UC Berkeley

UC Berkeley Previously Published Works

Title

Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects

Permalink

<https://escholarship.org/uc/item/3j62n213>

Journal

American Economic Review, 106(12)

ISSN

0002-8282

Authors

Green, B
Taylor, CR

Publication Date

2016-12-01

DOI

10.1257/aer.20151181

Peer reviewed

Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects*

Brett Green

UC Berkeley (Haas)
greenb@berkeley.edu
faculty.haas.berkeley.edu/bggreen

Curtis R. Taylor

Duke University
crtaylor@duke.edu
people.duke.edu/~crtaylor

February 14, 2016

Abstract

We study the optimal incentive scheme for a multistage project in which the agent privately observes intermediate progress. The optimal contract involves a *soft deadline* wherein the principal guarantees funding up to a certain date – if the agent reports progress at that date, then the principal gives him a relatively short *hard deadline* to complete the project – if progress is not reported at that date, then a probationary phase begins in which the project is randomly terminated at a constant rate until progress is reported. Self-reported progress plays a crucial (but non-stationary) role in implementation. We explore several variants of the model with implications for optimal project design. In particular, we show that the principal benefits by imposing a small cost on the agent in order to submit a progress report or by making the first stage of the project somewhat “harder” than the second. On the other hand, the principal does strictly worse by impairing the agent’s ability to observe his own progress.

JEL Classification: J41, L14, M55, D82

Keywords: Dynamic contract, Progress Report, Soft Deadline, Milestones, R&D

*We are grateful to Bruno Biais, Simon Board, Alessandro Bonatti, Barney Hartman-Glaser, R. Vijay Krishna, John Morgan, Sofia Moroni, Steve Tadelis, Andy Skrzypacz, Toby Stuart, and Felipe Varas for useful comments and suggestions. We thank conference participants at the Columbia/Duke/MIT/Northwestern IO Theory Conference, ESSFM Gersenzee, the Finance Theory Group (FTG) Workshop, Stanford Institute of Theoretical Economics (SITE), WFA-CFAR Conference on Corporate Finance, and seminar participants at ASU (Carey), HEC Paris, Humboldt University, MIT, Toulouse School of Economics, University of Bonn, University of Chicago, University of Miami, and UNC Kenan-Flagler for useful feedback. Andy Schwartz and Andrew Steck provided excellent research assistance.

1 Introduction

In this paper we analyze the following situation. A principal contracts with an agent to complete a project. The project requires the successful completion of two stages or *breakthroughs* in order to realize its benefits. The arrival rate of breakthroughs depends on the agent’s hidden action. In particular, the agent can secretly divert funds for private benefit, thereby reducing the rate of a breakthrough. The second breakthrough (i.e., the date at which the project is completed) is publicly observed, but the first breakthrough is privately observed by the agent and cannot be verified by the principal. Both players are risk neutral and the agent is protected by limited liability. We are interested in how the principal should optimally design the incentive scheme and to what extent it relies on communication from the agent to the principal about whether progress has been made.

Our investigation of this setting is motivated by three key ingredients often encountered in complex real-world projects involving research or implementation of novel technical designs. First, such projects are typically organized under an agency relationship because they require both a high level of technical expertise as well as substantial capital. Thus, there is generally some degree of separation between the expert individuals responsible for execution of a venture and the financial entity that provides its backing. Moreover, the preferences of the entrepreneur, contractor, or researcher (the “agent”) regarding the timing, intensity and direction of investment are unlikely to be perfectly aligned with the preferences of the financier, end user, or institution (the “principal”). For example, [Tirole \(2006\)](#) suggests several reasons why the relationship between researchers and their funding sources “is fraught with moral hazard.”

Second, complex projects often require completion of multiple sequential stages before their benefits can be realized. For example, developing a new drug requires identifying specific chemical compounds or molecules in vitro (i.e., in test tubes) and then demonstrating their efficacy in pre-clinical (i.e., animal) trials. After this phase is successfully completed, clinical (i.e., human) trials begin, progressing through small scale and then large-scale. Each of these stages must be successfully completed before a drug can be submitted to the FDA for approval, and only after approval can the drug be brought to market, generating revenue for its developer and health benefits for society. In a similar vein, large-scale software engineering projects are often purported to follow the celebrated “waterfall model” of development involving six sequential phases: conception, initiation, analysis, design, construction, testing, and implementation ([Royce, 1970](#)). Indeed, in 1985 the United States Department of Defense codified six similar phases in their standards for working with software development contractors (DOD-STD-2167A). Other examples of multistage projects are ubiquitous. For

example, most large-scale procurement projects (e.g., construction, defense) involve numerous sequential stages. In venture capital, entrepreneurs generally progress through numerous stages (e.g., patenting, prototype development, and manufacturing) before realizing profits. Basic research, almost by definition, requires numerous successive advancements before its societal benefits are realized.

Third, when and whether the agent makes progress on the project – that is, when and whether it transitions from one phase to the next – is often difficult, even impossible, for the principal to ascertain directly. This could obtain either because the principal lacks the technical proficiency to evaluate progress or because it is not possible to substantiate progress at a reasonable cost. In either case, project sponsors are often forced to rely on unverifiable progress reports made by the very individuals responsible for moving the project forward. Detecting fraudulent claims of progress in complex environments is, not surprisingly, notoriously difficult. Take, for example, the respective cases of extensive data falsification by Diederik Stapel, formerly a professor of social psychology at Tilburg University, and Anil Potti, formerly a cancer researcher at Duke University. The misconduct of both academics went undetected by their universities and funding agencies for a number of years. Similarly, in a 2006 survey of software practitioners, 86% of respondents reported having encountered false reports (Glass et al., 2008). The most common occurrences were in estimation and status reporting. “Respondents said that when lying happens, developers at the bottom level of the management hierarchy are most aware of the lying; they often know it’s happening even when their management doesn’t.”

These three common ingredients: agency, multiple stages, and intangible progress, give rise to the question at the heart of this paper. Specifically, if the principal cannot observe evolution of the project herself, can she nevertheless use unverifiable reports from the agent to monitor his progress and provide incentives for project advancement? Given the prevalence of misreporting, one might easily imagine that the answer to this question is ‘no.’ In fact, we find that self-reported progress is an essential component of an optimally designed incentive scheme for promoting project development. To ensure the veracity of progress reports, the principal must, however, use them judiciously.

To build intuition and a baseline of comparison for our main results, we start by analyzing a single-stage project in which only one breakthrough is needed to complete the project. In this case, the optimal incentive scheme can be implemented with a *simple contract*, which involves a single deadline T^* and a reward that depends only on the date at which the breakthrough arrives. If the agent realizes the breakthrough at $\tau \leq T^*$, he collects the reward of $R(\tau)$. If a breakthrough is not realized by the deadline, the principal terminates the project. While the first-best policy never involves project termination, the use of a deadline plays a

crucial role in the provision of incentives in our (second-best) setting. In choosing the optimal deadline the principal faces a trade off; a longer deadline yields a higher probability of project completion but also requires giving the agent more rents in order to prevent shirking.

Having established this single-stage benchmark, we analyze the situation of interest by adding a second stage to the project. Thus, there is a preliminary (e.g., research) stage, which must be successfully completed before moving on to the ultimate (e.g., development) stage. A breakthrough in the preliminary stage does not generate any direct benefits, is privately observed by the agent, and is unverifiable. On the other hand, project completion (i.e., when the second breakthrough has been made) is verifiable. In this setting, simple contracts are no longer optimal. If the principal uses a simple contract (with deadline T^*) then as the deadline approaches, an agent who has not yet made a breakthrough will “run out of steam” at some point $T_1 < T^*$ and begin shirking. Thus, a simple contract can be improved upon. If the principal could observe the project stage directly, then she would simply fire the agent for lack of progress at T_1 , thereby averting the cost of shirking *ex post* and enhancing incentives *ex ante*. When progress is intangible, however, the principal cannot threaten to fire the agent for lack of reported progress at T_1 without inducing him to make false reports.

Instead, the optimal contract involves a first-stage deadline that is not deterministic; funding is guaranteed up to some *soft deadline*. If the agent has not reported progress by the soft deadline, then a probationary phase ensues in which the principal randomizes over whether or not to terminate the project.¹ The soft deadline provides incentives for the agent not to shirk or falsely report progress in the first stage while not further reducing the probability of success in the second stage conditional on making a breakthrough. Moreover, an agent who has made progress strictly prefers to report truthfully at the soft deadline rather than risk being terminated while under probation. In short, the lottery over termination dates serves as a mechanism to induce both effort and truthful reporting while simultaneously maximizing the probability of project success.

Self-reported progress from the agent to the principal regarding the status of the project, therefore, plays an essential role in the optimal incentive scheme. Interestingly, the required communication is non-stationary over the duration of the project. Early in the life of the project, communication is unnecessary. If the project is completed before the soft deadline, then the agent’s compensation depends solely on the date of completion. If the project is not completed by the soft deadline, communication is required. More specifically, at the soft deadline the principal asks the agent to report whether he has made a breakthrough. If he answers “yes” then the principal gives him a relatively short hard deadline to complete the

¹The principal need not randomize over terminating the project at every instant. Rather, she may randomly and secretly draw a fixed date at which she will end probation and fire the agent for lack of progress.

project—if the project is not completed by the hard deadline then it is terminated. If the agent answers “no” then the probationary phase begins. During the probationary phase, the agent is asked to report any progress *immediately*, at which point he is given the same amount of time to complete the project as if he had reported “yes” at the soft deadline.

After characterizing the optimal contract and the set of equilibrium payoffs, we explore three extensions of the model with implications for optimal project design. First, we ask whether there is scope for making communication costly. We show that the principal can indeed benefit by imposing a small cost on the agent of reporting progress. Therefore, formal channels of communication that require time and effort (e.g., paperwork) can be useful even if the same information could be communicated at no cost. Second, we consider a project with asymmetric stages and show that *ceteris paribus* the principal’s payoff is higher when the first stage is moderately more difficult than the second stage. However, the principal does strictly worse by making the first stage too difficult relative to the second. Finally we consider a setting in which progress is unobservable to both players in order to see whether the principal can benefit from suppressing the agent’s access to information about the status of the project. We find that information suppression is suboptimal; the principal does better under the optimal contract with progress reports from the agent than in a setting where neither party observes the time of the first breakthrough. An alternative interpretation of this result is that it is better to divide up a project into two phases even if the agent privately observes when the first is completed.

We believe our findings are not only of theoretical interest, but also have practical relevance. For example, making continued funding contingent on periodic progress reports appears to be a common arrangement. The following text that appears on the website of the Amyotrophic Lateral Sclerosis (ALS) Association is broadly representative.

The ALS Association financial officer makes grant award payments to the Principal Investigator institution for disbursement for the project. Payments are made on a specified quarterly schedule and in the case of multi-year grants, after the first year, are contingent upon the receipt by ALS Association of satisfactory annual progress reports and documentation of research funds expended.

Regarding the use of probation, funding entities do not explicitly specify random deadlines (at least not to our knowledge). However, grant policies are often quite vague about the consequences for delays in reporting progress. For instance, the National Institute of Health’s (NIH) web site states “If your [progress] report is extremely late, *you risk losing funding* for the period of time between the end of the current budget period and when we finish

processing your report.”² Similarly, the National Science Foundations (NSF) Grant Policy Manual states, “NSF *reserves the right, ... to withhold future payments* after a specified date if the recipient fails to comply with the conditions of an NSF grant, including the reporting requirements.”³ While these policies do not specify random termination as such, it certainly seems more appropriate to view the indefinite penalties for late reporting as involving *soft* deadlines rather than *hard* ones.

The soft deadline and ensuing probationary phase of our optimal incentive scheme also resembles how “project slippage,” (i.e., missing a deadline) is handled in project management (e.g., [Ewusi-Mensah and Przasnyski, 1991](#)). In practice, “slippage time” (i.e., the amount of time the project is in a state of slippage) generally extends the deadline for completing the entire project. That is, contractors and employees are not generally required to recover slippage time in one phase of a project by shortening subsequent phases. Of course, the greater the slippage time, the more likely it is that the sponsor will cancel it. Our results suggest that some degree of slippage should, however, be tolerated on complex projects where sponsors are unable to assess progress directly.

The remainder of the paper proceeds as follows. In the next section, we discuss related literature. In Section 3, we present a single-stage version of the model as a benchmark. We introduce the second stage and present some preliminary findings in Section 4. We fully characterize the optimal contract and set of equilibrium payoffs in Section 5. Section 6 explores the role of costly reporting, projects with asymmetric stages, and the possibility of restricting the agent’s access to information. Concluding remarks appear in Section 7. Most proofs and several technical lemmas are located in the Appendix. Several proofs have been relegated to an online appendix.

2 Related Literature

There is a large and growing literature studying the optimal provision of incentives in dynamic environments.⁴ Our benchmark single-stage model is similar to [Shavell and Weiss \(1979\)](#) and [Hopenhayn and Nicolini \(1997\)](#),⁵ who looks at providing incentives to search for employment while simultaneously providing unemployment insurance. It is also similar to [Mason and Välimäki \(2015\)](#) who consider a dynamic moral hazard setting in which the principal must

²<http://www.niaid.nih.gov/researchfunding/qa/pages/pp.aspx#late>: emphasis added.

³http://www.nsf.gov/pubs/manuals/gpm05_131/gpm4.jsp: emphasis added.

⁴A non-exhaustive list includes [Green \(1987\)](#), [Spear and Srivastava \(1987\)](#), [Phelan and Townsend \(1991\)](#), [Quadrini \(2004\)](#), [Clementi and Hopenhayn \(2006\)](#), [DeMarzo and Sannikov \(2006\)](#), [DeMarzo and Fishman \(2007\)](#) and [Sannikov \(2008\)](#).

⁵See also [Lewis \(2012\)](#) who considers a delegated search model in which the optimal contract includes a deadline and a bonus for early completion.

provide incentives to the agent to complete a project.⁶ A novel aspect of our model is that we explicitly consider environments with multiple sequential stages in which the agent privately observes progress. Thus, we “sandwich” a hidden-information problem between two stages of a hidden-action problem.

Our multistage environment is related to [Biais et al. \(2010\)](#), who analyze a model in which large (observable) losses may arrive via a Poisson process, and an agent must exert unobservable effort in order to minimize the likelihood of their arrival. They allow for investment and characterize firm dynamics as well as asymptotic properties. Our model differs in that it (i) features only a finite number of arrivals, (ii) the arrival of a breakthrough is “good news”, and (iii) we consider the case in which arrivals are unobservable to the principal. Several other recent papers that involve observable Poisson arrivals include [Hoffmann and Pfeil \(2010\)](#), [Piskorski and Tchistyi \(2011\)](#), [DeMarzo et al. \(2014\)](#). Given the multistage setting, a key difference in this paper is that the agent’s continuation utility is not a sufficient state variable. [Toxvaerd \(2006\)](#) considers a setting in which a finite number of (observable) arrivals are needed in order to complete a project. In his setting, the agent is risk averse and the optimal contract trades off optimal risk-sharing for incentive provision, but does not involve deadlines or inefficient termination. Optimal dynamic mechanisms are explored in other settings by [Board \(2007\)](#), [Eso and Szentes \(2007\)](#), [Bergemann and Valimäki \(2010\)](#) and [Pavan et al. \(2014\)](#) among others.

In our model, the agent has access to private information that is persistent. Dynamic contracting with persistent private information has been studied in discrete type settings by [Fernandes and Phelan \(2000\)](#), [Battaglini \(2005\)](#), [Tchistyi \(2013\)](#), and with a continuum of types using a first order approach by [Williams \(2011\)](#) and [Edmans et al. \(2012\)](#). Our approach is most similar to [Zhang \(2009\)](#) and [Guo and Hörner \(2015\)](#). From a theoretical perspective, our work differs from this literature along several dimensions. One key difference in our environment is the presence of public (contractible) information (i.e., the ultimate success of the project), which can be used to screen the agent’s underlying private information. Another difference is that the transition probabilities across states are endogenously determined in our setting by the agent’s action.

In a working paper, [Hu \(2014\)](#) considers a discrete-time setting that is otherwise similar to ours. However, he restricts attention to deterministic deadlines and does not allow for communication between the principal and the agent, which leads him to conclude that the

⁶Though the problems are similar in spirit, the optimal contract in our one-stage benchmark looks quite different from [Mason and Välimäki \(2015\)](#). In our model, it is optimal to use termination deadlines to provide incentives to the agent, whereas in their setting, the use of termination is strictly suboptimal. This difference arises due to the nature of the moral hazard problem we consider (e.g., private benefit from shirking) rather than their costly effort model. See Remark 1 for further discussion.

optimal contract involves a single deadline in which an agent who has not made progress shirks as the deadline approaches. Our results show these restrictions are not without loss of generality and lead to substantively different findings. In particular, we show that without these restrictions, the optimal mechanism requires communication, involves random termination, and does not involve shirking.

