p(1 - p)(u_i + u_{-i}) dt$.¹⁰ Consider, then, decreasing u_i by ε and increasing u'_i by that amount. Note that, conditional on reaching $t + 2dt$ without a breakthrough, the resulting belief is unchanged, and, therefore, so is the continuation payoff. That is, $V_{i,t+2dt}$ is independent of ε .

Therefore, to the second order,

$$\frac{dV_{i,t}/d\varepsilon}{dt} = \underbrace{-((p - \alpha) - pV_{i,t})}_{\frac{dV_{i,t}}{du_i} \cdot \frac{du_i}{d\varepsilon}} + \underbrace{(p - \alpha) - pV_{i,t}}_{\frac{dV_{i,t}}{du'_i} \cdot \frac{du'_i}{d\varepsilon}} = 0.$$

To interpret this, note that increased effort affects the payoff in three ways: it increases the probability of a breakthrough, yielding a payoff of one, at a rate p ; it causes the loss of the continuation value $V_{i,t}$ at the same rate; lastly, it increases cost, at a rate α .

The upshot of this result is that the trade-off between effort today and tomorrow can only be understood by considering an expansion to the third order. Here we must recall that the probability of a breakthrough given effort level u is, to the third order, $pu dt - (pu)^2(dt)^2/2$ (see footnote 10); similarly, the continuation payoff is discounted by a factor $e^{-rdt} \approx 1 - rdt + r^2(dt)^2/2$.

Let us first expand terms in (5), which gives

$$\begin{aligned} V_{i,t} &= (pu - \alpha u_i) dt - (pu)^2(dt)^2/2 + (1 - (r + pu) dt + (r + pu)^2 dt^2/2) \\ &\times [(pu' - \alpha u'_i) dt - ((1 - p)u + pu'/2) pu' dt^2 \\ &+ [1 - (r + pu') dt + ((r + pu')^2/2 + p(1 - p)uu') dt^2] V_{i,t+2dt}], \end{aligned}$$

¹⁰ More precisely, for later purposes, $e^{-p(u_i + u_{-i}) dt} = 1 - p(u_i + u_{-i}) dt + (p^2/2)(u_i + u_{-i})^2(dt)^2 + o(dt^3)$.

where, for this equation and the next, $u := u_i + u_{-i}$, $u' := u'_i + u'_{-i}$. We then obtain, ignoring the second-order terms that, as shown, cancel out,

$$\begin{aligned} & \frac{dV_{i,t}/d\varepsilon}{dt^2} \\ & \quad - \frac{dV_{i,t}/du_i}{dt^2} \cdot \frac{du_i}{d\varepsilon} \\ & \quad \underbrace{= p^2u + p(pu'(1 - V_i) - \alpha u'_i - rV_i) - (r + pu)pV_i + p(1 - p)u'(1 - V_i)} \\ & \quad + \underbrace{p(r + pu')V_i - p(1 - p)u(1 - V_i) - p^2u' - (r + pu)(p(1 - V_i) - \alpha)} \\ & \quad \frac{dV_{i,t}/du'_i}{dt^2} \cdot \frac{du'_i}{d\varepsilon} \end{aligned}$$

Assuming that u_i and u_{-i} are continuous, almost all terms vanish. We are left with

$$\frac{dV_{i,t}/d\varepsilon}{dt^2} = \alpha p(u_i + u_{-i}) - r(p - \alpha) - \alpha p u_i.$$

This means that postponing effort to tomorrow is unprofitable if and only if

$$(6) \quad \alpha p u_i \geq \alpha p(u_i + u_{-i}) - r(p - \alpha).$$

Equation (6) admits a simple interpretation. What is the benefit of working a bit more today, relative to tomorrow? At a rate p (the current belief), working today increases the probability of an immediate breakthrough, in which event the agent will not have to pay the cost of the planned effort tomorrow (αu_i). This is the left-hand side. What is the cost? If the agent waited until tomorrow before working a bit harder, there is a chance that this extra effort would not have to be carried out. The probability of this event is $p \cdot (u_i + u_{-i})$, and the cost saved is α per unit of extra effort. This gives the first term on the right-hand side. Of course, there is also a cost of postponing, given that agents are impatient. This cost is proportional to the mark-up of effort, $p - \alpha$, and gets subtracted on the right-hand side.

First, observe that, as $p \rightarrow \alpha$, the right-hand side of (6) exceeds the left-hand side if u_{-i} is bounded away from zero. Effort tends to zero as p tends to α . Similarly, effort must tend to zero as $r \rightarrow 0$.

Second, assume for the sake of contradiction that agents stop working at some finite time. Then, considering the penultimate instant, it must be that, up to the second order, $p - \alpha = p(1 - p)(u_i + u_{-i})dt$, and so we may divide both sides of (6) by $u_i + u_{-i} = nu_i$, yielding

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

which is impossible, as dt is arbitrarily small. Therefore, not only does effort go to zero as p tends to α , but it does so sufficiently fast that the belief never reaches the

threshold α , and agents keep on working on the project forever, albeit at negligible rates.

It is now easy to guess what the equilibrium value of u_i must be. Given that agent i must be indifferent between exerting effort or not, and also exerting it at different instants, we must have

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

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

This simple logic relies on two substitutability assumptions: efforts of different agents are perfect substitutes, and the cost function is linear. Both assumptions will be relaxed later.

We emphasize that, because effort is not observed, players share a common belief only on the equilibrium path. For an arbitrary history, an agent's best-reply depends both on the public and on his private belief. Using dynamic programming is difficult, because the optimality equation is then a partial differential equation. Pontryagin's principle, on the other hand, is ideally suited, because the other agents' strategies can be viewed as fixed, given the absence of feedback.

