Breakthroughs, Deadlines and Severance: Contracting for Multistage Projects*

Brett Green
UC Berkeley (Haas)
greenb@berkeley.edu
faculty.haas.berkeley.edu/bgreen

Curtis R. Taylor
Duke University
crtaylor@duke.edu
people.duke.edu/∼crtaylor

October 9, 2014

Preliminary and Incomplete

Abstract

We study the optimal provision of incentives in a dynamic multistage agency setting. Specifically, there is a project, whose benefits are realized only after multiple sequential breakthroughs. The project requires funding from the principal and is operated by an agent who is protected by limited liability and can divert cash flows for private benefit. The nature of progress plays an important role in the optimal incentive scheme. When breakthroughs are *tangible* (i.e., observable and contractible), they serve as a useful monitoring device, and the principal controls incentives through a sequence of deadlines and a reward scheme that decreases over time; a breakthrough in one stage extends the amount of time the agent has to complete the next stage. When breakthroughs are *intangible* (i.e., privately observed by the agent), the principal must provide incentives for truthful self-reported progress. Our primary questions are whether, to what extent, and how this type of communication can be used in a meaningful way.

*We are grateful to Alessandro Bonatti, John Morgan and Steve Tadelis for useful comments and suggestions. We also thank conference participants at ESSFM Gersenzee and seminar participants at MIT Sloan and the University of Miami for useful feedback.

1 Introduction

Contracting environments often require completion of sequential stages before their benefits can be realized. Industries which are R&D intensive are especially relevant. To fix ideas, consider the stages involved in the development of a new drug. Doing so requires first identifying the specific chemical compounds or molecule to treat the disease in vitro (i.e., in test tubes), and then demonstrating its efficacy in pre-clinical (i.e., animal) trials. Once these two are successfully completed, clinical (i.e., human) trials progress through three sequential stages: dosing, small scale and large scale. All of these stages must be successfully completed before a drug can be brought to the FDA and considered for approval, and only after approval can the drug be brought to market and generate revenue for its developer and health benefits for society. Additional relevant applications of multistage projects are ubiquitous; construction, product and software development, venture capital, basic research and so on. Naturally, complex projects such as these require substantial amounts of funding in order to come to fruition and the preferences of the employee, researcher, or scientist (the "agent") regarding the timing, intensity and direction of investment are unlikely to be completely aligned with the preferences of the firm, institution, or funding entity (the "principal"). In this paper we develop a theory to investigate optimal provision of incentives in a class of multistage agency problems.

The environment we study features a project that requires successful completion of multiple stages, referred to as *breakthroughs*, in order for the principal to realize the project benefits. The model is set in continuous time with breakthroughs arriving randomly according to a Poisson process. The arrival rate of breakthroughs depends on whether the principal funds the project and the fraction of the funds the agent diverts for private benefit (or, equivalently, whether the agent shirks). Both principal and agent are risk-neutral — the agent, however, has limited funds to invest and is protected by limited liability.[1]

Within our multistage environment, there is a rich space of contracts and information structures to explore. We are especially interested in how the nature of *progress* influences the provision of incentives and the value of the project. In some contexts, progress is a tangible object that can be easily observed by the principal (e.g., whether a software subroutine works). In other contexts, progress is less well defined and infeasible for the principal to gauge, in which case it may become a source of private information for the agent (e.g., complex experimental data).[2] We say that progress is *tangible* if it is observable to both parties and

[1]Without limited liability, the first best outcome can be sustained by simply "selling" the project to the agent. This solution is practically infeasible for most of the relevant applications under consideration.