There is a rich existing literature exploring settings where parties *learn* about the value of a project over time.⁷ In these environments, lack of success typically indicates that the project is *bad*, and it is socially efficient to discontinue investment at some point. A common finding within this literature is that agency considerations may cause the principal to terminate the project earlier than socially optimal. By contrast, we focus on a setting in which the project is commonly known to be *good* from the outset (i.e., there is no learning and the project is always funded until completion under the first-best contract) in order to isolate the extent to which progress can be used to provide stronger incentives. [Bonatti and Hörner \(2011\)](#) study experimentation in teams for a project that requires a single breakthrough, which introduces a free-riding problem. They show that the equilibria of the game involve inefficient delays in effort provision and that deadlines, which terminate the project prior to the socially efficient time, are useful in mitigating delays despite forfeiting value when the deadline is reached. The free-riding problem also arises in [Moroni \(2015\)](#), who studies experimentation with multiple agents for a project that requires several observable breakthroughs.

[Holmstrom and Milgrom \(1991\)](#) and [Laux \(2001\)](#) look at settings with simultaneous tasks whereas in our setting the project involves stages that must be completed sequentially. [Lerner and Malmendier \(2010\)](#) investigate the role of property rights and contractibility in the design of research agreements within a multi-task setting and empirically document more prevalent use of termination options in settings where effort is not contractible. [Varas \(2015\)](#) studies a dynamic multi-tasking model in which an agent can complete the project faster by reducing its quality. Under certain conditions, he finds that a random termination policy is optimal. However, the underlying mechanism is somewhat different from ours. In his model, stochastic termination is used to prevent multi-tasking and arises as part of the optimal contract only when the agent is sufficiently more impatient than the principal. Moreover the use of a hard deadline is suboptimal in his model unlike in our multistage setting where the principal uses a combination of both soft and hard deadlines.

⁷See e.g., [Levitt and Snyder \(1997\)](#), [Bergemann and Hege \(1998, 2005\)](#), [Inderst and Mueller \(2010\)](#), [Manso \(2011\)](#), [Hörner and Samuelson \(2014\)](#), [Halac et al. \(2015\)](#).

3 A single-stage benchmark

A principal (she) contracts with an agent (he) to undertake a project. Time is continuous and the project can be operated over a potentially infinite horizon. The project requires the successful completion of a stage, also referred to as a *breakthrough*, in order for its benefits to be realized. Operating the project requires resources, which we model as a flow cost of c per unit time that the project is in operation. The principal has unlimited resources to fund the project. The agent has no funds and is protected by limited liability, but has the skills necessary to run the project. Both parties are risk-neutral and do not discount the future.⁸ The principal can terminate the project (i.e., discontinue paying the flow cost) at any point in time. Project termination is irreversible; if the principal terminates the project prior to the breakthrough, the project delivers no benefit and the game ends.

While the project is in operation, the agent chooses an action $a_t \in \{0, 1\}$, where $a_t = 1$ indicates that the agent appropriately invests or “works,” and $a_t = 0$ indicates that the agent diverts funds or “shirks” for private benefit. The arrival rate of a breakthrough is then given by λa_t . Thus, if the agent works over an interval of length dt , then the probability of a breakthrough in the interval is λdt . If the agent shirks, then the arrival rate of a breakthrough is zero but he receives a private flow benefit of ϕdt , where $\phi > 0$ measures the severity of the agency problem. The agent receives no intrinsic benefit from project success; he benefits solely from the compensation delivered by the principal and any private benefits from shirking.

Remark 1. *For many relevant applications, the most natural interpretation of the moral hazard problem is that the agent can secretly divert the principal’s investment for private benefit. For example, an entrepreneur can use venture capital funding for private consumption, or a scientist may fund a pet project not authorized under his current grant. Nevertheless, we will adopt the standard “shirk”/“work” terminology (e.g., [Tirole, 2006](#)).⁹*

The success of the project is publicly observed and contractible. Upon the arrival of a breakthrough the principal realizes a payoff $\Pi > 0$, makes any outstanding contractual payments to the agent and the game ends. Let τ denote the random variable representing the date of project success. Throughout our analysis, we employ the following assumptions.

⁸It is fairly straightforward to incorporate a common discount rate into the model. However, with discounting, closed-form solutions are no longer available in the multistage setting and therefore our proof method for some results does not generalize directly. In these situations, numerical examples support our findings.

⁹In the analog of our model where the agent incurs a cost of effort rather than a benefit from diversion (i.e., arrival rate of breakthroughs is zero without effort), the limited liability constraint has no bite and the first-best outcome is attainable. In a costly-effort model with discounting and a strictly positive arrival rate even when the agent shirks, it is possible to obtain results similar to those presented here under certain parametric restrictions.

Assumption 1. *The expected value of the project (absent agency costs) is strictly positive*

$$\lambda\Pi - c > 0.$$

Assumption 2. *Shirking is non-trivial and inefficient*

$$0 < \phi \leq c.$$

Remark 2. *Assumption 1 and 2 imply that under the first-best policy, the agent never shirks and the project is never terminated prior to success.*

At $t = 0$, the principal offers the agent a contract. We assume that the principal can fully commit to all terms. If the agent rejects the offer, then both parties receive their outside options normalized to zero.

As is typical of most principal-agent models, making a payment to the agent upon “failure” (i.e., termination prior to project success) is suboptimal. Moreover, because both parties have linear utility and are equally patient, it is without loss of generality to backload all monetary payments to the agent (see e.g., [Ray, 2002](#)). A contract can thus be denoted by a triple, $\Gamma \equiv \{a, R, T\}$, where a_t is the recommended action to the agent at time t , R_t is a monetary payment made to the agent for success at t , and T is the date at which the project is terminated absent a prior breakthrough.¹⁰ An action process, a , induces a probability distribution $\mathbb{P}^a$ over τ . Let $\mathbb{E}^a$ denote the corresponding expectation operator. Assuming for the moment that the agent always adheres to the recommended action, the principal’s (ex-ante) expected utility under any contract is given by

$$P_0(\Gamma) = \mathbb{E}^a \left[(\Pi - R_\tau) \cdot \mathbf{1}_{\{\tau \leq T\}} - \int_0^{T \wedge \tau} c \, dt \right], \quad (1)$$

and the agent’s expected utility is given by

$$U_0(\Gamma) = \mathbb{E}^a \left[R_\tau \cdot \mathbf{1}_{\{\tau \leq T\}} + \int_0^{T \wedge \tau} (\phi(1 - a_t)) \, dt \right]. \quad (2)$$

The contract Γ is said to be incentive compatible if a maximizes the agent’s expected utility (2) given (R, T) . The principal’s problem is to find an incentive compatible contract that maximizes (1) subject to delivering an expected utility to the agent of at least his outside option.

¹⁰In the multistage setting, the solution to the principal’s problem requires randomization over termination dates. However, randomization is unnecessary for a single-stage project so we delay introducing this possibility and additional notation until it is needed.

Given any contract Γ , let $U_t(\Gamma) \in \mathbb{R}_+$, denote the agent's continuation value after any non-terminal history when he acts optimally, i.e.,

$$U_t(\Gamma) \equiv \sup_a \mathbb{E}^a \left[R_\tau \cdot \mathbb{1}_{\{\tau \leq T\}} + \int_t^{T \wedge \tau} \phi(1 - a_s) ds | t \leq \tau \right]. \quad (3)$$

Assuming that $U_t \in \mathcal{C}^1$ for $t \in [0, T]$ (which will be verified later), the Hamilton-Jacobi-Bellman (HJB) equation for the agent's problem is given by¹¹

$$0 = U'_t + \sup_{a_t} \{ \phi(1 - a_t) + \lambda a_t (R_t - U_t) \}. \quad (4)$$

If the agent works over an interval dt , then he makes a breakthrough with probability λdt and gets the reward R_t . However, by doing so, he forgoes the private benefit of shirking ϕdt as well as the continuation utility, U_t , he would get if the breakthrough did not arrive. Thus, in order to give the agent incentives to work the principal must reward the agent with additional utility of at least ϕ/λ upon arrival of success. The incentive compatibility condition can therefore be summarized by the following lemma.

Lemma 3.1. *Given any contract Γ , the optimal action for the agent at time t is*

$$a_t = 1 \iff R_t \geq U_t(\Gamma) + \frac{\phi}{\lambda}. \quad (5)$$

Following [Spear and Srivastava \(1987\)](#), the principal's problem can be formulated recursively, where the state variable is the promised utility to the agent, denoted by $u \in \mathbb{R}_+$. Let $V(u)$ denote the principal's value function, which maximizes her expected utility, $P_0(\Gamma)$, subject to incentive compatibility (5) and the additional "promise keeping" condition

$$U_0(\Gamma) = u. \quad (6)$$

Notice that $U_t(\Gamma)$ is strictly positive for any $t < T$, therefore if $u = 0$, then the only possible solution is for the project to be terminated immediately and hence $V(0) = 0$. For

¹¹Expression (4) can be derived as the limit of the discrete-time Bellman equation

$$U_t = \sup_{a_t} \{ \phi(1 - a_t)dt + a_t (\lambda dt R_t + (1 - \lambda dt)U_{t+dt}) \}.$$

by using the first-order Taylor approximation $U_{t+dt} \approx U_t + U'_t dt$. Other HJB equations can be derived analogously

$u > 0$, the HJB equation for the principal's problem is

$$0 = \max_{R,a} \left\{ \lambda a (\Pi - R - V(u)) - c + V'(u) \frac{du}{dt} \right\}$$

$$\text{s.t. } \frac{du}{dt} = - \max_a \{ \phi(1-a) + \lambda a(R-u) \}.$$

Clearly, the solution to the principal's problem must involve $V'(u) \geq -1$ since the principal has the option to make direct payments to the agent. Hence, it is without loss to focus on contracts such that $a = 1$.¹² The first-order condition then requires that any solution to the principal's HJB involves $R = u + \frac{\phi}{\lambda}$ (i.e., the incentive compatibility condition binds), in which case (4) indicates that the agent's continuation value decreases at a constant rate ϕ prior to termination or success and the principal's HJB reduces to the ordinary differential equation

$$\lambda V(u) = \lambda \left(\Pi - u - \frac{\phi}{\lambda} \right) - c - \phi V'(u).$$

Using the boundary condition at $u = 0$ to pin down the constant, we arrive at the following candidate for the principal's value function

$$\bar{V}(u) \equiv \underbrace{\left(1 - e^{-\frac{\lambda u}{\phi}} \right)}_{\mathbb{P}(\tau \leq \frac{u}{\phi})} \left(\Pi - \frac{c}{\lambda} \right) - u. \quad (7)$$

The expression in (7) highlights the direct link between the promised utility (or rents) of the agent and the probability that the project ultimately succeeds.¹³

Proposition 3.2. *For a one-stage project, the principal's value function is given by $\bar{V}(u)$. Given any level of the agent's continuation utility $u \in \mathbb{R}_+$, this payoff can be attained under the contract with a deadline*

$$T^*(u) = \frac{u}{\phi}, \quad (8)$$

and a reward payment for success that is decreasing over time according to

$$R^*(\tau) = \phi \left(\frac{1}{\lambda} + T^*(u) - \tau \right), \quad \forall \tau \in [0, T^*(u)]. \quad (9)$$

The intuition behind this result is straightforward. The agent is given a deadline corresponding to the date at which his continuation utility will hit zero absent a breakthrough.

¹²If $a = 0$, the HJB is satisfied if and only if $V'(u) = -c/\phi \leq -1$ in which case a direct payment is a more efficient form of compensation than letting the agent shirk.

¹³Notice that $\bar{V}(u)$ is strictly concave, which confirms that randomization over the termination deadline is suboptimal and $\bar{V}'(u) \geq -1$, which confirms incentive compatibility binds.

If he innovates prior to that date, then he receives a payment compensating him for the instantaneous incentive to shirk ϕ/λ plus the opportunity cost of shirking for the remaining time on the clock, $\phi(T^*(u) - \tau)$. Thus, the sooner the agent makes a breakthrough, the larger is his share of the concomitant benefit.

In light of Proposition 3.2, the only thing left to pin down is the initial utility level for the agent, which is equivalent to the optimal termination date of the contract. Naturally, the division of surplus will depend on the relative bargaining power and outside options of each player. To fix ideas, throughout the paper we endow the principal with all of the bargaining power and set the agent's outside option to zero. We refer to the *optimal contract*, as the contract that maximizes the principal's payoff over all $u \in \mathbb{R}_+$. Therefore, let

$$u^* \equiv \arg \max_{u \in \mathbb{R}_+} \bar{V}(u).$$

We call a project *feasible* if the optimal contract does not involve immediate termination.

Corollary 3.3. *A single-stage project is feasible if and only if*

$$\lambda\Pi - c > \phi. \tag{C.1}$$

If (C.1) holds, then the optimal contract has a deadline $T^* \equiv \frac{1}{\lambda} \ln \left(\frac{\lambda\Pi - c}{\phi} \right)$.

The optimal deadline balances the probability that the project ultimately succeeds, $(1 - e^{-\lambda T^*})$, against the rents extracted by the agent, $u^* = \phi T^*$. The feasibility condition ensures that $T^* > 0$. Intuitively, the principal must give the agent ϕ per unit time in continuation value in order to prevent shirking. The condition therefore ensures that the principal's expected flow benefit, $\lambda\Pi$, outweighs her total flow cost of operating the project and inducing effort $(c + \phi)$. Notice that this inequality is stronger than would be needed in a first-best situation where the principal ran the project herself, namely $\lambda\Pi - c > 0$, in which case, as noted in Remark 2, the first-best policy is to invest indefinitely until the innovation arrives. Also notice that the optimal deadline increases as the agency problem becomes less severe, with $\lim_{\phi \rightarrow 0} T^* = \infty$. That is, the (second-best) outcome and principal payoffs converge to first-best as the agency conflict goes to zero. Hence, the deadline exists only to mitigate agency costs.

4 Two-stage projects

Having derived the optimal contract for a single-stage project in the previous section, we now introduce a second stage. Henceforth, the agent must make two breakthroughs in order for the principal to realize the project benefits. To simplify exposition, we assume that the

parameters (ϕ, c, λ) are the same for both stages.¹⁴ Analogous to Assumption 1, We assume that the expected value of the two-stage project is strictly positive; i.e., $\Pi - 2c/\lambda > 0$. Note that this assumption does *not* imply that the (two-stage) project is feasible. Indeed part of our interest is in characterizing the conditions under which the principal can profitably undertake a multistage project.

We let τ_1 and τ_2 denote the (random) times at which the first and second breakthrough occur. We distinguish between the *first stage* of the project, $t \in [0, \tau_1)$, and the *second stage*, $t \in [\tau_1, \tau_2)$. As before, the date at which the project ultimately succeeds, now denoted τ_2 , is publicly observed and can be directly contracted upon.¹⁵ For the majority of our analysis, we assume that intermediate progress (i.e., τ_1) is privately observed by the agent and cannot be verified by the principal.¹⁶ For this reason, we will refer to intermediate progress as being *intangible*.

Intangible progress can only be contracted upon indirectly via unverifiable “progress reports” from the agent. *A priori*, it is not clear that doing so has value to the principal. Recall that any information communicated is irrelevant for the socially optimal investment policy, hence reports are only beneficial to the principal if they reduce agency rents on net. Of course, in order to induce truthful reporting the agent must be given appropriate incentives. So, the question becomes whether the savings in rents associated with inducing more efficient investment (i.e., preventing shirking) justify the payment of the rents necessary to induce truthful reports.

4.1 Simple contracts

In a single-stage project, the principal can implement the optimal contract with a single (deterministic) deadline, T , and reward scheme that depends only on the project completion date. We refer to such contracts as *simple contracts*. Notice that simple contracts effectively preclude contracting on communication with the agent in any meaningful way. This leads to the undesirable feature that the agent will begin shirking as the deadline approaches if he has not made progress.

Proposition 4.1. *For any two-stage project and any simple contract with a bounded reward scheme, there exists a $\Delta > 0$ such that if the agent has not made the first breakthrough by $T - \Delta$, he will shirk for all $t > T - \Delta$.*

¹⁴We analyze projects with asymmetric stages in Section 6.

¹⁵This assumption is stronger than necessary. Our results continue to hold if τ_2 is privately observed by the agent but can be verified by the principal. This can be seen from the fact that the agent has no incentive to “hide” project success under the optimal contract when success is observable.

¹⁶In Section 6, we consider the case in which intermediate progress is unobserved by both the agent and principal.

Proof. The probability of completing the project by T given that the first breakthrough has not been made by $T - \Delta$ is

$$\Pr(\tau_2 \leq T | \tau_1 > T - \Delta) = 1 - e^{-\lambda\Delta} (1 + \lambda\Delta),$$

and as $\Delta \rightarrow 0$, the above expression, and hence the benefit to the agent of working, converges to zero at a rate proportional to Δ^2 , whereas the benefit of shirking during the time remaining is proportional to Δ , implying that for t close enough to T , the agent will prefer to shirk unless the reward for success is arbitrarily large. $\square$

In essence, an agent who has not yet made progress “runs out of steam” as the deadline approaches. A corollary of Proposition 4.1 is that simple contracts are not the most efficient way to provide incentives in the multistage setting. If τ_1 was publicly observable and contractible, the principal could improve upon a simple contract by incorporating a first-stage deadline whereby the project is terminated if the agent has not made a breakthrough by $T - \Delta$.¹⁷ However, because progress is intangible, the principal cannot employ a first-stage deadline and expect the agent to report truthfully. Doing so would induce the agent to make false claims of progress (just prior to $T - \Delta$) and then shirk.