The next theorem, proved in the Appendix, describes the strategy on the equilibrium path.¹¹

THEOREM 1: *There exists a unique symmetric equilibrium, in which, on the equilibrium path, the level of effort of any agent is given by*

$$(7) \quad u_{i,t}^* = \frac{r}{n-1} \frac{\alpha^{-1} - 1}{1 + \frac{(1-\bar{p})\alpha}{\bar{p}-\alpha} e^{\frac{n}{n-1}r(\alpha^{-1}-1)t}}, \quad \text{for all } t \geq 0.$$

If an agent deviated, what would his continuation strategy be? Suppose that this deviation is such that, at time t , the aggregate effort of agent i alone over the interval $[0, t]$ is lower than it would have been on the equilibrium path. This means that agent i is more optimistic than the other agents, and his private belief exceeds their common belief. Given that agent i would be indifferent between exerting effort or not if he shared the common belief, his optimism leads him to exert maximal effort until

¹¹ In the case $\bar{p} = 1$ that was ruled out earlier, the game reduces essentially to the static game. The effort level is constant and, because of free-riding, inefficiently low ($u_{i,t} = r(\alpha^{-1} - 1)/(n - 1)$).

the time 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 realized aggregate effort up to t is greater than in equilibrium, then he is more pessimistic than the other agents, and he will provide no effort until the common belief catches up with his private belief, if ever. This completes the description of the equilibrium strategy. In Section VA, we characterize asymmetric equilibria of the baseline model, and allow for asymmetries in the players' characteristics.

From (7), it is immediate to derive the following comparative statics. To avoid confusion, we refer to total effort at time t as the sum of instantaneous, individual effort levels at that time, and to aggregate effort at t as the sum (i.e., the integral) of total effort over all times up to t .

LEMMA 1: *In the symmetric equilibrium:*

- (i) *Individual effort decreases over time and increases \bar{p} .*
- (ii) *Aggregate effort decreases in α and increases in r . It also decreases in, but is asymptotically independent of, n : the probability of an eventual breakthrough is independent of the number of agents, but the distribution of the time of the breakthrough with more agents first-order stochastically dominates this distribution with fewer agents.*
- (iii) *The agent's payoff $V_i(\bar{p})$ is increasing in n and \bar{p} , decreasing in α , and independent of r .*

Total effort is decreasing in n for a given belief p , so that total effort is also decreasing in n for small enough t . However, this implies that the belief decreases more slowly with more agents. Because effort is increasing in the belief, it must then be that total effort is eventually higher in larger teams. Because the asymptotic belief is α , independent of n , aggregate effort must be independent of n as well. Ultimately, then, larger teams must catch up in terms of effort, but this also means that larger teams are slower to succeed.¹²

In particular, for teams of different sizes, the distributions of the random time τ of a breakthrough, conditional on a breakthrough occurring eventually, are ranked by first-order stochastic dominance. We define the expected cost of delay as $1 - E[e^{-r\tau} | \tau < \infty]$. It follows from Lemma 1 that the cost of delay is increasing in n . However, it is independent of r , because more impatient agents work harder, but discount the future more. As mentioned above, the agents' payoffs are also increasing in n . This is obvious for one versus two agents, because an agent may always act as if he were by himself, securing the payoff from a single-agent team. It is less obvious that larger, slower teams achieve higher payoffs. Our result shows that, for larger teams, the reduction in individual effort more than offsets the increased cost of delay. Figure 1 and the left panel of Figure 2 illustrate these results.

¹² As a referee observed, this is reminiscent of the bystander effect, as in the Kitty Genovese case: because more agents are involved, each agent optimally scales down his involvement, resulting in an outcome that worsens with the number of agents.




FIGURE 1. INDIVIDUAL AND TOTAL EFFORT




FIGURE 2. PAYOFFS AND COST OF DELAY

Note: Left panel: value per agent = 1; right panel: value per agent = $1/n$

Note that the comparative statics with respect to the number of agents hinge upon our assumption that the project's returns per agent were independent of n . If, instead, the total value of the project is fixed independently of n , so that each agent's share decreases linearly in n , aggregate effort decreases with the team's size, and tedious calculations show that each agent's payoff decreases as well. See the right panel of Figure 2 for an illustration.

C. A Comparison with the Observable Case

We now contrast the previous findings with the corresponding results for the case in which effort is perfectly observable. That is, we assume here that all agents' efforts are observable, and that agent i 's choice as to how much effort to exert at time t (hereafter, his effort choice) may depend on the entire history of effort choices up to time t . Note that the "cooperative," or socially optimal, solution is the same whether effort choices are observed or not: delay is costly, so that all work should be carried out as fast as possible; the threshold belief beyond which such work is unprofitable must be, as before, $p = \alpha/n$, which is the point at which marginal benefits of effort to the team are equal to its marginal cost.

Such a continuous-time game involves well-known nontrivial modeling choices. A standard way to sidestep these choices is to focus on Markov strategies. Here, the obvious state variable is the belief p . Unlike in the unobservable case, this belief is always commonly held among agents, even after histories off the equilibrium path.

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

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

Taking limits as $dt \rightarrow 0$ yields

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

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

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

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

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

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

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

It is standard to verify that the resulting u is the unique equilibrium strategy profile provided that \bar{p} is such that $u \leq \lambda$ for all $p < \bar{p}$. In particular, this is satisfied under assumption (4) on the discount rate, which we maintain henceforth. In the model without observability, recall that, in terms of the belief p , effort is given

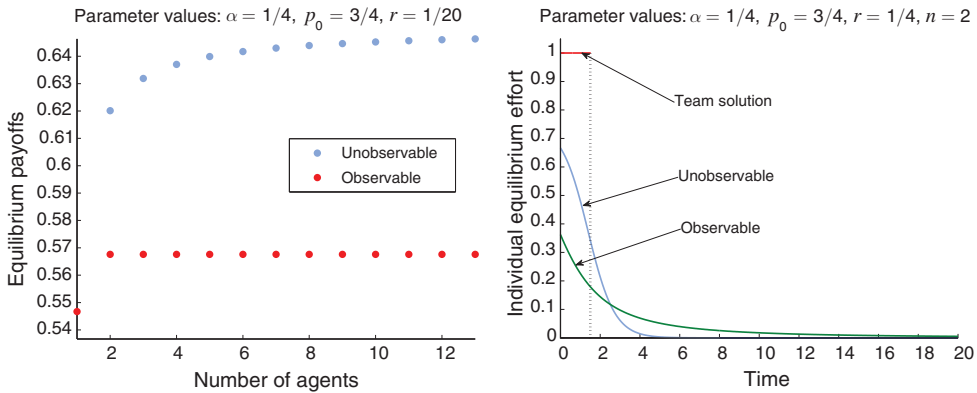


FIGURE 3. WELFARE AND EFFORT IN THE OBSERVABLE VERSUS NONOBSERVABLE CASE

by $u(p) = (r/(n-1))(\alpha^{-1} - p^{-1})$. As is clear from these formulas, the eventual belief is the same whether effort is observed or not, and so aggregate effort over time is the same in both models. However, delay is not.