[2]Recent examples of extensive data falsification and alleged falsification are the respective cases of Diederik Stapel, formerly a professor of social psychology at Tilburg University, and Anil Potti, formerly a cancer

can be contracted upon. On the other hand, *intangible* progress is privately observed by the agent and hence, cannot be directly contracted upon.[3] We investigate the following set of questions in particular: When progress is tangible, in what way is it used to incentivize the agent? Should the agent be given a limited amount of time to make each breakthrough? Or, does the principal control incentives solely through dynamically adjusting "reward" payments for ultimate success? When progress is intangible, does communication have value? More specifically, is it possible for the principal to elicit "self reports" from the agent and use these reports to provide stronger incentives and increase the value of the project? Or, can the principal do just as well by ignoring any self-reported progress made by the agent?

To build intuition for our results and provide a baseline for comparison, we start by analyzing a single-stage project in which there is one tangible breakthrough needed for the benefits to be realized. In this case, the optimal contract takes a fairly simple and intuitive form. The principle provides the agent with a deadline T^* and a reward schedule, which is decreasing over time. If the agent realizes the breakthrough at $\tau \leq T^*$, he collects the reward of $W(\tau)$. If a breakthrough is not realized by the deadline, the principal terminates the project (i.e., discontinues investment indefinitely) and both players get their outside option. Note that the use of a deadline plays an important role in the provision of incentives. This is despite the fact that the first-best outcome involves the agent working indefinitely until the breakthrough arrives.

We then incorporate 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 but it means the project progresses to the ultimate stage. If and when a breakthrough in the ultimate stage occurs, the project benefits are realized and thus a breakthrough in the ultimate stage is tangible.

When preliminary breakthroughs are also tangible, they can serve as a useful monitoring device. In the first stage, the principal provides incentives to the agent through three separate channels. The first channel is a first-stage deadline, denoted T_1. If the project does not progress to the second stage by the deadline, then the principle discontinues investment and the project is terminated. The first stage deadline need not be deterministic. For some parametric cases, the principal uses a *soft deadline*: for $t > T_1$, the principal randomizes over whether to terminate the project. Soft deadlines are useful because they allow the principal to provide strong incentives in the first stage, while simultaneously not reducing the probability

researcher at Duke University. Both Stapel's misconduct and Potti's alleged misconduct went undetected by their employers, funding agencies, and professions for a number of years.

[3] Yet another possibility is that progress is not directly observed, but it can be verified, perhaps at some cost (e.g., checking the proof of a theorem). See Section 5 for more discussion.

of success in the second stage conditional on making a breakthrough before the soft deadline.

The second channel is a second-stage deadline that (weakly) decreases with the amount of time it takes the agent to make his first breakthrough. That is, the earlier the agent makes the first breakthrough, the more time he gets to make the second one. In expectation, the second stage deadline, denoted $T_2(\tau_1)$ is strictly greater than T_1. Therefore, the optimal contract can be interpreted as putting the agent on a "short leash" until progress is made at which point, he is given a longer leash to complete the second stage. Finally, the third channel is a monetary reward schedule paid to the agent only upon the ultimate success of the project. The reward schedule is also (weakly) decreasing in both the time of the preliminary breakthrough τ_1 and the ultimate breakthrough τ_2.

For the case of intangible progress, our results are more preliminary and speculative at this point. Several preliminary results are as follows. First suppose that there is no communication between principle and agent regarding how the project has progressed. In this case, if the principle uses a deadline, some degree of cash flow diversion will occur. This is because an agent who has not yet made any progress has probability proportional to $(dt)^2$ of having two breakthroughs sufficiently close to the deadline, whereas diverting investment for private benefit allows him to realize a payoff proportional to (dt). Hence, an agent who has not yet had a breakthrough will "run out of steam" and begin diverting cash flows as the deadline approaches. However, cash flow diversion is socially inefficient and such a contract can be improved upon by enriching the contract space to allow for two additional features: communication and severance.

We introduce communication by allowing the agent to submit *self reports* telling the principal if and when he has made a breakthrough. An important observation is that the optimal contract that was derived in the case of tangible progress will not induce truthful reporting when the progress is intangible. To see why, suppose the agent has not yet had a breakthrough and time is nearing T_1. Rather than being fired at T_1, the agent strictly prefers to falsely report a breakthrough to "extend his clock" and, by doing so, obtain a payoff equal to the value of diverting investment for the rest of his tenure.