Nevertheless, a simple contract can be improved upon by enriching the contract space to allow for communication. To see this, suppose that upon reaching $T - \Delta$, the principal asks the agent “have you made a breakthrough yet?” If the agent reports “no”, then the project is terminated but the principal gives him a severance payment of $\phi\Delta$, which is exactly what the agent could obtain by diverting cash flows until T . If the agent answers “yes” then the principal continues funding the project until T with the same reward scheme. This arrangement induces truth-telling for an agent who has not made a breakthrough (since he is just indifferent). It also induces truthful reporting for an agent who has made a breakthrough provided that following a breakthrough working is incentive compatible under the original contract. Moreover, this modification does not change the probability of project success (since the agent without a breakthrough would have shirked anyway) and saves the principal $(c - \phi)\Delta \geq 0$ conditional on a “no” report. Thus, communication with the agent can weakly improve the principal’s expected payoff.

In fact, as we will see, the optimal contract has a similar structure to the one just described save for one major difference. Instead of deterministic termination and a severance payment following a negative report at $T - \Delta$, the principal begins a probationary phase in which she effectively “freezes” time (and the agent’s promised utility), but starts randomly terminating

¹⁷In Green and Taylor (2015), we show that optimal contract resembles this multi-deadline arrangement if progress is observable and contractible.

the project at a constant flow rate. By doing this, the principal still induces truthful reporting, but also induces an agent who has not made a breakthrough to continue working, thereby increasing the probability that the project ultimately succeeds relative to a contract with certain termination and a severance payment.

5 The optimal contract

In this section, we enrich the contract space to allow for communication and characterize the optimal contract and the set of equilibrium payoffs.

A contract is again denoted by $\Gamma = \{a, Y, T\}$, where a is the recommended action, T is the termination rule, and dY_t denotes the payment to the agent at time t .¹⁸ Each of the elements of the contract can depend on the history, which includes any reports made by the agent, whether the project has succeeded, and a public randomization device. Clearly, the principal cannot condition payments or the termination rule directly on τ_1 , rather, she can only condition on information communicated by the agent and τ_2 . Given an arbitrary contract, the agent's continuation value can depend on any information communicated and the actual stage of the project.¹⁹ However, by the revelation principle, in searching for the optimal mechanism it is without loss to focus on direct mechanisms that induce truthful reporting (see Myerson (1986) or Pavan et al. (2014) for further discussion). We can therefore restrict attention to direct mechanisms that induce the agent to report truthfully and immediately. We use $\hat{\tau}_1$ to denote the time at which the agent reports the first breakthrough in order to distinguish it from the time at which the first breakthrough is made.

As in other settings with persistent private information (e.g., Fernandes and Phelan (2000), Zhang (2009) and Guo and Hörner (2015)), it is convenient to formulate the principal's problem as a dynamic program using the vector of promised utilities as the state variable. Thus, we will make use of three state variables: the project stage $s \in \{1, 2\}$ (as reported by the agent) and the pair of promised continuation values to each type, $(u_1, u_2) \in \mathbb{R}_+^2$. Henceforth, we refer to the agent who has made a breakthrough as the “high” type and agent who has not yet made a breakthrough as the “low” type. For states in which $s = 2$, u_1 can be interpreted as the maximal payoff that a low type could obtain by falsely reporting a breakthrough. For states in which $s = 1$, u_2 can be interpreted either as the maximal payoff that a high type could obtain by not reporting progress or as the promised reward to the low type for making a breakthrough in that state. Importantly, these two payoffs must be

¹⁸Unlike the one-stage project, we cannot immediately rule out payments being made at dates other than that of project success. Though such payments are not part of the optimal contract, allowing for them facilitates the proof of several intermediate results (e.g., Lemma 5.3).

¹⁹Note that the actual time at which the breakthrough was made is irrelevant for the agent's *continuation value*. Only whether a breakthrough was made and if and when it was reported matters.

Figure 1: The set of implementable utilities corresponds to the shaded area including the 45-degree line but excluding the vertical axis.

identical given the agent is behaving optimally.

5.1 Implementable utility levels

It will be useful to characterize the set of utility pairs that are *implementable*. A utility pair $\vec{u} = (u_1, u_2) \in \mathbb{R}_+^2$ is implementable if there exists a mapping from the agent's type, s , to a contract Γ_s such that (i) each agent prefers to report his type truthfully, (ii) the recommended action is incentive compatible, and (iii) the contract delivers expected payoff of u_s to an agent who truthfully reports $s \in \{1, 2\}$. We denote the set of implementable utility levels by $\mathcal{U} \subset \mathbb{R}_+^2$. Of course, not all utility pairs are implementable. For example, the high type can always “mimic” the low type, and thus $u_1 > u_2$ is not implementable. The following lemma says that this is essentially the only restriction on the set of implementable utility pairs.

Lemma 5.1. *For a two-stage project, the set of implementable utility pairs is given by $\mathcal{U} = \{(u_1, u_2) \in \mathbb{R}_{++}^2 : u_2 \geq u_1\} \cup (0, 0)$.*

The proof is by construction and can be found in the online appendix. Let $\mathcal{L}_H \equiv \mathcal{U} \cap \{u_2 = u_1 + \phi/\lambda\}$ denote the line along which the incentive compatibility condition for the low-type agent holds with equality and define the two complimentary subregions $\mathcal{U}_L \equiv \mathcal{U} \cap \{u_2 < u_1 + \phi/\lambda\}$ and $\mathcal{U}_H \equiv \mathcal{U} \cap \{u_2 \geq u_1 + \phi/\lambda\}$. Notice that the low-type agent will strictly prefer to shirk for all $\vec{u} \in \mathcal{U}_L$ and will (weakly) prefer not to shirk for all $\vec{u} \in \mathcal{U}_H$. Finally, let $\mathcal{L}_L \equiv \mathcal{U} \cap \{u_1 = u_2\}$ denote the lower boundary of $\mathcal{U}$. These regions are illustrated in Figure 1.

5.2 Second stage problem

We can solve for the optimal contract using backward induction on the project stage. That is, given an arbitrary implementable utility pair, $(u_1, u_2) \in \mathcal{U}$, we first solve for the principal's

value function and optimal policy in the second stage and then use these payoffs to find the optimal first stage policy. We let $F_2 : \mathcal{U} \rightarrow \mathbb{R}$ denote the principal's second-stage value function, which maximizes her payoff subject to incentive compatibility (i.e., (11)), delivering the required promised utility u_2 to a high type (i.e., (12)) and delivering no more than u_1 to the low type (i.e., (13)).

Formally, F_2 solves

$$F_2(u_1, u_2) = \sup_{\Gamma} \mathbb{E}^a \left[\Pi \cdot \mathbf{1}_{\{\tau_2 \leq T\}} - \int_0^{T \wedge \tau_2} \{cdt + dY_t\} \mid s = 2 \right] \quad (10)$$

$$\text{s.t.} \quad a \in \arg \max_{\tilde{a}} \mathbb{E}^{\tilde{a}} \left[\int_0^{T \wedge \tau_2} \{\phi(1 - \tilde{a}_t)dt + dY_t\} \mid s = 2 \right] \quad (11)$$

$$u_2 = \mathbb{E}^a \left[\int_0^{T \wedge \tau_2} \{\phi(1 - a_t)dt + dY_t\} \mid s = 2 \right] \quad (12)$$

$$u_1 \geq \max_{\tilde{a}} \mathbb{E}^{\tilde{a}} \left[\int_0^{T \wedge \tau_2} \{\phi(1 - \tilde{a}_t)dt + dY_t\} \mid s = 1 \right]. \quad (13)$$

Proposition 5.2. *For any $\vec{u} \in \mathcal{U}$, the principal's value function in the second stage is given by*

$$F_2(u_1, u_2) = (1 - e^{-\lambda u_1 / \phi}) \left(\Pi - \frac{c}{\lambda} \right) - u_2. \quad (14)$$

Moreover, conditional on reaching the second stage with promised utilities (u_1, u_2) , the optimal continuation contract can be implemented with a simple contract with deadline u_1 / ϕ .

The expression for the principal's value function is intuitive in light of the single-stage benchmark (see equation (7)). The larger is u_1 , the longer is the deadline the principal can give the agent without violating (13), and therefore the higher is the probability of making the last breakthrough. However, in contrast to the single-stage benchmark, the intangible nature of progress leads to a decoupling of the link between the promised utility to the agent who has truthfully reported a breakthrough (i.e., u_2) and the maximum amount of time the principal can give the agent to complete the project (i.e., u_1 / ϕ).

5.3 First-stage problem

We can now turn to the principal's problem in the first stage. Prior to a reported breakthrough, the principal chooses a termination rule, reward scheme and recommended action, as well as how much utility to deliver to each type of agent upon reporting a breakthrough, denoted by W_1, W_2 . To induce truth telling, two additional constraints are required. We incorporate random termination (a necessary feature of the optimal contract) by letting the principal choose a distribution over termination dates denoted by S , which combined with a induces a

distribution over (τ_1, τ_2, T) .²⁰ Denote the corresponding expectation operator by $\mathbb{E}^{(a,S)}$. The principal's problem in the first stage can be written as follows.

$$\sup_{\Gamma} \mathbb{E}^{(a,S)} \left[F_2(W_1(\tau_1), W_2(\tau_1)) \mathbb{1}_{\{\tau_1 \leq T\}} - \int_0^{T \wedge \tau_1} \{c dt + dY_t\} \middle| s = 1 \right] \quad (\text{OBJ}_1)$$

subject to

$$a \in \arg \max_{\tilde{a}} \mathbb{E}^{(\tilde{a},S)} \left[W_2(\tau_1) \mathbb{1}_{\{\tau_1 \leq T\}} + \int_0^{T \wedge \tau_1} \{\phi(1 - a_t) dt + dY_t\} \middle| s = 1 \right] \quad (15)$$

and for all $t < T$, the truth-telling constraints are given by

$$W_1(t) \leq U_1(t) \equiv \mathbb{E}_t^{(a,S)} \left[W_2(\tau_1) \mathbb{1}_{\{\tau_1 \leq T\}} + \int_t^{T \wedge \tau_1} \{\phi(1 - a_t) dt + dY_t\} \middle| s = 1 \right] \quad (16)$$

$$W_2(t) \geq \max_{\tilde{a}, \tilde{\tau}_1 \geq t} \mathbb{E}_t^{(\tilde{a},S)} \left[W_2(\tilde{\tau}_1) \mathbb{1}_{\{\tilde{\tau}_1 \leq T\}} + \int_t^{T \wedge \tilde{\tau}_1} \{\phi(1 - \tilde{a}_t) dt + dY_t\} \middle| s = 2 \right]. \quad (17)$$

The first truth-telling constraint ensures that the low-type agent does not want to falsely report a breakthrough and the second ensures that the high-type agent cannot benefit from “hiding” a breakthrough from the principal. To solve the principal's first-stage problem, we first show that it is without loss to focus on contracts in which the agent does not shirk along the equilibrium path (Lemma 5.3). We then formulate a recursive version of the problem that relaxes (17) and hence only requires keeping track of the low-type agent's continuation value. The solution to the relaxed program is characterized in Proposition 5.4. Finally, we show that there exists a contract that satisfies the neglected constraint under which the principal obtains the same value as in the solution to the relaxed program (Proposition 5.5).

Lemma 5.3. *For any Γ satisfying (15)-(17) that involves the agent shirking over some interval of time, there exists a $\hat{\Gamma}$ that also satisfies (15)-(17) such that the agent does not shirk and $P_0(\hat{\Gamma}) \geq P_0(\Gamma)$.*

The intuition for this result is similar to how a simple contract can be improved upon as described in Section 4.1. That is, it is cheaper to compensate the agent directly with a payment than by letting him shirk.

To formulate the problem recursively, we incorporate the principal's ability to randomly terminate the project by letting σ denote the hazard rate of termination. The HJB for the

²⁰Naturally, we require S to be a right-continuous process and S_t to be measurable with respect to the principal's information set, including whether a breakthrough has been reported at (or prior to) time t .

relaxed problem is as follows,

$$0 = \max \left\{ \sup_{w_1, w_2, \sigma} \left\{ \lambda F_2(w_1, w_2) - (\lambda + \sigma) F_1(u_1) - c + F_1'(u_1) \frac{du_1}{dt} \right\}, \right. \\ \left. u_1 F_1'(u_1) - F_1(u_1) \right\}, \quad (\text{HJB})$$

subject to

$$u_1 \geq w_1 \quad (\text{NFP})$$

$$\lambda(w_2 - u_1) \geq \phi \quad (\text{IC})$$

$$\frac{du_1}{dt} = -\lambda(w_2 - u_1) + \sigma u_1 \quad (\text{PK})$$

$$F_1(0) = 0. \quad (\text{BC})$$

The no-false-progress constraint ([NFP](#)) requires that the “reward” to a low-type agent for falsely reporting a breakthrough (w_1) is no more than what she gets for being truthful (u_1). Because F_2 is increasing in w_1 ([Proposition 5.2](#)), this constraint binds; the principal would like to increase w_1 above u_1 following a reported breakthrough in order to give the agent more time to make the second breakthrough, but doing so would induce a false report. From [Proposition 5.2](#), we also know that F_2 is decreasing in w_2 . Increasing w_2 simply means giving more rents to a high-type agent who reports a breakthrough without increasing the probability of project success. Clearly, there is no reason for the principal to do this except to incentivize the low-type agent. Hence, ([IC](#)) also binds.

Having established these two intuitive features, we state the solution to the relaxed problem and then explain several of its additional characteristics.

Proposition 5.4. *The solution to ([HJB](#)) is given by*

$$F_1(u_1) = \begin{cases} \left(\Pi - \frac{2c}{\lambda} \right) \left[1 - \left(\frac{u_1 - u_s + \frac{2\phi}{\lambda}}{u_s + \frac{2\phi}{\lambda}} \right) e^{-\frac{\lambda}{\phi}(u_1 - u_s)} \right] - u_1 & \text{if } u_1 \geq u_s, \\ \left(\Pi - \frac{2c}{\lambda} \right) \left[\frac{u_1}{u_s + \frac{2\phi}{\lambda}} \right] - u_1 & \text{if } u_1 \in [0, u_s]. \end{cases} \quad (18)$$

where u_s is implicitly defined by

$$\left(\frac{\lambda u_s}{\phi} + 2 \right) e^{-\lambda u_s / \phi} \equiv \frac{\lambda \Pi - 2c}{\lambda \Pi - c}. \quad (20)$$

For all $u_1 \geq u_s$, the optimal policy involves $w_2 = u_1 + \phi/\lambda$, $w_1 = u_1$, $\sigma = \frac{\phi}{u_s} \mathbb{1}_{\{u_1 = u_s\}}$. For

$u_1 < u_s$, the optimal policy involves giving the low-type agent continuation value of u_s with probability u_1/u_s , and terminating the project with probability $1 - u_1/u_s$.

To interpret the principal's value function, recall from Lemma 5.3 that the agent never shirks under an optimal contract. The only source of inefficiency, therefore, is the potential for project termination. The social value of the contract, $F_1(u_1) + u_1$, thus equals the first-best value of the project, $\Pi - 2c/\lambda$, times the probability the project is successfully completed, which is the expression in square brackets of (18) and (19).

To see why the solution involves random termination, suppose instead that the optimal contract is deterministic and let $\hat{F}_1$ denote the candidate value function. Then, (HJB) becomes

$$\lambda \hat{F}_1(u_1) = \lambda F_2(u_1, u_1 + \frac{\phi}{\lambda}) - c - \phi \hat{F}_1'(u_1), \quad (21)$$

which has a solution of the form

$$\hat{F}_1(u_1) = \Pi - \frac{2c}{\lambda} - u_1 - u_1 e^{-\frac{\lambda u_1}{\phi}} \left(\frac{\lambda \Pi - c}{\phi} \right) + H_1 e^{-\frac{\lambda u_1}{\phi}}, \quad (22)$$

where H_1 is an arbitrary constant. Imposing the terminal boundary condition, $\hat{F}_1(0) = 0$, we get that $H_1 = \frac{2c}{\lambda} - \Pi$. Evaluating the second derivative at $u_1 = 0$ gives

$$\hat{F}_1''(0) = \frac{\lambda^2 \Pi}{\phi^2} > 0.$$

The convexity in $\hat{F}_1$ is illustrated by the dotted line in Figure 2. Clearly, the principal could improve her payoff (above $\hat{F}_1$) by randomizing when the agents continuation value is near zero, violating our supposition.