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

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

The individual payoff is independent of the number of agents $n \geq 2$ in the observable case. This is a familiar rent-dissipation result: when the size of the team increases, agents waste in additional delay what they save in individual cost. This can be seen directly from the formula for the level of effort, in that the total effort of all agents but i is independent of n . It is worth pointing out that this is not true in the unobservable case. This is one example in which the formula that gives the effort as a function of the common belief is misleading in the unobservable case: given p , the total instantaneous effort of all agents but i is independent of n here as well. Yet the

value of p is not a function of the player's information only: it is the common belief about the unobserved past total efforts, including i 's effort; hence, it depends on the number of agents. As we have seen, welfare is actually increasing in the number of agents in the unobservable case. The comparison is illustrated in the left panel of Figure 3.

In the observable case, there also exist asymmetric Markov equilibria, similar to those described in Keller, Rady, and Cripps (2005), in which agents "take turns" at exerting effort. In these Markovian equilibria, the "switching points" are defined in terms of the common belief. Because effort is observable, if agent i procrastinates, this "freezes" the common belief and therefore postpones the time of switching until agent i makes up for the wasted time. So, the punishment for procrastination is automatic. Taking turns is impossible without observability. Suppose that agent i is expected to exert effort alone up to time t , while another agent j exerts effort alone during some time interval starting at time t . Any agent working alone must be exerting maximal effort, if at all, because of discounting. Because any deviation by agent i is not observable, agent j will start exerting effort at time t no matter what. It follows that, at a time earlier than, but close enough to, t , agent i can procrastinate, wait for the time t' at which agent j will stop, and only then, if necessary, make up for this foregone effort (t' is finite because j exerts maximal effort). This alternative strategy is a profitable deviation for i if he is patient enough, because the induced probability that the postponed effort will not be exerted more than offsets the loss in value due to discounting. Therefore, switching is impossible without observability.

In the observable case, there also exist other, non-Markovian symmetric equilibria. As mentioned above, appropriate concepts of equilibrium have been defined carefully elsewhere (see, for instance, James Bergin and W. Bentley McLeod 1993). It is not difficult to see how one can define formally a "grim-trigger" equilibrium, for low enough discount rates, in which all agents exert effort at a maximal rate until time T_1 , at which $p = \alpha$, and if there is a unilateral deviation by agent i , all other agents stop exerting effort, leaving agent i with no choice but to exert effort at a maximal rate from this point on until the common belief reaches α . While this equilibrium is not first-best, it clearly does better than the Markov equilibrium in the observable case, and than the symmetric equilibrium in the unobservable case.¹³

IV. Deadlines

In the absence of any kind of commitment, the equilibrium outcome described above seems inevitable. Pleas and appeals to cooperate are given no heed and deadlines are disregarded. In this section, we examine whether a self-imposed time-limit can mitigate the problem. We assume that agents can commit to a deadline. We first treat the deadline as exogenous, and determine optimal behavior given this deadline. In a second step, we solve for the optimal deadline. In this section, we normalize the capacity λ to one.

¹³ It is tempting to consider grim-trigger strategy profiles in which agents exert effort for beliefs below α . We ignore them here, because such strategy profiles cannot be limits of equilibria of discretized versions of the game.

A. Equilibrium Behavior Given a 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

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

That is, if time T arrives and no breakthrough has occurred, the continuation payoff of the agents is nil. The baseline model of Section III is the special case in which $T = \infty$. The next lemma, which we prove in the Web Appendix, describes the symmetric equilibrium for $T < \infty$. Throughout this subsection, we maintain the restriction on the discount rate r , given by (4), that we imposed in Section IIIB.


LEMMA 2: *Given $T < \infty$, there exists a unique symmetric equilibrium, characterized by $\tilde{T} \in [0, T]$, in which the level of effort is given by*

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

where u_i^* is as in Theorem 1. The time \tilde{T} is nondecreasing in the parameter T and strictly increasing for T large enough. Moreover, the belief at time T strictly exceeds α .

See Figure 4. According to Lemma 2, effort is first decreasing over time, and over this time interval, it is equal to its value when the deadline is infinite. At that point, the deadline is far enough in the future not to affect the agents' incentives. However, at some point, the deadline looms large above the agents. Agents recognize that the deadline is near and exert maximal effort from then on. But it is then too late to catch up with the aggregate effort exerted in the infinite-horizon case, and $p_T > \alpha$. By waiting until time \tilde{T} , agents take a chance. It is not difficult to see that the eventual belief p_T must strictly exceed α : if the deadline were not binding, each agent would prefer to procrastinate at instant \tilde{T} , given that all agents then exert maximal effort until the end.

Figure 4 also shows the effort level in the symmetric Markov equilibrium with observable effort in the presence of a deadline (a Markov strategy is now a function of the remaining time and the public belief). (The analysis of this case can be found in the Web Appendix.) With a long enough deadline, equilibrium effort in the observable case can be divided into three phases. Initially, effort is low and declining. Then, at some point, effort stops altogether. Finally, effort jumps back up to the maximal level. The last phase can be understood as in the unobservable case; the penultimate one is a stark manifestation of the incentives to procrastinate under observability. Note, however, that the time at which effort jumps back up to the maximal level is a function of the remaining time and the belief, and the latter depends on the history of effort thus far. As a result, this occurs earlier under observability than nonobservability. Therefore, effort levels between the two scenarios cannot be compared pointwise, although the belief as the deadline expires is higher in the observable case (i.e., aggregate effort exerted is lower).

FIGURE 4. OPTIMAL STRATEGIES GIVEN A DEADLINE OF $T = 3$

One might suspect that the extreme features of the equilibrium effort pattern in presence of a deadline are driven by the linearity in the cost function. While the model becomes intractable even with cost functions as simple as power functions, i.e., $c(u_i) = c \cdot u_i^\gamma$, $\gamma > 1$, $c > 0$, the equilibrium can be solved for numerically. Figure 5 shows how effort changes over time, as a function of the degree of convexity, both in the model with infinite horizon (left panel), and with a deadline (right panel). As is to be expected, when costs are nonlinear, effort is continuous over time in the presence of a deadline, although the encouragement effect that the deadline provides near the end is clearly discernible. Convex costs induce agents to smooth effort over time, an effect that can be abstracted away with linear cost.