By the Revelation principle, it is sufficient to consider contracts in which truthful reporting is induced. To incentivize truth telling requires that two new constraints be satisfied. We refer to them as *No False Progress* (NFP) and *No Hidden Progress* (NHP). The former ensures a deviation of the form described above is not profitable and the latter ensures that an agent who has made a breakthrough prefers to report it immediately rather than "hide" it for some period of time.

We present some preliminary results about the structure of the optimal mechanism in the intangible progress setting and formulate several conjectures. For instance, the "short-leash, long-leash" scheme that the principal implements under tangible progress is unlikely to be optimal when progress is intangible. This is because long second-stage deadlines necessitate large first-stage severance payments in order to provide incentives for the agent to truthfully report lack of progress. Severance payments, of course, confer information rents to the agent at the principal's expense. Thus, we conjecture that when progress is intangible the optimal mechanism involves a relatively long initial stage and a relatively short final stage, reversing the optimal structure under tangible progress.

In the next section, we discuss related literature. In Section 2, we present a single-stage version of the model as a benchmark. In Section 3, we consider a two-stage project with tangible progress. Section 4 analyzes a two-stage project with intangible progress. Section 5 discusses several extensions. Section 6 concludes.

1.1 Related Literature

There is a large and growing literature studying the optimal provision of incentives in dynamic environments. A 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). Our single-stage model benchmark is most similar to Hopenhayn and Nicolini (1997), who looks at incentive provision to search for employment while simultaneously providing unemployment insurance, and Mason and Välimäki (2011) who consider the role of commitment in dynamic moral hazard problem.[4] A novel aspect of our model is that we explicitly consider environments with multiple sequential stages.

Our multistage environment with tangible progress is closely related to Biais, Mariotti, Rochet, and Villeneuve (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, Livdan, and Tchistyi (2014).

When progress is intangible, 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

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

continuum of types using a first order approach by Williams (2011) and Edmans, Gabaix, Sadzik, and Sannikov (2012). Our solution approach is most similar to Zhang (2009). One key difference in our environment is the presence of public (contractible) information (i.e., the ultimate success of the project), which is informative about the agent's underlying private information. Another difference from Zhang (2009), is that the transitions probabilities across states are endogenously determined in our setting. Optimal dynamic mechanisms are explored in Board (2007), Eso and Szentes (2007), Bergemann and Valimaki (2010) and Pavan, Segal, and Toikka (2014) among others. One of our key questions is to understand the value of communication (i.e., using a direct revelation mechanism).

There is also an extensive literature that studies the role of learning in contractual environments and how to incentivize experimentation in bandit type problems. Examples include Levitt and Snyder (1997) , Inderst and Mueller (2010), Manso (2011), Hörner and Samuelson (2014), Halac, Kartik, and Liu (2013). In our model, there is no learning or social value of information, the principal only uses information for the sole purpose of providing stronger incentives to the agent.

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.

Relational contracting and environments in which evaluations are either subjective or privately observed by the principal have been studied by Baker, Gibbons, and Murphy (1994), Baker (2000), Macleod (2003), Levin (2003), Fuchs (2007), and Maestri (2012). The primary difference in these models is that the principal has subjective information and must decide how to optimally use it whereas, in our setting, the agent endogenously acquires private information over time and the principal must decide how best to elicit it.

2 Single-Stage Project

A principal (she) contracts with an agent (he) to undertake a project over a potentially infinite time horizon. Both parties are risk-neutral and do not discount the future.[5] The principal is not financially constrained, while the agent has no financial resources and is protected by limited liability.[6]

[5]This assumption is for expositional purposes. Incorporating a common discount rate, $r > 0$, into the model is a fairly straightforward exercise. Throughout the paper, we remark on when results differ when discounting is incorporated.