For additional intuition, three observations are useful. First, the no-false-progress constraint prohibits the principal from granting the agent more (expected) time following a reported breakthrough. Second, a breakthrough in the first-stage has no value to the principal without its counterpart in the second stage. Third, when u_1 is sufficiently small, there is a very low probability that the agent will make both breakthroughs in the maximal remaining time that the principal can give the agent (i.e., $\frac{u_1}{\phi}$). Therefore, rather than continue to “let the clock run” (i.e., let u_1 smoothly decrease), the principal prefers to use random termination, which allows her to “pause the clock” and preserve time to make the second breakthrough conditional on making the first one before termination occurs.

To see how u_s is determined, observe that the net effect of increasing the rate of termination

Figure 2: This figure illustrates the solution to (HJB), which coincides with the principal’s value function under the optimal contract.

impacts the second term in (HJB), i.e.,

$$\underbrace{u_1 F_1'(u_1)}_{\text{Benefit of preserving } u_1} - \underbrace{F_1(u_1)}_{\text{Opportunity cost of termination}}$$

The first part of this expression is the benefit to the principal of “pushing” u_1 (and hence w_1) upward, which preserves additional time conditional on reaching the second stage. The second part is just the opportunity cost of terminating the project. Since the principal’s problem is linear in σ , the optimal point at which she begins using a flow rate of termination (i.e., u_s) must be such that the net effect is exactly zero. That is,

$$u_s F_1'(u_s) - F_1(u_s) = 0, \quad (23)$$

which can be interpreted as a local optimality condition. Note that for $u_1 > u_s$, F_1 satisfies (21) and therefore has the form in (22). Thus, the boundary condition in (23) implies a constant H_1 for each candidate u_s . It turns out that maximizing H_1 (and therefore F_1) over all possible u_s is equivalent to requiring that F_1 be twice differentiable at u_s , i.e.,

$$F_1''(u_1) = 0. \quad (24)$$

which is often referred to as the super-contact condition (Dumas, 1991). The value function stated in (18) of Proposition 5.4 is the solution to the free-boundary problem implied by (22)-(24).

Figure 3: This figure illustrates the optimal dynamics of continuation values under the implementation stated in Proposition 5.5. Prior to the soft deadline (i.e., κ^*), continuation values drift down along $\mathcal{L}_H$. Upon reaching κ^* , the evolution of continuation values depends on whether a breakthrough has been reported. If it has not, they stop drifting and either jump to the origin (if terminated) or remain at κ^* . If a breakthrough is reported at the soft deadline, then the continuation values drift down along the dotted line toward the origin.

5.4 Optimal dynamics

The optimal dynamics are illustrated in Figure 3. At $t = 0$, the agent's continuation utility starts at $u_I \equiv \arg \max_u F_1(u)$, the asterisk in the upper right corner of the figure. Prior to reporting a breakthrough, the continuation utilities drift down along $\mathcal{L}_H$ toward κ^* . If a breakthrough is reported prior to reaching κ^* then the optimal dynamics are not uniquely pinned down. That is, continuation values can continue evolving downward along $\mathcal{L}_H$ or they can travel to the interior of $\mathcal{U}_L$. Eventually, however, the state must travel toward the origin and if $\vec{u}$ reaches the origin prior to the second stage breakthrough, the project is terminated.

If a breakthrough is not reported prior to reaching κ^* , then the principal initiates the probationary phase in which she randomizes over terminating the project (the state jumps to the origin) and maintaining promised utilities at κ^* . Hence, κ^* , serves as a partially absorbing state until the probationary phase ends with either project termination or a reported breakthrough, at which point, the state variable again evolves toward the origin. Conditional on reaching κ^* and not being terminated, the promised utility to the agent once he (truthfully) reports a breakthrough is independent of when it is reported.

5.5 Implementation

To implement the optimal contract, communication between the principal and agent is critical. However, this communication is not required until the state reaches κ^* . That is, early on in the life of the project the principal need not be in communication with the agent regarding the status of the project. Instead, at $t = 0$, the principal simply gives the agent a future date

(i.e., the soft deadline T_s), at which a progress report is required. If the ultimate success of the project is realized prior to T_s then there is no need for the agent to make any report at all.²¹ The principal simply compensates him based on τ_2 . On the other hand, once the state reaches κ^* (i.e., at all $t \geq T_s$), the agent must report a breakthrough *as soon as it arrives* in order to avoid suboptimal termination.

We refer to the implementation described above as the *Minimally Optimal Communication Contract* (MOCC). It is the mechanism that minimizes the expected number of reports that the agent will be forced to make over the life of the project subject to delivering the maximal payoff to the principal (notice that if $\tau_2 < T_s$, then the agent does not make a report). Because of this property, the MOCC is uniquely optimal when the agent must incur (or the principal can impose) a small cost in order to report progress (see Section 6.1). Formally, it can be implemented as follows.

Proposition 5.5. *A two-stage project is feasible if and only if*

$$\lambda\Pi - 2c > 2\phi + \lambda\phi T^*. \quad (\text{C.2})$$

If (C.2) holds, the optimal contract can be implemented by use of a soft deadline T_s , a long clock $T_s + u_s/\phi$, a termination rate $\sigma = \phi/u_s$, and a reward function

$$R_s(\hat{\tau}_1, \tau_2) = \begin{cases} \phi \left(\frac{2}{\lambda} + \frac{1}{\sigma} + T_s - \tau_2 \right), & \text{if } \tau_2 \leq T_s \\ \phi \left(1 + \frac{\sigma}{\lambda} \right) \left(\frac{1}{\lambda} + \frac{1}{\sigma} + \max\{T_s, \hat{\tau}_1\} - \tau_2 \right) & \text{if } 0 < \tau_2 - \max\{T_s, \hat{\tau}_1\} \leq \frac{1}{\sigma} \\ 0 & \text{otherwise.} \end{cases}$$

such that:

- If the project is not completed prior to the soft deadline T_s , the principal asks the agent for a progress report at $t = T_s$.
 - If the agent reports that he has made the first breakthrough ($\hat{\tau}_1 \leq T_s$), then he is given the remaining time on the long clock, u_s/ϕ , to complete the project.
 - If the agent reports that he has not yet made the first breakthrough, then the principal stops the long clock and initiates a probationary phase in which the project is terminated at constant rate σ .

²¹While we have employed the revelation principle and solved for the optimal mechanism assuming that the agent reports progress immediately, the optimal contract does not require conditioning on the time at which the first breakthrough occurs if it is less than T_s . Therefore, the principal can achieve the same ex-ante payoff with no communication prior to T_s . Under this implementation, the principal's expected payoff prior to the soft deadline and project completion will depend on her belief about the project stage and will therefore be a weighted combination of F_1 and F_2 .

- *If the agent reports a breakthrough during the probationary phase, then he is given the remaining time on the long clock, u_s/ϕ , to complete the project.*
- *The agent gets the reward $R_s(\hat{\tau}_1, \tau_2)$ only if the project succeeds prior to being terminated.*

Even though the principal cannot substantiate the agent’s progress reports, such reports are an essential aspect of the optimal mechanism because they are used to govern the continuation contract: a report of “no” resulting in probation and one of “yes” resulting in a relatively short time to complete the project.²² Indeed, permitting the agent to “remain silent” regarding his progress would make the principal worse off as was demonstrated in Proposition 4.1.

Several other aspects of Proposition 5.5 also warrant discussion. First, instead of using a hard deadline and a severance payment to induce truthful reporting by the low-type agent (as discussed in Section 4.1), the principal finds it optimal to screen types with a soft deadline and no severance payment. In this way, the principal is able to induce the low-type agent to continue working, which increases the probability of project success relative to the severance contract. Second, to understand why the truth-telling constraints are satisfied, note that at every instant of the probationary phase the low-type agent is indifferent between honestly reporting his lack of progress (and facing continued probation while he works) and falsely reporting progress (and optimally shirking over the time left on the long clock). The high-type agent, however, strictly prefers reporting his true status at this point – that is, he strictly prefers having a short period of time to complete the project to facing probation. Intuitively, the low-type weakly prefers a lottery over a long deadline and termination to a short deadline because he is unlikely to finish the project with a short deadline. The high-type agent strictly prefers the short deadline since he has already made one breakthrough, he expects it will take him less time to complete the project and does not want to risk termination.

Finally, even under the MOCC, the reward schedule is not uniquely pinned down for $\tau_2 > T_s$. Any reward function (including those which are non-monotone in τ_2) satisfying promise keeping and incentive compatibility for the high-type agent will suffice. The reward function given in Proposition 5.5 is the unique one meeting these criteria that induces piecewise linear promised utility for the high-type agent.

6 Project Design

In this section we investigate three variations of the model that have implications for certain aspects of project design. First, we consider costly reporting and ask whether there is scope

²²Lemma A.1 in the appendix shows that $u_s/\phi \leq T^*$.

for imposing a cost on the agent in order to submit a progress report. Next, we explore the case of asymmetric stages and ask whether it can be beneficial for the principal to design one of the stages to be more difficult than the other. Finally, we consider the case in which progress is unobservable to both players in order to see whether the principal can benefit from suppressing the agent's access to information about the status of the project. Proofs of the formal results in this section are relegated to the online appendix.

6.1 Costly Reporting

Suppose there are two channels through which the agent may report progress: an informal one (e.g. verbal) and a formal one (e.g., written). Both types of reports can be contracted upon and both are falsifiable. The only operational difference between the two channels is that the formal channel requires incurring a cost $\rho > 0$ (e.g., the time and effort of filling out documentation), whereas the informal channel remains costless for the agent to use (as it was in Section 5).²³ We will assume that using the formal channel is equally costly for the agent whether he submits a true or false report though our results hold *a fortiori* if the cost is higher for a false report.

The presence of the informal channel allows us to again invoke the revelation principle and focus only on direct mechanisms in which the agent truthfully reports progress as soon as it arrives. If the principal elects to use only the informal channel, then the setting is identical to the one studied in Section 5. The question is whether it is ever advantageous for her to require the agent to use the costly channel for reporting progress.

Proposition 6.1. *If the project is feasible using only the informal channel and if ρ is sufficiently small, then:*

(i) *There exists $u_s(\rho)$ defined implicitly by*

$$\left(\frac{\lambda u_s(\rho)}{\phi} + 2 \right) e^{-\lambda u_s(\rho)/\phi} \equiv \frac{\lambda \Pi - 2c}{(\lambda \Pi - c)e^{-\lambda \rho/\phi} + \lambda \rho}$$

such that for $u_1 \geq u_s(\rho)$ the principal's value function in the first stage is

$$F_1(u_1; \rho) = \left(\Pi - \frac{2c}{\lambda} \right) \left[1 - \left(\frac{u_1 - u_s(\rho) + \frac{2\phi}{\lambda}}{u_s(\rho) + \frac{2\phi}{\lambda}} \right) e^{-\frac{\lambda}{\phi}(u_1 - u_s(\rho))} \right] - u_1 \quad (25)$$

(ii) *There exists a $T_s(\rho) > 0$ such that it is optimal for the principal to require the agent to report progress through the formal channel if the project has not succeeded by $T_s(\rho)$.*

²³To respect limited liability, ρ should be interpreted as a direct loss of utility rather than a monetary cost.

Thus, the optimal contract requires the high type to document his informal claim of progress with a formal report if he does not complete the project by $T_s(\rho)$, which serves as the soft deadline. Additionally, any claims of progress during the probationary phase must be made through the formal channel.

To understand this result, first observe that the principal's value function converges to the case of intangible progress investigated in Section 5 as ρ goes to zero; i.e.,

$$\lim_{\rho \rightarrow 0} u_s(\rho) = u_s \implies \lim_{\rho \rightarrow 0} F_1(u_1; \rho) = F_1(u_1).$$

For $\rho > 0$, however, the principal's payoff is impacted in two ways, one positive and one negative. On the up side, requiring formal reports at the soft deadline relaxes the no-false-progress constraint and allows the principal to add extra time of ρ/ϕ to the clock. On the down side, with probability $e^{-\lambda(u_1 - u_s(\rho))/\phi}$ the high-type agent will have to make a costly formal report, and promise keeping necessitates that the principal compensate him for this event. In the proof of the proposition, we show that for relatively small values of ρ , the positive effect dominates the negative effect. In other words, the principal is better off from an ex-ante perspective if she requires the agent to use the costly channel if the project is not completed before the soft deadline is reached. On the other hand, if ρ is sufficiently large, then the second effect dominates and it is optimal for the principal to use only the informal channel.

Essentially, imposing a small cost of reporting progress on the agent after the soft deadline helps the principal to more effectively screen types. Note that this holds even though we have assumed that the reporting cost is the same whether or not the agent submits an honest report. The intuition is that reporting costs relax the no-false-progress constraint for the low type with probability one, while the cost of additional compensation to the high type is incurred with probability less than one (i.e., an agent who has made the first breakthrough will avoid paying the reporting cost if he completes the project prior to the soft deadline). Reporting costs, therefore, have greater impact on an agent who has not yet made progress, and it is this differential impact that allows the principal to benefit from making communication costly. To illustrate this point further, it is worth noting that the principal is strictly worse off under a direct mechanism that uses only the formal channel than under a direct mechanism using only the informal one.

Finally, recall that the optimal contract can be implemented via the MOCC (see Section 5), where no communication takes place prior to the soft deadline. With costly reporting, communication prior to the soft deadline is also not required. That is, the optimal contract can be implemented with a termination policy and wage scheme that do not depend on reports

made using the informal channel.²⁴ Therefore, provided that ρ is not too large, the informal channel is unnecessary; the MOCC using only the formal channel can be used to implement the optimum. Further, if the informal channel is unavailable (i.e., any communication requires the agent to incur a cost ρ), then the MOCC is uniquely optimal since it minimizes the probability of incurring the cost.

6.2 Asymmetric Stages

In many (probably most) relevant applications, each stage of the project is different. For example, one stage may be expected to take more time (have a smaller λ), require more working capital (higher c), and/or yield greater private benefits to the agent from shirking (higher ϕ). In this subsection, we extend our analysis to a setting with stages that are not identical.

In general, a stage $k \in \{1, 2\}$ can be described by the pair $(\phi_k/\lambda_k, c_k/\lambda_k)$. To fix ideas, we set $\phi_1 = \phi_2 = \phi$ and $c_1 = c_2 = c$ and parameterize the asymmetry of stages by $\alpha \in [-1, 1]$, where

$$\frac{1}{\lambda_1} = \frac{1 + \alpha}{\lambda} \quad \text{and} \quad \frac{1}{\lambda_2} = \frac{1 - \alpha}{\lambda}$$

for some fixed λ . This parametrization maintains a fixed (first-best) project value as we vary α (i.e., $\Pi - 2c/\lambda$) therefore allowing us to isolate the effect of the asymmetry. Also, note that the probability distribution of τ_2 is symmetric in α . For $\alpha = 0$, the two stages are identical. For $\alpha > 0$, the first stage is expected to take more time and require a larger fraction of the total resources than the second stage. We therefore refer to the first stage as being *harder* if $\alpha > 0$, and *easier* if $\alpha < 0$.

In the online appendix, we extend the formal analysis from Section 5 to this setting. With asymmetric stages, the principal's value function and the optimal policy has the same structure as when the stages are symmetric. In particular, in the first stage there exists some level of agent continuation value above which the constraints bind and below which the principal uses random termination.

Proposition 6.2. *Provided the project is feasible, the following statements hold.*

- (i) *There exists an $\underline{\alpha} \in (0, 1)$ such that the principal's value function is strictly increasing in α for all $\alpha \in (-\underline{\alpha}, \underline{\alpha})$.*
- (ii) *There exists an $\bar{\alpha} \in (\underline{\alpha}, 1)$ such that the principal's ex-ante payoff under the optimal contract is strictly decreasing in α for $\alpha \in (\bar{\alpha}, 1)$ and strictly increasing for $\alpha \in (-\bar{\alpha}, -1)$.*

²⁴See footnote 21 for further discussion.

(iii) As $\alpha \rightarrow \pm 1$, the principal's value function converges (uniformly) to that of a one-stage project where the arrival rate is $\lambda/2$.

The first result shows that the principal is better off (worse off) if the first stage is moderately more difficult (easier) than the second. This is due to the fact that under the optimal contract with symmetric stages, the agent has more time (in expectation) to complete the first stage than he does to complete the second stage. If the first stage is harder, the principal can use this structure to her advantage. However, (ii) says that making the first stage too difficult is not advantageous for the principal as it reduces her ability to allocate the remaining (expected) time until termination in the most efficient manner. Note that (i) and (ii) imply that there is some interior level of difficulty, $\alpha^* \neq 0$, which maximizes the principal's ex-ante payoff.²⁵ Intuitively, as $\alpha \rightarrow \pm 1$, the project effectively has only one stage and the principal's payoff converges accordingly. Part (iii) confirms this intuition.

These results have several novel implications for the optimal design of projects. For example, it is better to design projects in a way that places more difficult are completed first. Yet, one should be careful not to make the earlier stages too difficult as doing so inhibits the ability to efficiently manage agency rents and allocate resources across stages. In other words, it is always better to split a single-stage project into a project with two asymmetric stages.