B. The Optimal Deadline

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

THEOREM 3: *The optimal deadline T is finite and is given by*

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




FIGURE 5. EFFORT WITH CONVEX COST WITHOUT AND WITH A DEADLINE

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

It is clearly suboptimal to set a deadline shorter than T : agents' incentives to work would not change, leading only to an inefficient reduction of aggregate effort. It is more surprising that the optimal deadline is short enough to eliminate the first phase of the equilibrium altogether, as $\tilde{T} = 0$. However, if the deadline were longer by one second, the time \tilde{T} at which patient agents exert maximal effort would be delayed by more than one second. Indeed, agents initially exert effort at the interior level $u_{i,t}^*$. This means they grow more pessimistic and will require a shorter remaining time $T - \tilde{T}$ in order to have strict incentives to work. The induced increase in the cost of delay more than offsets the gains in terms of aggregate effort.

The optimal deadline T is decreasing in n , because it is the product of two positive and decreasing functions of n . That is, tighter deadlines need to be set when teams are larger. This is a consequence of the stronger incentives to shirk in larger teams. Furthermore, nT decreases in n as well. That is, the total amount of effort is lower in larger teams. However, it is easy to verify that the agent's payoff is increasing in the team size. Larger teams are bad in terms of overall efficiency, but good in terms of individual payoffs.

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

It is a familiar theme in the contracting literature that ex post inefficiencies can yield ex ante benefits. However, what is driving this result here is not smoothing over time per se, but free-riding. Indeed, the single agent solution coincides with the planner's problem, so that any deadline in that case would be detrimental.¹⁴

¹⁴Readers have pointed out to us that this effort pattern in the presence of a deadline reminds them of their own behavior as a single author. Sadly, it appears that such behavior is hopelessly suboptimal for $n = 1$. We refer to Ted O'Donoghue and Matthew Rabin (1999) for a behavioral model with which they might identify.

V. Extensions

The analysis can be extended in several directions. It is of interest to: (i) characterize asymmetric equilibria of the baseline model, and allow for asymmetries in the players' characteristics; (ii) allow for more general mechanisms than deadlines; (iii) allow for synergies across agents; (iv) allow for incomplete information concerning the players' productivities; (v) allow for learning-by-doing; and (vi) consider multiple-task projects.

Each of these extensions is examined in the working paper. We summarize here the main findings.

A. Asymmetric Equilibria or Asymmetric Players

While the symmetric equilibrium of the baseline model is unique, there exist other, asymmetric, equilibria. For instance, suppose one agent is expected to behave as if he were by himself, so that he exerts maximal effort up to T_1 , at which point his belief reaches the level α . All other agents then find it optimal to exert no effort whatsoever, provided they are sufficiently patient. More generally, every 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 $t_1 \leq t_2 \leq \dots \leq t_n$, with $t_1 \in [0, T_1]$, $t_k \in \mathbb{R}_+ \cup \{\infty\}$ for $k \geq 2$ (with $t_1 = T_1 \Rightarrow t_2 = \dots = t_n = T_1$, while $t_1 < T_1 \Rightarrow t_n = \infty$), such that agent i exerts maximal effort by himself 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.¹⁵ The symmetric equilibrium obtains for $t_1 = \dots = t_{n-1} = 0$.

The extreme asymmetric equilibrium in which one player does all the work (i.e., $t_1 = T_1$) might be viewed as unfair, but it is also the best equilibrium in terms of the sum of the agents' payoffs. Further, it emerges as the unique equilibrium for some variations of the model. This is the case if agents have different prior beliefs \bar{p}_i , if one agent has a higher value for success, or if agents have different costs α_i . Then the most optimistic, or the most productive, agent (say, agent i) must exert effort all by himself. The same reasoning applies in both cases. The asymptotic belief at which agent i would stop working is strictly lower than the other agents' corresponding thresholds, implying that he must ultimately work by himself. At this stage, he will do so efficiently, by exerting maximal effort. This implies, however, that there cannot be a last instant at which some other agent exerts effort, because j 's effort at that instant, and i 's maximal effort at the next instant, are strategic substitutes.

On the other hand, the symmetric equilibrium is uniquely selected by other variations of the model, in particular, when there are complementarities across the agents' productivities.

¹⁵ Conversely, for any such collection of subsets of agents and switching times, there exists an equilibrium.

B. More General Mechanisms

The setting of a deadline is a rather extreme way in which the team can affect incentives. As a first alternative, the team may be able to commit to a common wage that is an arbitrary function of time, subject to the constraint that the budget must balance on average: given that the project is worth $v = n$ overall to the team, the expected wages paid should not exceed this amount. Note that a deadline is a special case in which this wage is a step function (of time). It is possible to solve explicitly for the optimal wage schedule. The wage is not constant over time. Rather, it varies continuously over time up to a time \hat{T} at which it drops to zero (or to any sufficiently low level). This wage is such that agents are just willing to exert maximal effort at all times up to \hat{T} , at which time effort stops altogether and forever. For low discount rates, the wage is decreasing over time: frontloading payments allows to incentivize agents, because the shrinking wage counteracts the incentive to procrastinate. Frontloading cannot be achieved with a deadline; hence, maximal effort is sustained here over a longer horizon, as $\hat{T} > T$.

As a second possibility, we might consider the problem from the perspective of a principal who designs and commits to a symmetric wage scheme that the agent or agents receive, as a function of the time of a breakthrough. Instead of maximizing the agents' welfare subject to budget balance, his objective is to maximize his profits (assuming that he collects v in case of a breakthrough). These problems are dual, inasmuch as, in both cases, the objective is to minimize the cost of providing incentives for a given level of aggregate effort. As a result, the principal's wage scheme follows the same dynamics as the wage scheme that is optimal for the team. For both problems, it is optimal to induce maximal effort as cheaply as possible, as long as any effort is worthwhile, and to stop at some point. The date at which effort stops being worthwhile for the monopolist, T^* , occurs earlier than for the team: $T^* < \hat{T}$. Aggregate effort is decreasing in the number of agents, as is the payoff to the principal.