[6]Without financial constraints and limited liability for the agent, the first-best outcome can easily be sustained by simply "selling" the project to the agent. However, this solution is practically infeasible for

The project requires funding from the principal, at flow cost c, in order to operate. Provided the principal has not yet terminated the project, the agent chooses whether to divert funds for private benefit at each instant of time. That is, the agent chooses an action $a_t \in [0,1]$ which represents the fraction of the funds that are invested in the project and $(1-a_t)$ represents the fraction that are diverted for benefit. The arrival rate of success is then given by λa_t. Thus, if the agent does not divert any funds over an interval of length Δ, then the probability that the project is successful during the interval is $1-e^{-\lambda\Delta}$. If the principal terminates the project, the arrival rate of success is zero forever after. Upon the arrival of success the principal realizes a benefit $b>0$, makes any outstanding promised payments to the agent and the game ends. The agent receives no intrinsic benefit from a success, he benefits solely from the compensation delivered by the principal.

The principal cannot observe whether the agent invests or diverts the funds she supplies. If the agent diverts c dollars for a period of length Δ then he receives a private benefit of $\phi c\Delta$, where $\phi \in (0,1)$ measures the severity of the agency problem.[7] For notational convenience define $\mu \equiv \frac{1}{\lambda}$ to be the expected amount of time before experiencing success if the agent invests indefinitely. We assume the parameter values satisfy

$$b - c\mu > \phi c\mu. \tag{1}$$

Intuitively, the principal (essentially) must give the agent ϕc per period in continuation value in order to prevent diversion in addition to the flow cost c of investment. Equation 1, therefore, ensures that the principal's benefit, b, outweighs her expected cost $(1+\phi)c\mu$. Note that this inequality is stronger than would be needed in a first-best situation where the principal ran the project herself, namely $b - c\mu > 0$. Further, note that the first-best policy (absent the agency conflict) is to invest indefinitely since the expected flow benefit is larger than the cost.

At $t=0$, the principal offers the agent a contract. If the agent rejects the contract offer, then both parties receive their outside options. A contract is a triple, and denoted $\Gamma = \{a, W, T\}$, where a_t is the recommended action to the agent at date t, dW_t is a monetary payment made to the agent at time t, T is a termination rule (perhaps random or infinite).

Limited liability requires that payments to the agent must be non-negative, $dW_t \geq 0$. Given a recommended action profile, $a = \{a_t, 0 \leq t \leq \infty\}$ let τ denote the random variable for the arrival time of success.

most of the relevant applications under consideration.

[7] We model the agency environment as a cash flow diversion (CFD) problem. It is well-known (e.g., DeMarzo and Fishman (2007)) that CFD models are formally equivalent to hidden effort models in which the agent must be incentivized to exert effort.

The principal faces a constrained maximization problem. Her objective function is

$$\max_{\{\tau,W,a\}} \mathbb{E}^a \left[b \cdot \mathbb{1}_{\{\tau \le T\}} - \int_0^{T\wedge\tau} (cdt + dW_t) \right]$$

Subject to the incentive compatibility constraint

$$a \in \arg\max_{\tilde{a}} \mathbb{E}^{\tilde{a}} \left[\int_0^{T\wedge\tau} (\phi c(1-\tilde{a}_t)dt + dW_t) \right],$$

and the participation constraint

$$\mathbb{E}^a \left[\int_0^{T\wedge\tau} (\phi c(1-\tilde{a}_t)dt + dW_t) \right] \ge \underline{U}. \tag{2}$$

Lemma 2.1. *For a single-stage project, there exists a solution to the principal's problem such that (i) T is deterministic and finite, (ii) $a_t = 1$ for all $0 \le t \le T$, and (iii) $dW_t = 0$ for all $t \ne \tau$.*