6.3 Unobservable Progress

Because the agent privately observes the state of the project, the principal must give the agent information rents in order to induce truthful reporting. However, the principal also uses this information to more efficiently provide incentives. In considering how to design projects, it is then natural to ask whether these benefits outweigh the costs. That is, if the principal cannot observe progress, is it better to also restrict the agent's ability to do so?

Our results show that the answer to the question posed above is “no.” That is, the principal benefits more from effectively using this information than the rents she gives up to acquire it. One implication is that, in designing projects and their incentives schemes, the principal should focus on how best to illicit and effectively use the agent's private information (as was done in Section 5) rather than trying to suppress his access to it.

By way of terminology, we refer to progress as being unobservable if τ_1 is not observed by either party and hence cannot be contracted upon either directly or indirectly.

Proposition 6.3. *When progress is unobservable to both the principal and the agent, the*

²⁵Furthermore, our numerical examples suggest that the principal is always better off by putting the harder stage first, which suggests that $\alpha^* \in (0, 1)$.

principal's ex-ante payoff under the optimal contract is given by

$$\max_u \left(\Pi - \frac{2c}{\lambda} \right) (1 - e^{-\lambda u/\phi}) - \left(\Pi - \frac{c}{\lambda} \right) \frac{\lambda u}{\phi} e^{-\lambda u/\phi} - u,$$

which can be implemented using a simple contract.

When progress is privately observed by the agent, the principal can do strictly better than a simple contract by using a more sophisticated termination policy in which the project is terminated at a constant rate if the agent has not reported progress by the soft deadline. This allows the principal to give a low-type agent more time to complete the project conditional on having a breakthrough prior to being terminated. Why is it that the principal cannot benefit from a similar strategy with unobservable progress? The reason, of course, is that the principal has no way of knowing whether the project is still in the first stage or has progressed to the second. Any policy which uses random termination must do so indiscriminately, which means terminating projects that have already progressed to the next stage. This highlights the cost of unobservable progress relative to progress privately observed by the agent.

The benefit of unobservable progress is that the agent also does not know whether progress has been made, and it is therefore easier to induce effort.²⁶ For example, it is no longer the case that the agent necessarily begins shirking as the deadline approaches (i.e., Proposition 4.1 no longer holds), which is why a simple contract remains optimal. Having derived the principal's maximal payoff in both scenarios, we can easily compare them.

Corollary 6.4. *For any two-stage project that is feasible, under the respective optimal contracts, the principal is strictly better off when progress is privately observed by the agent than when it is unobservable to both the principal and the agent.*

Thus the costs of suppressing the agent's access to information are greater than the benefits; i.e., it is better for the principal to ask a privately-informed agent for progress reports than to contract with an agent who is unable to observe the evolution of the project.

7 Concluding Remarks

In this paper we study the optimal provision of incentives for two-stage projects in which the agent privately observes intermediate progress. We characterize the optimal contract and explore the implications for optimal project design.

The optimal contract exhibits features commonly observed in the execution of real-world multistage projects. Specifically, progress reports by the individuals responsible for moving the

²⁶This benefit is similar to the optimal contract in Fuchs (2007), in that by providing the agent with less information, there are fewer incentive compatibility constraints to satisfy.

venture forward are often required, although the consequences for delayed (reported) progress are often vague or imprecise. The optimal contract with the minimal amount of communication involves the use of a soft deadline, prior to which the agent is not required to communicate with the principal and after which breakthroughs must be reported immediately in order to avoid suboptimal termination. This feature closely resembles the project management concept of “slippage,” wherein later stages of the project are delayed until either progress is reported or the project is canceled.

Regarding the design of projects, we demonstrate three results. First, the principal achieves a higher payoff by imposing a small cost to the agent for reporting a breakthrough after some point. Second, the principal benefits from making the first stage somewhat more difficult than the second. Third, the principal does better if the agent privately observes progress and makes reports than if he does not observe progress himself.

There are numerous avenues for future work in this area. For instance, progress need not be completely unobservable and unverifiable but might be imperfectly observed by the principal, perhaps as the result of a costly audit. Also, we have focused on a setting with two discrete stages. Natural extensions would be to consider a setting with more stages (or model progress as a continuous process) and allow for the possibility of setbacks along the path to project completion. Additionally, in order to isolate progress as an instrument for providing incentives, we have suppressed uncertainty about the underlying value of the project. It would be edifying to study the role of progress in an environment where parties learn about the value of the project. Finally, our analysis is couched in a setting with full commitment by the principal. Relaxing this assumption may well shed important light on other aspects of the subject.

References

- Battaglini, M. (2005). Long-term contracting with Markovian consumers. *American Economic Review* 95, 637–658.
- Bergemann, D. and U. Hege (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance* 22(6-8), 703–735.
- Bergemann, D. and U. Hege (2005). The Financing of Innovation: Learning and Stopping. *Rand Journal of Economics* 36(4), 719–752.
- Bergemann, D. and J. Valimaki (2010). The Dynamic Pivot Mechanism. *Econometrica* 78(2), 771–789.
- Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large risk, limited liability and dynamic moral hazard. *Econometrica* 78, 73–118.
- Board, S. (2007). Selling options. *Journal of Economic Theory* 136, 324–340.
- Bonatti, A. and J. Hörner (2011). Collaborating. *American Economic Review* 101(2), 632–663.
- Clementi, G. L. and H. A. Hopenhayn (2006). A Theory of Financing Constraints and Firm Dynamics. *Quarterly Journal of Economics* 121, 229–265.
- DeMarzo, P. M. and M. J. Fishman (2007). Optimal long-term financial contracting. *Review of Financial Studies* 20, 2079–2128.
- DeMarzo, P. M., D. Livdan, and A. Tchisty (2014). Risking Other People’s Money: Gambling, Limited Liability, and Optimal Incentives. *Working Paper*.
- DeMarzo, P. M. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance* 61, 2681–2724.
- Dumas, B. (1991). Super contact and related optimality conditions. *Journal of Economic Dynamics and Control* 15, 675–685.
- Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012). Dynamic CEO Compensation. *Journal of Finance* 67(5), 1603–1647.
- Eso, P. and B. Szentes (2007). Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies* 74, 705–731.
- Ewusi-Mensah, K. and Z. H. Przasnyski (1991). On Information Systems Project Abandonment : An Exploratory Study of Organizational Practices. *MIS Quarterly* (March), 67–87.
- Fernandes, A. and C. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory* 91, 223–247.

- Fuchs, W. (2007). Contracting with Repeated Moral Hazard and Private Evaluations. *American Economic Review* 97, 1432–1448.
- Glass, R. L., J. Rost, and M. S. Matook (2008). Lying on software projects. *IEEE Software* 25(6), 90–95.
- Green, B. and C. Taylor (2015). On Breakthroughs, Deadlines, and the Nature of Progress: Optimal Contracts for Multistage Projects.
- Green, E. J. (1987). Lending and the Smoothing of Uninsurable Income. In E. C. P. Wallace and Neil (Eds.), *Minnesota Studies in Macroeconomics, Volume I*. University of Minnesota Press.
- Guo, Y. and J. Hörner (2015). Dynamic Mechanisms without Money. *Working Paper*.
- Halac, M., N. Kartik, and Q. Liu (2015). Optimal Contracts for Experimentation. *Working Paper*.
- Hoffmann, F. and S. Pfeil (2010). Reward for luck in a dynamic agency model. *Review of Financial Studies* 23, 3329–3345.
- Holmstrom, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics and Organization* 7, 24–52.
- Hopenhayn, H. A. and J. P. Nicolini (1997). Optimal Unemployment Insurance. *Journal of Political Economy* 105, 412.
- Hörner, J. and L. Samuelson (2014). Incentives for Experimenting Agents. *RAND Journal of Economics*.
- Hu, Z. (2014). Financing Innovation with Unobserved Progress. *Working Paper*.
- Inderst, R. and H. M. Mueller (2010). CEO replacement under private information. *Review of Financial Studies* 23, 2935–2969.
- Laux, C. (2001). Limited-liability and incentive contracting with multiple projects. *RAND Journal of Economics* 32, 514–526.
- Lerner, J. and U. Malmendier (2010). Contractibility and the Design of Research Agreements. *American Economic Review* 100, 214–246.
- Levitt, S. D. and C. M. Snyder (1997). Is no news bad news? Information transmission and the role of ‘early warning’ in the principal-agent model. *RAND Journal of Economics* 28, 641–661.
- Lewis, T. R. (2012). A theory of delegated search for the best alternative. *RAND Journal of Economics* 43, 391–416.
- Manso, G. (2011). Motivating Innovation. *Journal of Finance* 66, 1823–1860.

- Mason, R. and J. Välimäki (2015). Getting it done: dynamic incentives to complete a project. *Journal of the European Economic Association* 13(1), 62–97.
- Moroni, S. (2015). Experimentation in Organizations. *Working Paper*.
- Myerson, R. B. (1986). Multistage Games with Communication. *Econometrica* 54(2), 323–358.
- Pavan, A., I. Segal, and J. Toikka (2014). Dynamic Mechanism Design: A Myersonian Approach. *Econometrica* 82(2), 601–653.
- Phelan, C. and R. M. Townsend (1991). Computing Multi-Period, Information-Constrained Optima. *Review of Economic Studies* 58, 853.
- Piskorski, T. and A. Tchistyi (2011). Stochastic house appreciation and optimal mortgage lending. *Review of Financial Studies* 24, 1407–1446.
- Quadrini, V. (2004). Investment and liquidation in renegotiation-proof contracts with moral hazard. *Journal of Monetary Economics* 51, 713–751.
- Ray, D. (2002). The Time Structure of Self-Enforcing Agreements. *Econometrica* 70, 547–582.
- Royce, W. W. (1970). Managing the Development of Large Software Systems. *Ieee Wescon* (August), 1–9.
- Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies* 75, 957–984.
- Shavell, S. and L. Weiss (1979). The Optimal Payment of Unemployment Insurance Benefits over Time. *Journal of Political Economy* 87, 1347.
- Spear, S. E. and S. Srivastava (1987). On Repeated Moral Hazard with Discounting. *The Review of Economic Studies* 54, 599.
- Tchistyi, A. (2013). Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing. *Working Paper*.
- Tirole, J. (2006). *The Theory of Corporate Finance*, Volume 8. Princeton University Press.
- Toxvaerd, F. (2006). Time of the essence. *Journal of Economic Theory* 129(1), 252–272.
- Varas, F. (2015). Contracting Timely Delivery with Hard to Verify Quality. *Working Paper*.
- Williams, N. (2011). Persistent Private Information. *Econometrica* 79(4), 1233–1275.
- Zhang, Y. (2009). Dynamic contracting with persistent shocks. *Journal of Economic Theory* 144, 635–675.

A Appendix

Proof of Lemma 3.1. The HJB equation for the agent's problem can be derived in the usual way.

$$U_t = \sup_{a_t} (\lambda a_t R_t + \phi(1 - a_t))dt + (1 - \lambda a_t dt)U_{t+dt}$$

Using a Taylor expansion $U_{t+dt} = U_t + U'_t dt + o(dt)$, canceling U_t on both sides, dividing by dt and taking the limit as $dt \rightarrow 0$, we obtain (4). The lemma follows because (i) the HJB is a necessary condition for a to solve (3) and (ii) it is satisfied if and only if (5) holds. $\square$

Proof of Proposition 3.2. Fix an arbitrary $u \in \mathbb{R}_+$, we first show that $V(u) \leq \bar{V}(u)$. Since $\max_a U_0(\Gamma) \geq \phi T$, incentive compatibility and promise keeping requires that $T \leq u/\phi$. Because $\lambda \Pi > c \geq \phi$, given any finite deadline T satisfying $T \leq u/\phi$, total surplus is maximized by $T = u/\phi$ and $a_t = 1$ for all $t \in [0, u/\phi]$. That is,

$$V(u) + u \leq \max_{a, T \leq u/\phi} \mathbb{E}^a \left[\Pi \cdot \mathbf{1}_{\{\tau \leq T\}} - \int_0^{T \wedge \tau} c dt \right] = \left(1 - e^{-\lambda \frac{u}{\phi}}\right) \left(\Pi - \frac{c}{\lambda}\right).$$

Hence,

$$V(u) \leq \left(1 - e^{-\lambda \frac{u}{\phi}}\right) \left(\Pi - \frac{c}{\lambda}\right) - u = \bar{V}(u). \quad (\text{A.1})$$

Next, let $\Gamma(u)$ denote the contract with deadline $T_u = u/\phi$, reward scheme $R_u(t) = \phi(\frac{1}{\lambda} + u/\phi - t)$ and $a_t = 1$ for all $t \in [0, T(u)]$. Notice that $P_0(\Gamma(u)) = \bar{V}(u)$, $U_0(\Gamma(u)) = u$, and $\Gamma(u)$ is incentive compatible. Hence $V(u) \geq \bar{V}(u)$, which combined with (A.1) implies $V(u) = \bar{V}(u)$. $\square$

Proof of Proposition 5.2. We first show that the expression on the right side of (14) is an upper bound. For $\vec{u} \in \mathcal{U}_L \cup \mathcal{L}_H$, we then show there exists a simple contract that achieves this bound. For $\vec{u} \in \mathcal{U}_H \setminus \mathcal{L}_H$, we construct a sequence of contracts under which the principal's value converges to the right-hand side of (14). To derive the bound, note that

$$\begin{aligned} \max_{\tilde{a}} \mathbb{E}^{\tilde{a}} \left[\int_0^{T \wedge \tau_2} (\phi(1 - \tilde{a}_t) dt + dY_t) \mid s = 1 \right] &\geq \mathbb{E}^{\tilde{a}=0} \left[\int_0^{T \wedge \tau_2} (\phi(1 - \tilde{a}_t) dt + dY_t) \mid s = 1 \right] \\ &\geq \mathbb{E}^{\tilde{a}=0} \left[\int_0^{T \wedge \tau_2} \phi dt \mid s = 1 \right] = \phi \mathbb{E}[T]. \end{aligned}$$

Therefore, to satisfy (13), the termination policy must be such that $\mathbb{E}[T] \leq u_1/\phi$. The promise keeping constraint requires

$$u_2 - \mathbb{E}^a \left[\int_0^{T \wedge \tau_2} \phi(1 - a_t) dt \mid s = 1 \right] = \mathbb{E}^a \left[\int_0^{T \wedge \tau_2} dY_t \mid s = 1 \right],$$

and substituting into (10), we get that

$$F_2(u_1, u_2) \leq \sup_{a, T} \mathbb{E}^a \left[\Pi \cdot \mathbf{1}_{\{\tau_2 \leq T\}} - \int_0^{T \wedge \tau_2} (c - \phi(1 - a_t)) dt \mid s = 1 \right] - u_2, \quad \text{s.t. } \mathbb{E}[T] \leq u_1/\phi.$$

The right-hand side is increasing in a_t , hence

$$F_2(u_1, u_2) \leq \sup_T \left(\Pi - \frac{c}{\lambda} \right) \mathbb{E} \left[1 - e^{-\lambda T} \right] - u_2, \quad \text{s.t. } \mathbb{E}[T] \leq u_1/\phi.$$

By Jensen's inequality we have that $\mathbb{E}[1 - e^{-\lambda T}] \leq 1 - e^{-\lambda \mathbb{E}[T]}$, which is strictly increasing in $\mathbb{E}[T]$, and therefore the constraint binds. Inserting $\mathbb{E}[T] = u_1/\phi$ completes the proof of the bound.