Figure 6 illustrates how the wage scheme varies with the parameters. Effort stops when the wage w_t is such that $p_t w_t = \alpha$, or $w_t = \alpha/p_t$. Note that this time 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. This means that the total amount of effort is independent of the discount rate, and so is efficiency.

As the discount rate tends to zero, the wage becomes an affine function. To provide incentives when agents are perfectly patient, the wage must decay at the rate of the marginal cost (or be constant when $n = 1$).

Note that this wage pattern is reminiscent of Mason and Välimäki (2008). In our environment, though, incentives to reduce effort over time are driven by learning, not by smoothing over time. In particular, unlike in their model, a very impatient team commands increasing wages over time, to compensate for the growing pessimism.

C. Synergies

There are several ways to model synergies. Agents might be more effective by working together than on their own, as is the case if

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




FIGURE 6. OPTIMAL WAGES IN THE PRINCIPAL-AGENT PROBLEM

so that agents' effort levels are complements. Alternatively, it might be that agents' skills differ, so that, depending on the type of project, both, one or the other, or neither agent can succeed (i.e., the project's type is such that, with probability \bar{p}^0 , both agents can succeed; with probability \bar{p}^i , only agent i can succeed, with $\bar{p}^0 + \bar{p}^1 + \bar{p}^2 < 1$). Both alternatives include the baseline model as a special case, but they yield different predictions.

When agents' efforts are complements, aggregate effort increases in the number of agents, because larger teams are more productive. As complementarities become strong enough ($\rho < 1/2$), the symmetric equilibrium emerges as the unique equilibrium, and aggregate effort tends to the efficient level as $\rho \rightarrow 0$. As in the baseline model, effort dwindles over time and remains always positive. There is wasteful delay, but not as much, for any given belief, as if effort were observable. If agents have different unit costs, it is no longer the case that the more productive agent works by himself. Initially, the more efficient agent works more than his counterpart, but the opposite may be true from some point on.

When agents have different, but symmetric, skills ($\bar{p}^1 = \bar{p}^2$), the unique equilibrium is symmetric, and exhibits the same features as in the baseline model: aggregate effort is inefficiently low, and provided inefficiently late. When $\bar{p}^i \rightarrow 0$, we are back to the baseline model, while when $\bar{p}^0 \rightarrow 0$, the agents' skills become independent. Yet free-riding persists, and aggregate effort is low, because each agent takes into account the lower probability that his skill is appropriate.

D. Incomplete Information

In many applications, uncertainty pertains not only to the quality of the project, but also to the productivity of agents. We may model this by assuming that $\alpha \in \{\alpha_L, \alpha_H\}$ is private information, and 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.

Thus, agents have beliefs about the quality of the project and the productivity of their opponent, and updating depends on the effort choice, which is private information. The following constitutes a symmetric equilibrium path. The low-type (high-productivity) agent starts by exerting effort by himself. As time passes, he becomes more pessimistic, while the high-type agent does not revise his belief downward as fast (the lack of success is not as surprising to him, because he does not work). That is, the private beliefs of the two agents' types about the quality of the project diverge. At some time, the relative optimism of the high-cost agent more than offsets his cost disadvantage, and he starts exerting effort. Simultaneously, the low-cost agent stops working once and for all. The high-cost agent's effort then dwindles over time and his belief converges asymptotically to α_H , his Marshallian threshold. Because of his initial effort level, the private belief of the low-cost agent remains forever below the high-cost agent's and converges asymptotically to a level below his own Marshallian threshold, namely α_L . The effort trajectories in such an equilibrium are shown in Figure 7.¹⁶

E. Learning by Doing

In practice, agents do not only learn from their past effort whether they can succeed, but also how they can succeed. Such learning by doing can be modeled as in Ulrich Doraszelski (2003), by assuming that each agent i accumulates knowledge according to

$$\dot{z}_{i,t} = u_{i,t} - \delta z_{i,t},$$

with $z_{i,0} = 0$. If the project is good, a breakthrough occurs with instantaneous probability $\sum_i h_{i,t}$, where

$$h_{i,t} = u_{i,t} + \rho z_{i,t}^\phi.$$

The baseline model obtains if we let $\delta \rightarrow \infty$, or $\rho \rightarrow 0$. While the first-order conditions given by Pontryagin's theorem cannot be solved in closed form, they can be solved numerically. It is no longer the case that effort is positive forever (at least, if ϕ is not too large). This should not be too surprising, because accumulated knowledge is a substitute for actual effort, so that it serves as its proxy once its stock is sufficiently large relative to the public belief. The probability of a breakthrough evolves

¹⁶Note that effort can be increasing over time. Because the probability attached by agent $-i$, to the event that agent i is of the low type, decreases over time, the effort agent $-i$ needs to exert to keep player i indifferent between exerting effort or not might actually increase.





FIGURE 7. EFFORT UNDER INCOMPLETE INFORMATION

as in the baseline model. It decreases over time, and remains always positive (which is obvious, since accumulated knowledge never fully depreciates). Effort decreases continuously and reaches zero in finite time. The asymptotic belief is now lower than α : although effort may (or may not) stop before this threshold is reached, the belief keeps decreasing afterward, because of the accumulated knowledge. Figure 8 depicts the locus of beliefs and knowledge stocks at which effort stops, and shows one possible path for the public belief, from $z_{i,0} = 0$ to the point at which all effort stops, for two possible values of ϕ . The dotted lines represent the evolution of p and z once effort has stopped. As one would expect, time until effort stops grows without bound as we approach the baseline model (i.e., if $\rho \rightarrow 0$ or $\delta \rightarrow \infty$).

F. Multiple Tasks

Most projects involve multiple stages, or multiple tasks. Depending on the application, it is possible to tackle these tasks simultaneously or not. Further, it might be that tasks are conjunctive (both need to be successfully completed), disjunctive (completing one is enough), or additive (completing both doubles the payoff). Obviously then, a thorough analysis of this problem warrants a paper on its own, and we shall settle for a few remarks here.

If tasks must all be completed, it is socially optimal to start with the one that is hardest (high cost α) or most unlikely to succeed (low prior belief \bar{p}). This is because it makes little sense to waste effort in completing an easy task while significant uncertainty lingers regarding the feasibility of completing the project. The opposite is true if tasks are disjunctive, as always working on the task with the larger mark-up $p - \alpha$ is optimal. (Once these mark-ups are equal, effort should be split so as to maintain this indifference.)