Lemma 2.1 narrows the set of contracts one needs to consider and simplifies the formulation of the principal's problem. Following standard techniques, we formulate it recursively as a dynamic programming problem where the state variable is given by the agent's continuation value. For any $t < T$, let U_t be the expected continuation utility of the agent under the contract if he has not succeeded by date t, let $w(U_t)$ denote the promised utility (i.e., monetary payment) to the agent for success, and let $V(U_t)$ denote the principal's expected continuation payoff. Given any incentive compatible contract, for $t \le T$, the agents continuation value evolves according to[8]

$$\lambda U_t = \lambda w(U_t) + U_t'. \tag{3}$$

The incentive compatibility constraint, which ensures that locally the agent prefers investing to diverting funds for any $t \le T$ is given by

$$\lambda \left(w(U_t) - U_t \right) \ge \phi c. \tag{4}$$

Finally, that the project is terminated at T implies that both principal and agent get their respective outside options at that point. We normalize both outside options to zero for notational convenience

$$V(U_T) = 0, \;\; U_T = 0. \tag{5}$$

Note that (3)-(5) immediately implies that $U_t \ge 0$ for all $t \in [0, T)$ and hence the agent's

[8] Expression (3) as well as all other differential equations appearing below can be derived as the limit of a discrete-time model in which the probability of succeeding in a period of length Δt is $\lambda \Delta t$.

participation constraint is satisfied. The principal's value function, $V(U_t)$, solves the Bellman equation

$$\begin{aligned} \lambda V(U_t) &= \max_{W, U_t'} \lambda(b - w(U_t)) - c + V'(U_t)\lambda(U_t - w(U_t)) \\ & \quad \text{s.t. } (\mathrm{IC}_1)\text{-}(5). \end{aligned} \tag{6}$$

Note that $V'(u) \geq -1$ for all u, since the principal can always make direct monetary payments. It is then straightforward to show that the incentive compatibility condition, equation (IC_1), must bind with equality for all U_t; if it were slack, the principal could simply reduce w and increase T. This modification leads to a strict improvement in the principal's payoff. That (IC_1) binds everywhere pins down w. Moreover, it implies that prior to a breakthrough, the agent's continuation utility falls proportional to time: $U_t' = -\phi c$. Combining this fact with the boundary condition in equation (5), equation (6) implies that the principal's value function is given by

$$V^*(u) = \left(b - \frac{c}{\lambda} - u\right) - \left(b - \frac{c}{\lambda}\right) e^{-\lambda u/\phi c}. \tag{7}$$

The first term on the right-hand side represents the first best value of the project net of delivering u to the agent. The second term is the *agency cost* required to prevent cash flow diversion. Notice that $V^*(u)$ is strictly concave, which confirms that randomization over the termination deadline is suboptimal.

The only choice left to the principal is to select the optimal initial utility for the agent, or (equivalently) the optimal termination date of the contract. Hence, she solves

$$\max_{u \in \mathbb{R}_+} V^*(u). \tag{8}$$

Proposition 2.2 (Single-Stage Optimal Contract)**.** *The solution to the principal's problem can be implemented with a deadline*

$$T^* \equiv \mu \ln\left(\frac{b - c\mu}{\phi c \mu}\right), \tag{9}$$

and a reward payment for success that is decreasing over time according to $w_t = \phi c(\mu + T^* - t)$.[9]

The parametric restriction in (1) ensures that $T^* > 0$. Note that the optimal deadline is increases as the agency problem becomes less severe, with $\lim_{\phi \to 0} T^* = \infty$. That is, the (second-best) outcome and principal payoffs converge to first-best as the agency conflict goes to zero. As we noted earlier, if the principal could run the project herself at the same

[9] Interestingly, the same expression (9) holds if the principal and agent both discount the future at rate $r > 0$.

cost, then her optimal policy is to invest c until the project is successful regardless of how long it takes. Hence, the finite deadline T^* exists only to mitigate agency costs. Because success essentially ends the contractual relationship, the agent's reward upon succeeding must compensate him not only for the instantaneous incentive to divert funds, but for the expected value of the option to divert funds for the rest of the contractual relationship. The deadline T^* ensures that the option value to the agent of diverting funds is finite and falls over time.