For $\vec{u} \in \mathcal{U}_L \cup \mathcal{L}_H$, the bound can be achieved by a simple contract with a deadline of $T = u_1/\phi$ and $R(t) = \phi(\frac{1}{\lambda} + T - t) + q$ where $q = \frac{u_2 - u_1}{1 - e^{-\lambda u_1/\phi}}$. For $\vec{u} \in \mathcal{U}_H \setminus \mathcal{L}_H$, consider the sequence of simple contracts indexed by S and defined as follows:

$$\begin{aligned} R_S(t) &= \begin{cases} U_2(t) + S(U_2(t) - U_1(t)), & t \in [0, \Delta_S] \\ U_2(t) + \frac{\phi}{\lambda}, & t \in (\Delta_S, T_S] \end{cases} \\ \Delta_S &= \frac{1}{\lambda(S-1)} \ln \left(\frac{u_2 - u_1}{\frac{\phi}{\lambda}} \right) \\ T_S &= \frac{U_1(\Delta_S)}{\phi}, \end{aligned}$$

where $U_2(t)$ and $U_1(t)$ are given by (A.8) and (A.9) respectively. By construction, we have that

$$\begin{aligned} u_2 &= \int_0^{T_S} \lambda e^{-\lambda t} R_S(t) dt \\ u_1 &= \int_0^{\Delta_S} \lambda^2 t e^{-\lambda t} R_S(t) dt + e^{-\lambda \Delta_S} (1 + \lambda \Delta_S) \phi (T_S - \Delta_S) \\ R_S(t) &\geq U_2(t) + \frac{\phi}{\lambda}. \end{aligned}$$

Therefore, it is incentive compatible for the high-type agent to work for all $t \in [0, T_S]$ and his expected payoff from doing so is exactly u_2 . Furthermore, given that $U_2(t) = U_1(t) + \frac{\phi}{\lambda}$ for all $t \geq \Delta_S$, the maximal payoff to the low type is exactly u_1 . Finally, the expected payoff to the principal under this contract is given by $P_S = (1 - e^{-\lambda T_S}) (\Pi - \frac{c}{\lambda}) - u_2$ and $\lim_{S \rightarrow \infty} T_S = u_1/\phi$ implies that $\lim_{S \rightarrow \infty} P_S = F_2(u_1, u_2)$. $\square$

Proof of Lemma 5.3. First, note that Proposition 5.2 implies it is optimal to induce effort conditional on reaching the second stage. Therefore, it suffices to prove the same result holds in the first stage. To obtain a contradiction, let (t_1, t_2) denote an arbitrary interval of time over which the agent shirks in the first stage under Γ , where $0 \leq t_1 < t_2 \leq T$. To conserve notation, we assume that Γ involves a deterministic termination rule (the arguments in the proof can easily be extended). Define $\Delta = t_2 - t_1$. Let $\hat{\Gamma}$ be such that prior to a reported breakthrough:

- (i) $\hat{T} = T - \Delta$,
- (ii) For all $t < t_1$, $(\hat{a}_t, \hat{Y}_t, \hat{W}_1(t), \hat{W}_2(t))$ is identical to $(a_t, Y_t, W_1(t), W_2(t))$,
- (iii) For $t \in (t_1, \hat{T})$, let $(\hat{a}_t, \hat{Y}_t, \hat{W}_1(t), \hat{W}_2(t)) = (a_{t+\Delta}, Y_{t+\Delta}, W_1(t+\Delta), W_2(t+\Delta))$,
- (iv) At time t_1 , the principal makes an unconditional payment to the agent in the amount of $d\hat{Y}_{t_1} = \phi\Delta + \mathbb{E}^{a=0} \left[\int_{t_1}^{t_2} dY_t | s = 1 \right]$.

It is straightforward to check that if Γ satisfies (15)-(17) for all $t < T$ then $\hat{\Gamma}$ satisfies (15)-(17) for all $t < \hat{T}$. Prior to t_1 , both the agent's action and the principal's payoff conditional on a breakthrough is the same under both contracts. If a breakthrough does not happen prior to t_1 , then $\mathcal{P}_{t_1}(\hat{\Gamma}) = \mathcal{P}_{t_1}(\Gamma) + (c - \phi)\Delta \geq \mathcal{P}_{t_1}(\Gamma)$. Hence, $P_0(\hat{\Gamma}) \geq P_0(\Gamma)$. $\square$

Proof of Proposition 5.4. First, we construct the value function under the stated policy. Using the boundary conditions (23) and (24), we pin down u_s and show that the value function under the stated policy indeed has the form given by (18)-(19). We then verify that, given the u_s implied by the boundary conditions, the value function in (19)-(18) solves (HJB). That $\max_u F_1(u)$ is an upper bound on the solution to (OBJ₁) is immediate (since it relaxes (17)). In the Proof of Proposition (5.5), we show there exists a contract satisfying this additional constraint that achieves the bound.

For $u > u_s$, F_1 evolves according to (21) and therefore has a solution of the form (22). For $u < u_s$, the principal's value under the stated policy is given by

$$F_1(u_1) = \frac{u_1}{u_s} F_1(u_s) \quad (\text{A.2})$$

There are two unknowns to pin down: (u_s, H_1) . Solving (23) for H_1 gives

$$H_1 = \frac{\lambda u_s \left(-c\lambda \left(\Pi \phi \left(e^{\frac{\lambda u_s}{c\phi}} - 2 \right) + u_s \right) + \Pi \lambda^2 u_s + 2c^2 \phi \left(e^{\frac{\lambda u_s}{c\phi}} - 1 \right) \right)}{c\phi(c\phi + \lambda u_s)^2}.$$

Plugging this expression into (24), (or, equivalently, maximizing H_1 over all possible u_s), we get that u_s is defined implicitly by (20) and hence

$$H_1 = \left(\frac{\lambda u_s}{\phi} - 2 \right) \left(\Pi - \frac{c}{\lambda} \right). \quad (\text{A.3})$$

Substituting (A.3) into (22), we get (18). Then (19) follows from (A.2), which verifies that the value function under the stated policy for u_s given by (20) indeed has the stated form.

Finally, we verify that F_1 solves (HJB) subject to the four constraints. Given u_s as defined implicitly by (20), one can easily check that $F_1(0) = 0$ and thus (BC) is satisfied. Using arguments already given in the text, subject to (IC), (PK), and (NFP), we have that

$$\begin{aligned} \sup_{w_1, w_2, \sigma} \left\{ \lambda F_2(w_1, w_2) - (\lambda + \sigma) F_1(u_1) - c + F_1'(u_1) \frac{du_1}{dt} \right\} \\ = \underbrace{\lambda (F_2(u_1, u_1 + \phi/\lambda) - F_1(u_1)) - c - \phi F_1'(u_1)}_{L_s(u_1)} + \sup_{\sigma} \left\{ \sigma (u_1 F_1'(u_1) - F_1(u_1)) \right\}, \end{aligned}$$

and by construction, $L_s(u_1) = 0$ for $u_1 \geq u_s$ and $u_1 F_1'(u_1) - F_1(u_1) = 0$ for $u_1 \leq u_s$. Hence, it is sufficient to show that (i) $L_s(u_1) \leq 0$ for $u_1 < u_s$, and (ii) $u_1 F_1'(u_1) - F_1(u_1) \leq 0$ for $u_1 > u_s$. For (i), notice from (18) that L_s is concave and hence L_s' is decreasing for all $u_1 \in [0, u_s]$. That $L_s'(u_s) = 0$ implies L_s is increasing below u_s , and that $L_s(u_s) = 0$ then gives the result. For (ii), notice from (18) that $F_1''(u_1) < 0$ and hence $u_1 F_1'(u_1) - F_1(u_1)$ is decreasing for all $u_1 > u_s$. The result then follows since $u_s F_1'(u_s) = F_1(u_s)$. $\square$

The following lemma will be used in the proof of Proposition 5.5.

Lemma A.1. Define $u_I \equiv \arg \max_{u \geq 0} F_1(u)$.

1. If (C.2) holds strictly, then $u_s < \phi T^* < u_I$ and $F_1(u_I) > 0$.
2. If (C.2) holds with equality, then $u_s = \phi T^* = u_I$ and $F_1(u_I) = 0$.
3. If (C.2) is violated, then $F_1(u_I) < 0$.

Proof. A non-trivial solution to the contract-design problem exists under intangible progress if and only if $u_I > u_s$ satisfies the first-order condition

$$\left(\frac{\lambda\Pi - c}{\phi}\right) \left(\frac{\lambda(u_I - u_s)}{\phi} + 1\right) e^{-\lambda u_I/\phi} - 1 = 0, \quad (\text{A.4})$$

or

$$u_s = u_I + \frac{\phi}{\lambda} \left(1 - e^{\lambda(u_I/\phi - T^*)}\right).$$

The right hand side is maximized when $u_I = \phi T^*$, in which case $u_s = \phi T^*$. Therefore $u_I > u_s$ iff $\phi T^* > u_s$. To determine when $\phi T^* > u_s$, consider the function

$$J(u) \equiv \Pi - \frac{2c}{\lambda} - \left(\Pi - \frac{c}{\lambda}\right) \left(\frac{\lambda u}{\phi} + 2\right) e^{-\lambda u/\phi}.$$

Observe that $J(0) = -\Pi$, $\lim_{u \rightarrow \infty} J(u) = \Pi - 2c/\lambda$, and $J'(u) > 0$ for all $u < \infty$. Moreover, from (20), $J(u_s) = 0$. Therefore, $\phi T^* > u_s \iff J(\phi T^*) > 0$. The latter is equivalent to (C.2). $\square$

Proof of Proposition 5.5. Let Γ_s denote the contract stated in the proposition with

$$T_s \equiv \max \left\{ \frac{u_I - u_s}{\phi}, 0 \right\}.$$

We first show that Γ_s induces the prescribed behavior by the agent (i.e., truth telling and no shirking). Start from any $t \geq \tau_1$ (i.e., after a breakthrough has been made), and let $U_2(\tau_1, t)$ denote the high-type agent's equilibrium continuation value at time t from following the prescribed behavior in Γ_s .

- If $\tau_1 \geq T_s$, then

$$U_2(\tau_1, t) = \int_t^{\tau_1 + 1/\sigma} \lambda e^{-\lambda(\tau_2 - t)} R_s(\tau_1, \tau_2) d\tau_2 = \phi \left(1 + \frac{\sigma}{\lambda}\right) \left(\frac{1}{\sigma} + \tau_1 - t\right). \quad (\text{A.5})$$

To verify it is optimal for the agent to work (conditional on reporting truthfully) until making the second breakthrough or until running out of time, notice that

$$\lambda(R_s(\tau_1, t) - U_2(\tau_1, t)) = \phi \left(1 + \frac{\sigma}{\lambda}\right) > \phi.$$

Hence, working is strictly optimal. To verify that it is optimal for the agent to report progress immediately (i.e., that (17) holds), note that for any $t \geq T_s$,

$$W_2(t) = U_2(t, t) = \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma}\right).$$

Due to the stationarity of the continuation contract, if the agent prefers to delay reporting progress at $t = \tau_1$, then he prefers to delay reporting progress indefinitely (i.e., to never report progress). If he delays a report then it is strictly optimal to shirk (since the reward is zero if $\tau_2 \in (T_s, \hat{\tau}_1)$) and his expected payoff is

$$\int_t^\infty \sigma e^{-\sigma(s-t)} \phi(s-t) ds = \frac{\phi}{\sigma} < U_2(t, t).$$

Hence, the agent strictly prefers to report progress as soon as it arrives for all $\tau_1 \geq T_s$.

- If $\tau_1 < T_s$, then

$$U_2(\tau_1, t) = \int_t^{T_s} \lambda e^{-\lambda(\tau_2 - t)} R_s(\tau_1, \tau_2) d\tau_2 + \int_{T_s}^{T_s + 1/\sigma} \lambda e^{-\lambda(\tau_2 - t)} R_s(\tau_1, \tau_2) d\tau_2 = \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} + T_s - t \right).$$

Hence, $\lambda(R_s(\tau_1, t) - U_2(\tau_1, t)) = \phi$ and $a_t = 1$ is weakly optimal, regardless of whether progress is reported (since R_s is independent of $\hat{\tau}_1$ for $\tau_2 \leq T_s$). Further, the same argument as in the bullet above shows that the agent prefers to report progress at $t = T_s$ and work from that point forward.

Now consider any $t \leq \tau_1$ (i.e., prior to a breakthrough being made). Let $U_1(t)$ denote the low-type agent's equilibrium continuation value at time t under Γ_s .

- If $t \geq T_s$, then $W_2(t) = \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} \right)$ and thus

$$U_1(t) = \int_t^\infty \lambda e^{-(\lambda + \sigma)(\tau_1 - t)} W_2(\tau_1) d\tau_1 = \frac{\phi}{\sigma}. \quad (\text{A.6})$$

Since $\lambda(W_2(t) - U_1(t)) = \phi$, working is (weakly) optimal for the agent. Next, we verify that the agent does not want to falsely report progress at any $t \geq T_s$. Due to the stationarity of the continuation contract, if the agent prefers to falsely report progress at $t > T_s$, then he prefers to do so at $t = T_s$. Thus, suppose he falsely reports progress at the soft deadline. Let $\tilde{U}(t)$ be his expected payoff at $t \geq T_s$ from acting optimally henceforth. Consider the following chain

$$\begin{aligned} T_s < t &\iff \phi \left(\frac{\sigma}{\lambda} \right) \left(\frac{1}{\sigma} + T_s - t \right) < \frac{\phi}{\lambda} \\ &\iff \phi \left(1 + \frac{\sigma}{\lambda} \right) \left(\frac{1}{\sigma} + T_s - t \right) - \phi \left(\frac{1}{\sigma} + T_s - t \right) < \frac{\phi}{\lambda} \\ &\iff U_2(\tau_1, t) - \phi \left(\frac{1}{\sigma} + T_s - t \right) < \frac{\phi}{\lambda}, \end{aligned}$$

where the last line follows from (A.5). Next note that

$$\phi \left(\frac{1}{\sigma} + T_s - t \right) \leq \tilde{U}(t),$$

because the left side is the payoff to the agent from falsely reporting progress at T_s , having no breakthrough, and shirking from date t to $T_s + \frac{1}{\sigma}$ and the right side is his expected payoff from falsely reporting progress at T_s , having no breakthrough, and acting optimally from date t to $T_s + \frac{1}{\sigma}$. Thus we have

$$U_2(t) - \tilde{U}(t) < \frac{\phi}{\lambda},$$

which says that it is suboptimal for the agent to work at any $t > T_s$ after falsely reporting progress at T_s . Thus, if the agent falsely reports progress at T_s , then it is optimal for him to shirk until time runs out; i.e., $\tilde{U}(t) = \phi \left(\frac{1}{\sigma} + T_s - t \right)$. Finally, recall that ϕ/σ is the agent's expected payoff from honestly reporting no progress at T_s and then working while on probation. Thus, at T_s the agent is indifferent between honestly reporting no progress and then (optimally) working while on probation and falsely reporting progress and then

(optimally) shirking until time runs out. Therefore, honestly reporting no progress at T_s is weakly optimal.

- Next, consider $t < T_s$. Noting that

$$W_2(t) = \begin{cases} \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} + T_s - t \right), & \text{if } t < T_s \\ \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} \right), & \text{if } t \geq T_s. \end{cases}$$

and

$$U_1(t) = \int_t^{T_s} \lambda e^{-\lambda(\tau_1-t)} W_2(\tau_1) d\tau_1 + \int_{T_s}^{\infty} \lambda e^{-(\lambda+\sigma)(\tau_1-t)} W_2(\tau_1) d\tau_1 = \phi \left(\frac{1}{\sigma} + T_s - t \right), \quad (\text{A.7})$$

we get that $\lambda(W_2(t) - U_1(t)) = \phi$ for all $t < T_s$, which shows that working is weakly optimal. Finally note that working until T_s is also optimal for the agent if he plans to falsely report progress at T_s , that is $\tilde{U}(t) = U_1(t)$ for $t \leq T_s$.

Thus we have shown that Γ_s induces truth-telling and no shirking. Next we verify that the principal's expected payoff under Γ_s corresponds to $F_2(\tilde{U}(t), U_2(t))$ in the second stage and $F_1(U_1(t))$ in the first stage.

Second Stage: Suppose the agent has made the first breakthrough. By the above analysis, he immediately and truthfully reports τ_1 and continues to work until he makes the second breakthrough or runs out of time.

For any $t \in [\max\{T_s, \tau_1\}, \max\{T_s, \tau_1\} + 1/\sigma]$ the principal's expected payoff is thus

$$\begin{aligned} & \int_t^{\max\{T_s, \tau_1\} + 1/\sigma} \lambda e^{-\lambda(\tau_2-t)} \left(\Pi - \frac{c}{\lambda} - R_s(\tau_1, \tau_2) \right) d\tau_2 \\ &= \left(\Pi - \frac{c}{\lambda} \right) \left(1 - e^{-\lambda(1/\sigma + \max\{T_s, \tau_1\} - t)} \right) - \phi \left(1 + \frac{\sigma}{\lambda} \right) \left(\frac{1}{\sigma} + \max\{T_s, \tau_1\} - t \right) = F_2(\tilde{U}(t), U_2(t)). \end{aligned}$$

Similarly for $t \in [\tau_1, T_s]$, the principal's expected payoff is

$$\begin{aligned} & \int_t^{T_s} \lambda e^{-\lambda(\tau_2-t)} \left(\Pi - \frac{c}{\lambda} - R_s(\tau_1, \tau_2) \right) d\tau_2 + \int_{T_s}^{T_s + 1/\sigma} \lambda e^{-\lambda(\tau_2-t)} \left(\Pi - \frac{c}{\lambda} - R_s(\tau_1, \tau_2) \right) d\tau_2 \\ &= \left(\Pi - \frac{c}{\lambda} \right) \left(1 - e^{-\lambda(1/\sigma + T_s - t)} \right) - \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} + T_s - t \right) = F_2(\tilde{U}(t), U_2(t)). \end{aligned}$$

First Stage: Now suppose the agent has not yet made the first breakthrough. For $t \geq T_s$, Γ_s calls for random termination at rate σ , so the principal's expected payoff is

$$\begin{aligned} & \int_t^{\infty} \lambda e^{-(\lambda+\sigma)(\tau_1-t)} \left(\left(\Pi - \frac{c}{\lambda} \right) \left(1 - e^{-\lambda/\sigma} \right) - \frac{c}{\lambda} - \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} \right) \right) d\tau_1 \\ &= \left(\frac{\lambda}{\lambda + \sigma} \right) \left(\Pi - \frac{2c}{\lambda} - \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda/\sigma} \right) - \frac{\phi}{\sigma} \\ &= F_1(U_1(T_s)), \end{aligned}$$

where the last equality follows from $u_s = \frac{\phi}{\sigma}$ and substituting from (20). For $t < T_s$ the principal's

expected payoff is

$$\begin{aligned}
& \int_t^{T_s} \lambda e^{-\lambda(\tau_1-t)} \left(\left(\Pi - \frac{c}{\lambda} \right) \left(1 - e^{-(1/\sigma+T_s-\tau_1)} \right) - \phi \left(\frac{1}{\lambda} + \frac{1}{\sigma} + T_s - \tau_1 \right) - \frac{c}{\lambda} \right) d\tau_1 + F_1(U_1(T_s))e^{-\lambda(T_s-t)} \\
&= \left(\Pi - \frac{2c}{\lambda} \right) \left(1 - e^{-\lambda(T_s-t)} \right) + \lambda \left(\frac{1}{\sigma} + T_s - t \right) \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda(1/\sigma+T_s-t)} - \phi \left(\frac{1}{\sigma} + T_s - t \right) \\
&= F_1(U_1(t)),
\end{aligned}$$

where the last equality follows from $U_1(t) = \phi \left(\frac{1}{\sigma} + T_s - t \right)$, $u_s = \frac{\phi}{\sigma}$, and using (20) to substitute. Further, $U_1(0) = \phi(T_s + 1/\sigma) = u_I$, so that $F_1(U_1(0)) = F_1(u_I)$, which we have already argued is an upper bound on the solution to the principal's problem (see Proof of Proposition 5.4).