FIGURE 8. PUBLIC BELIEF AND ACCUMULATED KNOWLEDGE, AS A FUNCTION OF ϕ

Because the completion of tasks is observable, it is possible to have agents exert maximal effort, as long as they are assigned different tasks and tasks are conjunctive. In that case, each agent knows that completing his task is necessary for success, and that the other agent will not help him out with the task that he is assigned. On the other hand, the incentives to procrastinate are exacerbated when tasks are disjunctive, because an agent who shirks does not become more pessimistic regarding the possibility of succeeding with the task that he is responsible for.

With two agents, if tasks are sequential and conjunctive, in any equilibrium in which the last task is performed by a single agent exerting maximal effort by himself, it must be that agents alternate working on tasks. To see this, consider, for instance, the penultimate task. The agent who will idle once this task is completed has a strictly higher continuation payoff than the other agent. Therefore, the game involving the penultimate task can only be viewed as a game in which players have different characteristics (namely, a higher payoff from success for one of the agents), and as we have mentioned above, such a variant of the baseline model admits a unique equilibrium, in which the agent with the higher value for success exerts effort all by himself.

VI. Concluding Remarks

We have shown that moral hazard distorts not only the amount of effort, but also its timing. Agents work too little, too late.

Downsizing the team might help, provided that agents' skills are not too complementary. Setting an appropriate deadline is also beneficial, inasmuch as the reduction in delay more than offsets the further reduction in effort. Setting appropriate deadlines and choosing the right team size have long been discussed in the management literature. See, for instance, Tom DeMarco (1997) for an entertaining account on those aspects of project management.

Our finding that increasing transparency might aggravate the delay appears to be new, and evaluating its empirical relevance is an interesting and open problem. On the other hand, the importance of staggered prizes as a way of providing incentives, in the spirit of the optimal wage dynamics that we discussed in Section VB, is used

in practice. For instance, it is applied by the X Prize Foundation (see, for example, the Google Lunar X Prize).

Among the many features that are still missing from our model, perhaps the most important one involves extending our model to encompass multiple stages, and the assignment of tasks to specific team members, as a function of their perceived characteristics and ongoing performance. While Section V provides a few ideas in this respect, most of the work remains to be done.

APPENDIX

PROOF OF THEOREM 1:

Preliminaries: Observe first that, since $\dot{p}_t = -p_t(1 - p_t)\sum_i u_{i,t}$, we have $p_t\sum_i u_{i,t} = -\dot{p}_t/(1 - p_t)$. It follows that $p_t\sum_i u_{i,t} = d\log(1 - p_t)/dt$. We can then rewrite the discount factor $\exp(-\int_0^t (p_s\sum_i u_{i,s} + r)ds)$ in expression (3) as $\exp(-rt)(1 - \bar{p})/(1 - p_t)$, and the objective function as

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

where $u_{-i,t} := \sum_{j \neq i} u_{j,t}$. Applying integration by parts to the objective and ignoring all irrelevant terms (those that do not depend on u_i or x), we obtain

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

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

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

over functions $u_{i,t}$ in $[0, \lambda]$, given the function $u_{-i,t}$.

The Hamiltonian for this problem is

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

It is easy to see that no agent exerts effort if $p_t < \alpha$ (consider the original objective function: if $p_t < \alpha$, then choosing $u_{i,t} = 0$ is clearly optimal). We therefore assume that $\bar{p} > \alpha$, which is equivalent to $x_0 < \ln \beta$, where $x_0 := \ln((1 - \bar{p})/\bar{p})$. Assumption (4) on the discount rate is equivalent to $1 + e^{-x_0}(\lambda/r - \beta) > 0$.

Necessary Conditions: Define $\gamma_{i,t} := \hat{\gamma}_{i,t}e^{rt}$. By Pontryagin's principle, there must exist a continuous function γ_i such that, for each i ,

- (i) Maximum principle: For each $t \geq 0$, $u_{i,t}$ maximizes $\gamma_{i,t}(u_{i,t} + u_{-i,t})$;

- (ii) Evolution of the co-state variable: The function γ satisfies $\dot{\gamma}_{i,t} = r\gamma_i + 1 + e^{-x_t}(u_{-i,t}/r - \beta)$;
- (iii) Transversality condition: 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 transversality condition follows here from Takashi Kamihigashi (2001). Since there is a co-state variable γ_i for each player, we are led to consider a phase diagram in R^{n+1} , with dimensions representing $\gamma_1, \dots, \gamma_n$, 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

$$(8) \quad x_t^* = \ln \beta - \ln(1 + (\beta e^{-x_0} - 1)e^{-(n/(n-1))r\beta t}).$$

This implies $u_{i,t}^* = (r\beta/(n-1))/((\beta e^{-x_0} - 1)^{-1}e^{(n/(n-1))r\beta t} + 1)$, which corresponds to expression (7) 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:

- (i) 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.
- (ii) Consider paths that start with $\gamma_i \leq 0$ for all i , with strict inequality $\gamma_i < 0$ for all but one agent j . We then have $u_{-j} = 0$. Since $\bar{p} > \alpha$ implies that $r\gamma_j + 1 - \beta e^{-x_0} < 0$, it follows that $\dot{\gamma}_j < 0$, and we might as well assume that $\gamma_i < 0$ for all i . So, we have $u_i = 0$ for all i , and x remains constant, and all γ_i diverge to $-\infty$. Since x_0 is less than $\ln \beta$, the limit of our reference trajectory, this again violates the transversality condition. The same argument rules out any path that enters this subset of the state space, provided it does so for $x_t < \ln \beta$. However, we do not rule out the case of $\gamma = (\gamma_1, \dots, \gamma_n) \leq 0$ with two or more indices j s.t. $\gamma_j = 0$ and $u_{-j} > 0$.
- (iii) 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 such that $\gamma_j > 0$, 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.