3 Two-Stage Project with Tangible progress

With the single-stage benchmark in hand, let us now introduce a second stage. Henceforth, suppose the agent must make two breakthroughs in order for the principal to realize the project benefits (i.e., the principal obtains benefit b only after the second breakthrough has been achieved). We will distinguish between the first stage of the project (before the first breakthrough) and the second stage (after the first breakthrough and before the second). To simplify notation and to facilitate comparison, we assume that λ, c, and ϕ are the same in both stages. Finally, in this section we assume that progress is *tangible*. That is, the time of the first (and second) breakthroughs is observable to both players and can be contracted upon. We take up the case of *intangible* progress in Section 4.

As before, a contract consists of $\Gamma = \{a, W, T\}$, where each of the elements can depend on the entire *history*, which we denote by h_t, is summarized by the triple (t, τ_1, τ_2): t denotes the total time elapsed; $\tau_s \in \emptyset \cup [0, t)$ denotes the elapsed time in stage s at which a breakthrough arrives, where $\tau_s = \emptyset$ indicates that a breakthrough in stage s has not occurred.

Since progress is tangible (i.e., there is no persistent private information), the optimal contract can be derived recursively using the state variable, (s, U_t), where s denotes the current stage and U_t denotes the promised utility to the agent. Analogous to Lemma 2.1, one can show that the optimum can be implemented without cash flow diversion (i.e., $a = 1$ prior to termination). For any Γ, we let $w(s, u)$ denote the implied *transitional* promised utility delivered to the agent conditional on realizing a breakthrough in stage s when current promised utility is u. Similar to the single-stage case, providing incentives not to divert requires that

$$\lambda \left(w(s, U_t) - U_t \right) \geq \phi c. \tag{IC}$$

Unlike Lemma 2.1, when multiple stages remain, the use of random project termination or *soft deadlines* may be required to achieve the optimum. Let $\sigma(s, u)$ denote the rate at which the project is terminated (i.e., $\sigma(s, u) = \lim_{dt \to 0} \frac{\Pr(T \in (t, t+dt) | s, u)}{dt}$). Omitting arguments, the promise keeping condition requires

$$U_t' = (\lambda + \sigma)U_t - (\lambda w + \sigma P), \tag{PK}$$

where P denotes the severance payment to the agent upon termination.

We solve for the optimal contract by backward induction on the number of stages remaining. Note the optimal contract in the second stage (conditional on u) is the same as that from the single stage problem. Specifically, it involves (i) $w(2, U_t) = \phi c\mu + U_t$ and (ii) $U_t' = -\phi c$, until $U_t = 0$ at which point the project is terminated. We therefore take the second stage value function as given and use it to solve the principal's problem in first stage. To do so, define $V_2(u) \equiv V(u)$ (from equation (7)) to be the principal's value function in the second stage. The principal's value function in the first stage solves

$$\begin{aligned} \lambda V_1(u) = \max_{\{w,\sigma,P\}\geq 0} & \lambda V_2(w) - c - \sigma(V_1(u) + P) + V_1'(u)U_t' \\ \text{s.t. } & \text{(IC), (PK)} \end{aligned} \tag{10}$$

Note that the principals objective is weakly decreasing in P and any contract in which $P > 0$ can be weakly improved upon by increasing w and setting $P = 0$ without violating either constraint (recall that $V_2' \geq -1$). Therefore, as in Section 2, using a severance payment is suboptimal. Moreover, randomization is optimal (i.e., $\sigma > 0$) if and only if the value principal's value function has a convex portion when randomization is not permitted. We can therefore proceed by assuming $\sigma(1, u) = 0$ for all u and then check whether the resultant value function is concave.