Thus, we can conclude that Γ_s is an optimal contract. Finally, that (C.2) is necessary and sufficient for a feasible contract to exist follows immediately from Lemma A.1. $\square$

B Supplemental Appendix (for online publication only)

B.1 Omitted Proofs

Proof of Lemma 5.1. Any implementable $\vec{u}$ such that $u_1 = 0$ must involve immediate termination with probability one, otherwise the low-type agent can shirk and obtain a strictly positive expected payoff. Hence, if $u_1 = 0$ then it must also be that $u_2 = 0$. Also, clearly the high type can achieve a weakly higher payoff than the low type given any Γ . Hence, any implementable pair must satisfy $u_2 \geq u_1$. Therefore, $\mathcal{U} \subseteq \{(u_1, u_2) \in \mathbb{R}_{++}^2 : u_2 \geq u_1\} \cup (0, 0)$. The rest of the proof is by construction.

Any $\vec{u} \in \mathcal{U}_L$ can be implemented by providing the agent a choice between two contracts. The first contract terminates the project immediately with a severance payment, $P = u_1$. The second contract has a deadline $T = u_1/\phi$ and a reward function $R(t) = \phi \left(\frac{1}{\lambda} + T - t \right) + q$ where $q = \frac{u_2 - u_1}{1 - e^{-\lambda u_1/\phi}}$. Clearly, the first contract delivers u_1 to both types. Under the second contract, the low type will strictly prefer to shirk (yielding a payoff of u_1) and the high type will strictly prefer to work (yielding an expected payoff of $u_2 \geq u_1$). Hence, each type weakly prefers to report truthfully and any $\vec{u} \in \mathcal{U}_L$ is implementable.

To prove that any $\vec{u} \in \mathcal{U}_H$ can be implemented, we use a contract that is independent of the agent's report for a period of length Δ at which point the continuation utilities are $(\hat{u}_1, \hat{u}_2) \in \mathcal{U}_L$ and the contract in the above paragraph is implemented for $t > \Delta$. To do so, define $S = \frac{u_2}{u_1} > 1$ (since $(u_1, u_2) \in \mathcal{U}_H$). Let $(\hat{u}_1, \hat{u}_2)$ be the point in $\mathcal{L}_H$ that intersects the ray with slope S that goes through the origin and (u_1, u_2) :

$$S\hat{u}_1 = \hat{u}_2 + \frac{\phi}{\lambda} \iff \hat{u}_1 = \frac{\phi}{\lambda} \left(\frac{u_1}{u_2 - u_1} \right), \text{ and } \hat{u}_2 = \frac{\phi}{\lambda} \left(\frac{u_2}{u_2 - u_1} \right).$$

Let $U_2(t), U_1(t)$ be the functions jointly satisfying

$$\begin{aligned} \frac{dU_2}{dt} &= \lambda S(U_1(t) - U_2(t)), & U_2(0) &= u_2 \\ \frac{dU_1}{dt} &= \lambda(U_1(t) - U_2(t)), & U_1(0) &= u_1, \end{aligned}$$

which have unique solutions given by

$$U_2(t) = \frac{1}{S-1} \left(S \left(u_1 + (u_2 - u_1)e^{-\lambda t(S-1)} \right) - u_2 \right) \quad (\text{A.8})$$

$$U_1(t) = \frac{1}{S-1} \left(S u_1 - u_2 + (u_2 - u_1)e^{-\lambda t(S-1)} \right). \quad (\text{A.9})$$

Note that $U_2(t) - U_1(t) = (u_2 - u_1)e^{-\lambda t(S-1)}$, so define Δ such that $U_2(\Delta) - U_1(\Delta) = \frac{\phi}{\lambda}$. That is

$$\Delta = \frac{1}{\lambda(S-1)} \ln \left(\frac{u_2 - u_1}{\phi/\lambda} \right).$$

Finally, let $\hat{R}(t) = U_2(t) + S(U_2(t) - U_1(t))$ for all $t < \Delta$. Note that for $t < \Delta$, neither type shirks since $\hat{R}(t) \geq U_2(t) + \frac{\phi}{\lambda}$ and $U_2(t) \geq U_1(t) + \frac{\phi}{\lambda}$. By construction, the contract that (i) pays $\hat{R}(t)$ for ultimate success prior to Δ and (ii) implements the contract from the first paragraph at $t = \Delta$ (at which point $(U_1(\Delta), U_2(\Delta)) = (\hat{u}_1, \hat{u}_2) \in \mathcal{U}_L$) delivers the desired expected utilities and induces truth telling, which completes the proof. $\square$

Proof of Proposition 6.1. The proof takes several steps. In Steps 1 through 3 we assume that the

principal *must* use the formal channel, but that she does so optimally. In Step 4 we prove that for ρ sufficiently small it is, in fact, optimal for the principal to use the formal channel.

Step 1. First observe that the only possible reason for requiring a costly report is to relax the no-false-progress constraint so as to add more time to the clock following the first reported breakthrough. This implies that the formal communication channel should only be used to report progress (not lack of progress). Next, observe that it cannot be optimal for the principal to require the costly formal report at date t and take no action until date $t' > t$, because this is dominated by waiting until t' to require the report. (The project might be completed between t and t' .) Thus, it is optimal to put off requiring a formal report as long as possible. Let y be the highest value of the low type's continuation utility at which the principal requires a formal report, and let $F_2(u_1, u_2; \rho)$ denote the principal's value function in the second stage. This value function is given by

$$F_2(u_1, u_2; \rho) = \left(\Pi - \frac{c}{\lambda} \right) \left(1 - e^{-\lambda(u_1 + \rho)/\phi} \right) - u_2 - \mathbb{1}_{\{u_1 \leq y\}} \rho e^{-\lambda(u_1 - y)/\phi}. \quad (\text{A.10})$$

This is established using the same method of proof as for Proposition 5.2 with two straightforward alterations. First, to satisfy (13), the termination policy must be such that $\mathbb{E}[T] \leq (u_1 + \rho)/\phi$. Second, for $u_1 \leq y$, there is a chance that the high type will have to pay the reporting cost. Therefore, the promise keeping constraint necessitates

$$u_2 + \rho e^{-\lambda(u_1 - y)/\phi} = \mathbb{E}^{a=1} \left[\int_0^{T \wedge \tau_2} dY_t | s = 1 \right],$$

Because $F_2(u_1, u_2; \rho)$ is concave in u_1 and linear in u_2 , $\sigma = 0$ is optimal.

Step 2. We solve for the first-stage value function assuming $\sigma = 0$ and check for concavity. The HJB is

$$\lambda F_1(u_1; \rho) = \max_{y \geq 0} \lambda F_2(u_1, u_1 + \phi/\lambda; \rho) - c - \phi \frac{dF_1}{du_1}.$$

Because the principal prefers to delay formal reports as long as possible, $y = 0$ is optimal. Therefore we have

$$\frac{\lambda}{\phi} F_1(u_1; \rho) + F_1'(u_1; \rho) = \frac{\lambda}{\phi} \left[\Pi - \frac{2c}{\lambda} - \frac{\phi}{\lambda} - u_1 - \left(\rho + \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda\rho/\phi} \right) e^{-\lambda u_1/\phi} \right].$$

This has a solution of the form

$$F_1^c(u_1; \rho) = \Pi - \frac{2c}{\lambda} - u_1 - \left(\rho + \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda\rho/\phi} \right) \frac{\lambda u_1}{\phi} e^{-\lambda u_1/\phi} + K e^{-\lambda u_1/\phi}. \quad (\text{A.11})$$

Using the boundary condition $F_1(0; \rho) = 0$ gives

$$F_1^c(u_1; \rho) = \left(\Pi - \frac{2c}{\lambda} \right) \left(1 - e^{-\lambda u_1/\phi} \right) - u_1 - \left(\rho + \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda\rho/\phi} \right) \frac{\lambda u_1}{\phi} e^{-\lambda u_1/\phi}.$$

For small ρ , this is convex for u_1 sufficiently close to zero, necessitating random termination at some point $u_s(\rho)$.

Step 3. Using the same analysis as in the proof of Proposition 5.4 yields (25), where $y = u_s(\rho)$ is implicitly defined by

$$\Pi - \frac{2c}{\lambda} - \left(\rho + \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda\rho/\phi} \right) \left(2 + \frac{\lambda u_s(\rho)}{\phi} \right) e^{-\lambda u_s(\rho)/\phi} \equiv 0.$$

Step 4. Define $u_I(\rho) \equiv \arg \max_{u \geq 0} F_1(u; \rho)$. We wish to show that for ρ sufficiently small, $F_1(u_I(\rho); \rho) > F_1(u_I)$. Because $\lim_{\rho \rightarrow 0} u_I(\rho) = u_I$, the claim will follow if we show $\frac{\partial F_1(u_I; 0)}{\partial \rho} > 0$. Applying the Envelope Theorem

$$\frac{\partial F_1(u_I; 0)}{\partial \rho} = u'_s(0) \frac{\lambda}{\phi} \left(\Pi - \frac{c}{\lambda} \right) e^{-\lambda u_I / \phi} + \frac{\lambda}{\phi} \left(\Pi - \frac{c}{\lambda} - \frac{\phi}{\lambda} \right) \left(2 + \frac{\lambda(u_I - u_s)}{\phi} \right) e^{-\lambda u_I / \phi}.$$

Moreover

$$u'_s(0) = - \frac{\left(\Pi - \frac{c}{\lambda} - \frac{\phi}{\lambda} \right) \left(2 + \frac{\lambda u_s}{\phi} \right)}{\left(\Pi - \frac{c}{\lambda} \right) \left(1 + \frac{\lambda u_s}{\phi} \right)}.$$

Therefore

$$\frac{\partial F_1(u_I; 0)}{\partial \rho} = \left[\frac{\lambda}{\phi} e^{-\lambda u_I / \phi} \left(\Pi - \frac{c}{\lambda} - \frac{\phi}{\lambda} \right) \right] \left[2 + \frac{\lambda(u_I / u_s)}{\phi} - \frac{2 + \frac{\lambda u_s}{\phi}}{1 + \frac{\lambda u_s}{\phi}} \right].$$

The first term in brackets is evidently strictly positive. Moreover, it follows from $u_I > u_s > 0$ and straightforward algebra that the second term in brackets is also strictly positive. $\square$

Proof of Proposition 6.2. First, we extend the analysis from Section 5.3 to characterize F_1^α for an arbitrary α . Using the same arguments as in Proposition 5.2, we have that

$$F_2^\alpha(u_1, u_2) = \left(1 - e^{-\frac{\lambda u_1}{(1-\alpha)\phi}} \right) \left(\Pi - \frac{c(1-\alpha)}{\lambda} \right) - u_2$$

Replacing F_2 with F_2^α in (HJB) and the appropriately modified (binding) constraints, we find that F_1^α has the form

$$\hat{F}_1^\alpha(u_1) = \left(\Pi - \frac{2c}{\lambda} - u_1 \right) + \left(\frac{1-\alpha}{2\alpha} \right) \left(\Pi - \frac{(1-\alpha)c}{\lambda} \right) e^{-\frac{\lambda u_1}{(1-\alpha)\phi}} + C_1^\alpha e^{-\frac{\lambda u_1}{(1+\alpha)\phi}}. \quad (\text{A.12})$$

If the terminal boundary condition is imposed ($\hat{F}_1^\alpha(0) = 0$) then $C_1^\alpha = \frac{-(1+\alpha)(\lambda\Pi - c(1+\alpha))}{2\alpha\lambda}$ and $\hat{F}_1^{\alpha''}(0) = \frac{\lambda^2\Pi}{(1-\alpha^2)\phi} > 0$. Hence, there exists some $u_s(\alpha)$ such that random termination is optimal for $u \in (0, u_s(\alpha)]$. Let $c^\alpha(u)$ denote the constant in the principal's value function that satisfies the smooth-pasting condition (i.e., (23)) at an arbitrary $u > 0$. That is,

$$c^\alpha(u) \equiv \frac{(\alpha+1)e^{\frac{2\alpha\lambda u}{(\alpha^2-1)\phi}} \left(c \left(\phi \left(\alpha^2 - 2\alpha + 4\alpha e^{\frac{\lambda u}{\phi-\alpha\phi}} + 1 \right) + u(\lambda - \alpha\lambda) \right) - \lambda\Pi \left(-\alpha\phi + 2\alpha\phi e^{\frac{\lambda u}{\phi-\alpha\phi}} + \lambda u + \phi \right) \right)}{2\alpha\lambda(\alpha\phi + \lambda u + \phi)}. \quad (\text{A.13})$$

Twice differentiability at $u_s(\alpha)$ (i.e., (24)) is equivalent to $u_s(\alpha) = \max_u c^\alpha(u)$, which requires the first-order condition

$$\frac{e^{\lambda u_s(\alpha)/\phi(1-\alpha)}}{\lambda u_s(\alpha) + 2\phi} = \frac{\lambda\Pi - c(1-\alpha)}{\phi(1-\alpha)(\lambda\Pi - 2c)}. \quad (\text{A.14})$$

The right-hand side of the above expression is strictly greater than $1/2\phi$ for all $\alpha \in (-1, 1)$. The left-hand side is equal to $1/2\phi$ at $u_s(\alpha) = 0$, strictly increasing in $u_s(\alpha)$ and unbounded. This guarantees the existence of a unique $u_s(\alpha)$ satisfying (A.14), which completes the characterization of F_1^α . To summarize,

- For $u \geq u_s(\alpha)$, F_1^α is of the form given in (A.12) with $C_1^\alpha = c^\alpha(u_s(\alpha))$, where u_s^α is the unique solution to (A.14).

- For $u \in [0, u_s(\alpha))$, $F_1^\alpha(u) = \frac{u}{u_s(\alpha)} F_1^\alpha(u_s(\alpha))$.

To prove (i), first note that by the envelope theorem $\frac{d}{d\alpha} c^\alpha(u_s(\alpha)) = \frac{\partial}{\partial \alpha} c^\alpha(u_s(\alpha))$. Using this fact, evaluating the derivative and taking the limit as $\alpha \rightarrow 0$, we get that for $u \geq u_s(0) = u_s$,

$$\lim_{\alpha \rightarrow 0} \left(\frac{d}{d\alpha} F_1^\alpha(u) \right) = \left(\frac{e^{-\frac{\lambda u}{\phi}} (u(\lambda u_s + \phi) - \lambda u_s^2)}{\phi^2 (\lambda u_s + \phi)^2} \right) \times \\ \left(\Pi \lambda \left(\lambda^2 u_s^2 - \phi^2 \left(e^{\frac{\lambda u_s}{\phi}} - 1 \right) + \lambda u_s \phi \right) - c \left(\lambda^2 u_s^2 - 2\phi^2 \left(e^{\frac{\lambda u_s}{\phi}} - 1 \right) + 2\lambda u_s \phi \right) \right).$$

The first-term on the right hand side is clearly positive for $u \geq u_s$. Using (A.14), the second term reduces to $\lambda(\lambda u_s + \phi)((\lambda \Pi - c)u_s - \Pi \phi)$, which is also clearly positive if $\lambda u_s/\phi > \frac{\lambda \Pi/c}{\lambda \Pi/c - 1}$. We now claim that if u_s solves (A.14) for $\alpha = 0$, then this latter inequality must hold. Let $x \equiv \lambda u_s/\phi \geq 0$, $y \equiv \lambda \Pi/c - 2 > 0$, and $\alpha = 0$. The claim is that

$$\frac{e^x}{x+2} = \frac{y+1}{y} \implies x > \frac{y+2}{y+1}.$$

To see that this is true, suppose that $\frac{e^x}{x+2} = \frac{y+1}{y}$ and $x \leq \frac{y+2}{y+1}$. Note that $\frac{e^x}{x+2}$ is strictly increasing. Therefore,

$$\frac{e^x}{x+2} \leq \frac{e^{\frac{y+2}{y+1}}}{\frac{y+2}{y+1} + 2} < \frac{y+1}{y},$$

which gives the contradiction. We have thus shown that at $\alpha = 0$, the derivative of $F_1^\alpha(u)$ w.r.t. α is strictly positive for all $u \geq u_s(\alpha)$. That the same statement is true for $u \in (0, u_s)$ is immediate by the linearity of the value function below u_s . Since F_1^α is also continuously differentiable in both of its arguments, it must be strictly increasing in a neighborhood around $\alpha = 0$ for all $u > 0$, which completes the proof of (i).