- (iv) Consider paths that start with some $\gamma_i < 0$ for all $i \neq j$, and with $\gamma_j > 0$. Assume further that $r\gamma_j + 1 - \beta e^{-x_0} < 0$. Since $u_j > 0$ as long as $\gamma_j > 0$, the trajectory must eventually leave this subset of parameters, either because $\gamma_j \leq 0$, and then we are back to case 2, or because $\dot{\gamma}_j \geq 0$, and then we are back to case 3. If the trajectory enters one of the previous subsets, it is not admissible for the reasons above. Therefore, the only case left is if this trajectory hits $(\gamma_1, \dots, \gamma_n, x) \leq (0, \dots, 0, \ln \beta)$ with at least two indices j such that $\gamma_j = 0$. This is case 5. Notice that if there were more than one index j for which $\gamma_j > 0$, 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.
- (v) Consider paths that start with $(\gamma_1, \dots, \gamma_n, x) \leq (0, \dots, 0, \ln \beta)$ with at least two j such that $\gamma_j = 0$. Let $A_t = \{j: \gamma_{j,t} = 0\}$. Then there is a unique solution $u_{j,t}^*(|A|) := r(\beta - e^{x_t})/(|A| - 1)$, such that $\dot{\gamma}_j = 0$ for all $j \in A$ (as long as $x_t < \ln \beta$, since $u_{j,t}^* = 0$ when $x_t = \ln \beta$). Along this trajectory, $x_t \rightarrow \ln \beta$. Furthermore, the effort levels must switch to $u_{j,t}^*(|A_t| + 1)$ for all $j \in A \cup \{i\}$ whenever $\gamma_i = 0$ for $i \notin A_t$. It is similar 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 agent i exerting $u_i = \lambda$, then switching (before the beliefs have reached $\ln \beta$) to two or more agents (including i) who play the reference strategy $u_{i,t}^*(|A_t|)$, as if only agents $i \in A_t$ were present in the team. At any point in time before the beliefs have reached the threshold α , 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 exerts effort alone. In the subsequent phases, all (active) players exert effort at equal levels, adding new players at any point in time. Of course, there are extreme cases in which some phase is nonexistent. Therefore, the only symmetric equilibrium is one in which $|A_0| = n$, that is, all players exert effort $u_j^*(n)$ from the start.

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

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

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

PROOF OF LEMMA 1:

- (i) From expression (7), it is clear that individual effort is decreasing in t , and that for a fixed t , $u_{i,t}^*$ is increasing \bar{p} .
- (ii) Aggregate effort is measured by x_t^* , since we know $\dot{x}_t = \sum_i u_{i,t}$. Differentiating expression (8), it follows that the equilibrium x_t^* is decreasing in α and in n , and that $\lim_{t \rightarrow \infty} x_t^* = \ln \beta$ for all n .

Given the equilibrium strategies, the probability of a success occurring is given by

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

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

$$G_n(t) := \int_0^{\tilde{s}(t,n)} f(s) ds \Big/ \int_0^\infty f(s) ds,$$

where $\tilde{s}(t, n) := nr\beta t/(n-1)$. Since \tilde{s} is decreasing in n , the probability of a success arriving before time t is also decreasing in n . In other words, the conditional distributions of arrival times $G_n(t)$ are ranked by first-order stochastic dominance. As a consequence, the conditional expected time of a breakthrough is increasing in n .

- (iii) Substituting expressions (8) and (7) for x_t^* and $u_{i,t}^*$, the equilibrium payoffs in (3) can be written as

$$V = \frac{r^2 k \beta}{(1 + \beta)(\beta + 1 + k)} \int_0^\infty \left(1 + \frac{\beta}{1 + ke^{-\frac{nrt\beta}{n-1}}} + \frac{\beta kn/(n-1)}{k + e^{\frac{nrt\beta}{n-1}}} \right) e^{-\frac{nrt\beta}{n-1} - rt} dt,$$

where $k := (\beta e^{-x_0} - 1)$. The change of variable $y = \exp(-(n/(n-1))r\beta t)$ allows us to write the payoff as

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

which is increasing in n , since $y \in [0, 1]$ implies the integrand is decreasing in n . Furthermore, the unnormalized payoff is independent of r . The other comparative statics follow upon differentiation of V .

PROOF OF THEOREM 2:

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

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

This term is positive. To see this, 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 $(1-\alpha)^{-1} > 0$. Thus, f is increasing in p over the range $[\alpha, 1]$, and it is equal to zero at $p = \alpha$, so it is positive over this range.

REFERENCES

- Admati, Anat R., and Motty Perry. 1991. "Joint Projects without Commitment." *Review of Economic Studies*, 58(2): 259–76.
- Alchian, Armen A., and Harold Demsetz. 1972. "Production, Information Costs, and Economic Organization." *American Economic Review*, 62(5): 777–95.
- Bag, Parimal Kanti, and Santanu Roy. 2008. "Repeated Charitable Contributions under Incomplete Information." *Economic Journal*, 118(525): 60–91.
- Bergemann, Dirk, and Ulrich Hege. 2005. "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics*, 36(4): 719–52.
- Bergin, James, and W. Bentley MacLeod. 1993. "Continuous Time Repeated Games." *International Economic Review*, 34(1): 21–37.
- Bliss, Christopher, and Barry Nalebuff. 1984. "Dragon-Slaying and Ballroom Dancing: The Private Supply of a Public Good." *Journal of Public Economics*, 25(1–2): 1–12.
- Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–74.
- Bonatti, Alessandro, and Johannes Hörner. 2009. "Collaborating." Cowles Foundation Discussion Paper 1695.
- Compte, Olivier, and Philippe Jehiel. 2004. "Gradualism in Bargaining and Contribution Games." *Review of Economic Studies*, 71(4): 975–1000.
- DeMarco, Tom. 1997. *The Deadline: A Novel about Project Management*. New York: Dorset House.
- Doraszelski, Ulrich. 2003. "An R&D Race with Knowledge Accumulation." *RAND Journal of Economics*, 34(1): 20–42.
- Doz, Yves L. 1996. "The Evolution of Cooperation in Strategic Alliances: Initial Conditions or Learning Processes?" *Strategic Management Journal*, 17(S1): 55–83.
- Doz, Yves L., and Gary Hamel. 1998. *Alliance Advantage: The Art of Creating Value through Partnering*. Boston: Harvard Business School Press.
- Fershtman, Chaim, and Shmuel Nitzan. 1991. "Dynamic Voluntary Provision of Public Goods." *European Economic Review*, 35(5): 1057–67.
- Gong, Yaping, Oded Shenkar, Yadong Luo, and Mee-Kau Nyaw. 2007. "Do Multiple Parents Help or Hinder International Joint Venture Performance? The Mediating Roles of Contract Completeness and Partner Cooperation." *Strategic Management Journal*, 28(10): 1021–34.
- Gradstein, Mark. 1992. "Time Dynamics and Incomplete Information in the Private Provision of Public Goods." *Journal of Political Economy*, 100(3): 581–97.
- Gulati, Ranjay. 1995. "Social Structure and Alliance Formation Patterns: A Longitudinal Analysis." *Administrative Science Quarterly*, 40(4): 619–52.
- Gulati, Ranjay, and Harbir Singh. 1998. "The Architecture of Cooperation: Managing Coordination Costs and Appropriation Concerns in Strategic Alliances." *Administrative Science Quarterly*, 43(4): 781–814.
- Hamilton, Barton H., Jack A. Nickerson, and Hideo 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): 465–97.
- Hergert, Michael, and Deigan Morris. 1988. "Trends in International Collaborative Agreements." In *Cooperative Strategies in International Business*, ed. Farok J. Contractor and Peter Lorange, 99–109. Lexington, MA: Lexington.
- Holmström, Bengt. 1982. "Moral Hazard in Teams." *Bell Journal of Economics*, 13(2): 324–40.