We prove (ii) and (iii) for the case of $\alpha \rightarrow 1$, the proof for $\alpha \rightarrow -1$ follows a similar argument. We first show that $\lim_{\alpha \rightarrow 1} u_s(\alpha) = 0$. To do so, rewrite (A.14) as

$$\phi(1-\alpha)e^{\lambda u_s(\alpha)/\phi(1-\alpha)} = \frac{(\lambda \Pi - c(1-\alpha))(\lambda u_s(\alpha) + 2\phi)}{\lambda \Pi - 2c}.$$

Suppose $u_s(1) \equiv \lim_{\alpha \rightarrow 1} u_s(\alpha) \in (0, \infty)$. Then $\lim_{\alpha \rightarrow 1} \phi(1-\alpha)e^{\lambda u_s(\alpha)/\phi(1-\alpha)} = \infty > \frac{(\lambda \Pi - c(1-\alpha))(\lambda u_s(1) + 2\phi)}{\lambda \Pi - 2c}$, a contradiction. Also, clearly $u_s(1) < \infty$ otherwise the principal's value function would be arbitrarily negative. The only remaining possibility is $u_s(1) = 0$.

From (A.12), we have that for $u \geq u_s^\alpha$,

$$\lim_{\alpha \rightarrow 1} \left(\frac{d}{d\alpha} F_1^\alpha(u) \right) = e^{-\frac{\lambda u}{2\phi}} \lim_{\alpha \rightarrow 1} \left(\frac{\lambda u}{4\phi} c^\alpha(u_s(\alpha)) + \frac{\partial}{\partial \alpha} c^\alpha(u_s(\alpha)) \right).$$

Notice from (A.13) that $\lim_{\alpha \rightarrow 1} (\lim_{u \rightarrow 0} c^\alpha(u)) = \lim_{u \rightarrow 0} (\lim_{\alpha \rightarrow 1} c^\alpha(u)) = -(\Pi - 2c/\lambda)$. Therefore, we can conclude that $\lim_{\alpha \rightarrow 1} c^\alpha(u_s(\alpha)) = -(\Pi - 2c/\lambda)$. Hence, to prove (ii), it is sufficient to show that $\lim_{\alpha \rightarrow 1} \frac{\partial}{\partial \alpha} c^\alpha(u_s(\alpha)) = 0$, for this implies $\frac{d}{d\alpha} F_1^\alpha(u) \approx -e^{-\frac{\lambda u}{2\phi}} (\Pi - 2c/\lambda) \frac{\lambda u}{4\phi} < 0$ for $u \geq u_s(\alpha)$ and α sufficiently close to 1. To see that $\lim_{\alpha \rightarrow 1} \frac{\partial}{\partial \alpha} c^\alpha(u_s(\alpha)) = 0$, first notice from (A.14) that

$$\lim_{\alpha \rightarrow 1} \frac{e^{-\frac{\lambda u_s(\alpha)}{\phi(1-\alpha)}}}{(1-\alpha)} \in (0, \infty), \quad (\text{A.15})$$

implying that $u_s(\alpha)$ is $O((1-\alpha)\ln(1-\alpha))$ as $\alpha \rightarrow 1$. Differentiating (A.13) with respect to α and omitting the argument of $u_s(\alpha)$, we get that

$$\begin{aligned} \frac{\partial}{\partial \alpha} c^\alpha(u_s(\alpha)) = & \frac{e^{\frac{2\alpha\lambda u_s}{(\alpha^2-1)\phi}}}{2(1-\alpha)^2\alpha^2(\alpha+1)\lambda\phi(\alpha\phi+\lambda u_s+\phi)^2} \times \left[\lambda\Pi\left((1-\alpha)^2(\alpha+1)^3\phi^3+2\lambda^3u_s^3\alpha(\alpha^2+1)\right.\right. \\ & + \lambda^2u_s^2\phi\left[3\alpha+2\alpha^4e^{\frac{\lambda u_s}{\phi(1-\alpha)}}+\alpha^3\left(5-4e^{\frac{\lambda u_s}{\phi(1-\alpha)}}\right)+\alpha^2\left(2e^{\frac{\lambda u_s}{\phi(1-\alpha)}}-1\right)+1\right]+2(1-\alpha)(\alpha+1)^2\lambda u_s\phi^2) \\ & - (1-\alpha)c\left[(\alpha-1)^2(\alpha+1)^4\phi^3+2\alpha(\alpha^2+1)\lambda^3u_s^3-\lambda^2u_s^2\phi\left(\alpha^4-4\alpha+4\alpha^3\left(e^{\frac{\lambda u_s}{\phi(1-\alpha)}}-1\right)\right.\right. \\ & \left.\left.-4\alpha^2e^{\frac{\lambda u_s}{\phi(1-\alpha)}}-1\right)+2(1-\alpha)(\alpha+1)^3\lambda u_s\phi^2\right] \left. \right]. \end{aligned}$$

Using (A.15), we know that $e^{\frac{2\alpha\lambda u_s(\alpha)}{(\alpha^2-1)\phi}}$ is $O(1-\alpha)$ as $\alpha \rightarrow 1$. Therefore, any term inside the outermost brackets that goes to zero faster than $O(1-\alpha)$ will converge to zero when scaled by the fraction outside the brackets. By inspection, the only terms that do not clearly go to zero faster than $O(1-\alpha)$ are

$$\lambda^2u_s(\alpha)^2\phi\left[2\alpha^4e^{\frac{\lambda u_s(\alpha)}{\phi(1-\alpha)}}-4\alpha^3e^{\frac{\lambda u_s(\alpha)}{\phi(1-\alpha)}}+2\alpha^2e^{\frac{\lambda u_s(\alpha)}{\phi(1-\alpha)}}\right].$$

Thus, we get have that

$$\begin{aligned} \lim_{\alpha \rightarrow 1} \left(\frac{\partial}{\partial \alpha} c^\alpha(u_s(\alpha)) \right) &= \lim_{\alpha \rightarrow 1} \left(\frac{e^{\frac{2\alpha\lambda u_s(\alpha)}{(\alpha^2-1)\phi}} e^{\frac{\lambda u_s(\alpha)}{\phi(1-\alpha)}}}{(1-\alpha)^2(\alpha+1)(\alpha\phi+\lambda u_s(\alpha)+\phi)^2} \lambda u_s(\alpha)^2 [\alpha^2-2\alpha+1] \right) \\ &= \lim_{\alpha \rightarrow 1} \left(\frac{e^{\frac{\lambda u_s(\alpha)}{(1+\alpha)\phi}}}{(\alpha+1)(\alpha\phi+\lambda u_s(\alpha)+\phi)^2} \lambda u_s(\alpha)^2 \right) \\ &= \frac{\lambda}{4\phi^2} \left(\lim_{\alpha \rightarrow 1} u_s(\alpha) \right)^2 \\ &= 0, \end{aligned}$$

which completes the proof of (ii). For (iii), we have

$$\begin{aligned} \hat{F}_1^\alpha(u) - \bar{V}^{\frac{\lambda}{2}}(u) &= \left(\frac{1-\alpha}{2\alpha} \right) \left(\Pi - \frac{(1-\alpha)c}{\lambda} \right) e^{-\frac{\lambda u}{(1-\alpha)\phi}} + C_1^\alpha e^{-\frac{\lambda u}{(1+\alpha)\phi}} - \left(\Pi - \frac{2c}{\lambda} \right) e^{-\frac{\lambda u}{2\phi}} \\ &\leq \left(\frac{1-\alpha}{2\alpha} \right) \left(\Pi - \frac{(1-\alpha)c}{\lambda} \right) + \left| C_1^\alpha - \left(\Pi - \frac{2c}{\lambda} \right) \right| + \left| C_1^\alpha \left(e^{-\frac{\lambda u}{\phi(1+\alpha)}} - e^{-\frac{\lambda u}{2\phi}} \right) \right| \\ &\leq \left(\frac{1-\alpha}{2\alpha} \right) \Pi + |C_1^\alpha - (\Pi - 2c/\lambda)| + \left| C_1^\alpha \left(e^{-\frac{x^*(\alpha)}{2}} - e^{-\frac{x^*(\alpha)}{(1+\alpha)}} \right) \right|, \end{aligned}$$

where the first inequality uses the triangle inequality and $e^{-|x|} \leq 1$ and the second uses the fact that $e^{-\frac{x}{2}} - e^{-\frac{x}{(1+\alpha)}}$ is hump-shaped in x and achieves its maximum at $x^*(\alpha) \equiv 2(1+\alpha)\ln(\frac{1+\alpha}{2})/(\alpha-1)$. Clearly all three terms converge to 0 as $\alpha \rightarrow 1$. $\square$

Proof of Proposition 6.3. Consider first any simple contract with deadline T . If the agent does not shirk,²⁷ then the probability that the project succeeds at $t \in [0, T]$ is given by $\lambda^2 t e^{-\lambda t}$ and the probability that the project does not succeed prior to T is given by $e^{-\lambda T}(1 + \lambda T)$. Therefore, the total surplus is given by

$$S(T) \equiv \int_0^T \lambda^2 t e^{-\lambda t} (\Pi - ct) dt + e^{-\lambda T} (1 + \lambda T) (-cT).$$

In order to induce the agent to work, he must be given rents in the amount of at least $u = \phi T$, otherwise he can do better by shirking. Making the change of variables from T to u , we have that the principal's ex-ante expected payoff under a simple contract with deadline $T = u/\phi$ is bounded above by

$$G(u) \equiv S(u/\phi) - u = \left(1 - e^{-\lambda u/\phi} (1 + \lambda u/\phi)\right) \Pi - \frac{2c}{\lambda} \left(1 - e^{-\lambda u/\phi} (1 + \lambda u/2\phi)\right) - u. \quad (\text{A.16})$$

Note that $G(u)$ is not the principal's value function under the optimal contract, since her belief about the project stage changes over time. We will construct the value function shortly. To prove the proposition, it suffices to show that

- (i) The principal's ex-ante payoff for a project with unobservable progress under *any contract* is bounded above by $\max_u G(u)$.
- (ii) There exists an incentive-compatible simple contract under which the principal's ex-ante expected payoff is $\max_u G(u)$.
- (iii) For all $u > 0$, $G(u) < F_1(u)$. Therefore, the principal does strictly better with intangible progress than she does with unobservable progress.

For (i), let $w^* = \arg \max_u G(u)$, which is generically unique. We have already argued that the principal's maximal payoff under a simple contract is bounded above by $G(w^*)$. Given that neither player has any information about the status of the project, the only possibility is that the principal randomizes over the termination date. It therefore suffices to show that the principal cannot benefit from such randomization, or equivalently, that the principal's value function (under this simple contract with the optimally chosen deadline) is globally concave in the agent's continuation value.

Suppose that the principal can implement a simple contract in which the incentive compatibility condition holds with equality for all t (this is the best possible case for the principal). It is most intuitive to construct this value function from the pair of value functions that are conditional on s (i.e., whether a breakthrough has been made) and weight them appropriately by the probability that the principal assigns to each. Given u , the principal's payoff conditional on being in the first stage ($s = 1$) is $G(u)$; i.e.,

$$\begin{aligned} F^{unobs}(u|s=1) &= \int_0^{u/\phi} \lambda^2 t e^{-\lambda t} (\Pi - ct) dt - e^{-\lambda u/\phi} (1 + \lambda u/\phi) cu/\phi - u \\ &= (1 - e^{-\lambda u/\phi}) (\Pi - 2c/\lambda) - \frac{\lambda u}{\phi} e^{-\lambda u/\phi} (\Pi - c/\lambda) - u. \end{aligned}$$

Conditional on being in the second stage, the principal's value function is the benchmark payoff

²⁷Arguments similar to those made for a single-stage project can be used to confirm shirking is suboptimal.

$\bar{V}(u)$:

$$F^{unobs}(u|s=2) = \int_0^{u/\phi} \lambda e^{-\lambda t} (\Pi - ct) + e^{-\lambda u/\phi} (-cu/\phi) - u = (1 - e^{-\lambda u/\phi})(\Pi - c/\lambda) - u.$$

Over time, the principal's beliefs will evolve about the state of the project. Conditional on reaching state $u < w^*$ prior to project success, a period of time of length $t(u; w^*) = \frac{w^* - u}{\phi}$ has elapsed. Therefore, the principal's beliefs are given by

$$\mu(u; w^*) = \Pr(\tau_1 \leq t(u; w^*) | \tau_2 > t(u; w^*)) = \frac{\lambda \left(\frac{w^* - u}{\phi} \right)}{1 + \lambda \left(\frac{w^* - u}{\phi} \right)}.$$

The principal's value function for $u \leq w^*$ is therefore given by

$$F^{unobs}(u; w^*) = \mu(u; w^*) F^{unobs}(u|s=2) + (1 - \mu(u; w^*)) F^{unobs}(u|s=1).$$

We will now verify this value function is concave for all $u \leq w^*$. Using the functional forms given above and twice differentiating $F^{unobs}(u; w^*)$, we get that

$$\begin{aligned} \frac{d^2}{du^2} F^{unobs}(u; w^*) &= \frac{-\lambda e^{-\lambda u/\phi}}{\phi^2 (\lambda(w^* - u) + \phi)^3} \left[(\lambda^3 w^* (w^* - u)^2 + \lambda w^* \phi^2) (\lambda \Pi - c) \right. \\ &\quad \left. + \lambda^2 \phi (w^* - u)^2 (\lambda \Pi - 2c) + \phi^3 (\lambda \Pi + 2(e^{\lambda u} - 1)) \right]. \end{aligned}$$

All three terms inside the brackets are clearly positive, implying the value function is concave in u for all $u \leq w^*$, which completes the proof of (i).

We prove (ii) by showing that for any T , there exists a $w : [0, T] \rightarrow \mathbb{R}_+$ such that (1) it is incentive compatible for the agent to work for all $t \in [0, T]$, and (2) the agent's continuation utility at date t is $u(t) = \phi(T - t)$. Let $u_2(t)$ be the promised continuation value conditional on being in the second stage and $u(t)$ be the unconditional continuation value at t . Promise keeping requires that

$$\lambda u(t) = \lambda(\mu(t)w(t) + (1 - \mu(t))u_2(t)) + u'(t).$$

Conditional on progress, the evolution of u_2 is given by

$$\lambda u_2(t) = \lambda w(t) + u'_2(t). \tag{A.17}$$

We want to find $w(t)$ such that $u(t) = \phi(T - t)$ for all $t \in [0, T]$. Note that this implies that $u'(t) = -\phi$. Using the promise keeping condition,

$$\phi(T - t + 1/\lambda) = (1 - \mu(t))u_2(t) + \mu(t)w(t).$$

Substituting for $w(t)$ from (A.17), we get that

$$\begin{aligned} \phi \left(T - t + \frac{1}{\lambda} \right) &= (1 - \mu(t))u_2(t) + \mu(t) \left(u_2(t) - \frac{u'_2(t)}{\lambda} \right) \\ &= u_2(t) - \frac{t}{1 + \lambda t} u'_2(t). \end{aligned}$$

Imposing the boundary condition $u_2(T) = 0$, we arrive at a unique solution for $u_2(t)$, which we can

then substitute back into (A.17), to arrive at

$$w(t) = \phi \left(T - t + \frac{1}{\lambda} + \frac{e^{-\lambda(T-t)}}{\lambda^2 T} + \frac{e^{\lambda t}}{\lambda} (q(-\lambda T) - q(-\lambda t)) \right),$$

where $q(z) = -\int_{-z}^{\infty} e^{-x}/x dx$. It is straightfoward to check that $w(t) > 0$ for all $t \in [0, T]$, which completes the proof of (ii). For (iii), note that $G(u)$ is identical to $\hat{F}_1(u)$ with $H_1^c = \frac{2c}{\lambda} - \Pi$ (see (22)). The result then follows from the fact that F_1 has the same form as $\hat{F}_1$, but with a constant strictly larger than $\frac{2c}{\lambda} - \Pi$. $\square$