- Hopenhayn, Hugo, and Francesco Squintani.** 2004. "Preemption Games with Private Information." ESRC Centre for Economic Learning and Social Evolution Working Paper 108.
- Kamihigashi, Takashi.** 2001. "Necessity of Transversality Conditions for Infinite Horizon Problems." *Econometrica*, 69(4): 995–1012.
- Karau, Steven J., and Kipling D. Williams.** 1993. "Social Loafing: A Meta-Analytic Review and Theoretical Integration." *Journal of Personality and Social Psychology*, 65(4): 681–706.
- Keller, Godfrey, Sven Rady, and Martin Cripps.** 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73(1): 39–68.
- Klein, Nicolas, and Sven Rady.** 2008. "Negatively Correlated Bandits." Center for Economic and Policy Research Discussion Paper 6983.
- Latané, Bibb, Kipling Williams, and Stephen Harkins.** 1979. "Many Hands Make Light the Work: The Causes and Consequences of Social Loafing." *Journal of Personality and Social Psychology*, 37(6): 822–32.
- Legros, Patrick, and Steven A. Matthews.** 1993. "Efficient and Nearly-Efficient Partnerships." *Review of Economic Studies*, 60(3): 599–611.
- Levi, Daniel.** 2007. *Group Dynamics for Teams*. Thousand Oaks, CA: Sage Publications.
- Levin, Jonathan, and Steven Tadelis.** 2005. "Profit Sharing and the Role of Professional Partnerships." *Quarterly Journal of Economics*, 120(1): 131–71.
- Lewis, Tracy R., and Marco Ottaviani.** 2008. "Search Agency." <http://www.kellogg.northwestern.edu/Faculty/ottaviani/homepage/Papers/Search%20Agency.pdf>.
- Lockwood, Ben, and Jonathan P. Thomas.** 2002. "Gradualism and Irreversibility." *Review of Economic Studies*, 69(2): 339–56.
- Luo, Yadong.** 2002. "Contract, Cooperation, and Performance in International Joint Ventures." *Strategic Management Journal*, 23(10): 903–19.
- Ma, Ching-To, John Moore, and Stephen Turnbull.** 1988. "Stopping Agents from 'Cheating.'" *Journal of Economic Theory*, 46(2): 355–72.
- Madhok, Anoop.** 1995. "Revisiting Multinational Firms' Tolerance for Joint Ventures: A Trust-Based Approach." *Journal of International Business Studies*, 26(1): 117–37.
- Marx, Leslie M., and Steven A. Matthews.** 2000. "Dynamic Voluntary Contribution to a Public Project." *Review of Economic Studies*, 67(2): 327–58.
- Mason, Robin, and Juuso Välimäki.** 2008. "Dynamic Moral Hazard and Project Completion." Center for Economic and Policy Research Discussion Paper 6857.
- McGrath, Joseph E.** 1991. "Time, Interaction, and Performance (TIP): A Theory of Groups." *Small Group Research*, 22(2): 147–74.
- Murto, Pauli, and Juuso Välimäki.** 2008. "Learning and Information Aggregation in an Exit Game." Helsinki Center for Economic Research Discussion Paper 235.
- O'Donoghue, Ted, and Matthew Rabin.** 1999. "Incentives for Procrastinators." *Quarterly Journal of Economics*, 114(3): 769–816.
- Olson, Mancur.** 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Presman, Ernst L.** 1990. "Poisson Version of the Two-Armed Bandit Problem with Discounting." *Theory of Probability and Its Applications*, 35(2): 307–17.
- Presman, Ernst L., and Isaac N. Sonin.** 1990. *Sequential Control with Incomplete Information: The Bayesian Approach to Multi-Armed Bandit Problems*. Translated and edited by E. A. Medova-Dempster and M. A. H. Dempster. London: Academic Press.
- Radner, Roy, Roger Myerson, and Eric Maskin.** 1986. "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria." *Review of Economic Studies*, 53(1): 59–69.
- Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille.** 2007. "Social Learning in One-Arm Bandit Problems." *Econometrica*, 75(6): 1591–1611.
- Seierstad, Atle, and Knut Sydsæter.** 1987. *Optimal Control Theory with Economic Applications*. New York: Elsevier Science.
- Strausz, Roland.** 1999. "Efficiency in Sequential Partnerships." *Journal of Economic Theory*, 85(1): 140–56.
- Toxvaerd, Flavio.** 2007. "A Theory of Optimal Deadlines." *Journal of Economic Dynamics and Control*, 31(2): 493–513.
- Tuckman, Bruce W., and Mary Ann C. Jensen.** 1977. "Stages of Small-Group Development Revisited." *Group Organization Management*, 2(4): 419–27.
- Williams, Steven R., and Roy Radner.** 1995. "Efficiency in Partnership when the Joint Output Is Uncertain." In *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability*, ed. John O. Ledyard, 79–99. Boston: Kluwer Academic.
- Winter, Eyal.** 2004. "Incentives and Discrimination." *American Economic Review*, 94(3): 764–73.
- Wuchty, Stefan, Benjamin F. Jones, and Brian Uzzi.** 2007. "The Increasing Dominance of Teams in Production of Knowledge." *Science*, 316(5827): 1036–39.