We adopt a conjecture and verify technique by supposing that (IC) always binds in stage 1 and then verifying that this cannot be improved upon. The supposition is $w(1, u) = u + \phi c\mu$, which combined with (PK) implies $U_t' = -\phi c$. Hence, our candidate V_1 satisfies

$$\lambda V_1(u) = \lambda V_2(u + \phi c\mu) - c + V_1'(u)(-\phi c),$$

which has a solution of the form

$$b - \frac{2c}{\lambda} - u - \frac{\lambda u}{\phi c}\left(b - \frac{c}{\lambda}\right) e^{-(1+\frac{\lambda}{\phi c}u)} + K_1 e^{-\frac{\lambda}{\phi c}u}.$$

The terminal boundary condition $V_1^c(0) = 0$, gives $K_1 = \frac{2c - \lambda b}{\lambda}$. And thus we are left with a

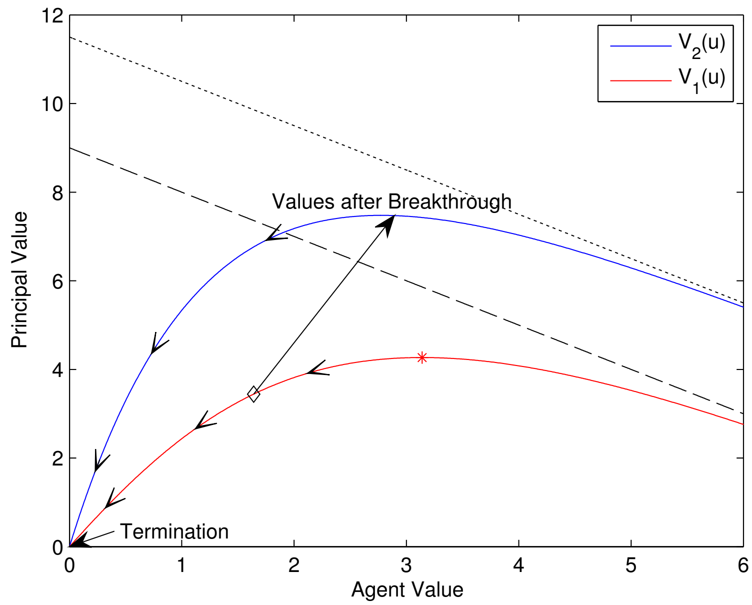

Figure 1: This figure illustrates the value functions and optimal dynamics for the case when (12) holds. Starting from the red asterisk, the continuation values follow the left pointing arrows on the red line until either (1) a breakthrough occurs (e.g., at the black diamond) at which point they jump upward to the blue line as indicated by the solid arrow or (2) the origin is reached at which point the project is terminated.

candidate value function

$$V_1^c(u) = \left(b - \frac{2c}{\lambda} - u\right) - \left(b - \frac{2c}{\lambda}\right) e^{-\frac{\lambda}{\phi c}u} - \frac{\lambda u}{\phi c}\left(b - \frac{c}{\lambda}\right) e^{-(1+\frac{\lambda}{\phi c}u)}. \tag{11}$$

Similar to the value function in equation (7), the first term on the right-hand side above is the first-best value of the project net of delivering u to the agent, while the remaining terms represent the (expected) agency costs. Given the closed form expression for the candidate value function, it is easy to check for its concavity.

Lemma 3.1. *V_1^c is concave for all $u \geq 0$ if and only if*

$$\lambda b \geq \frac{2c(e-1)}{e-2}. \tag{12}$$

When V_1^c is concave, the principal cannot benefit from randomizing. Therefore, if (12) holds then (11) characterizes the principal's value function in the first stage under the optimal contract.

Figure 1 illustrates the value function and optimal dynamics in this case. The optimal incentive scheme can be implemented with a remarkably simple structure. At the beginning of the contract the principal sets two clocks, a short one with length T_1 and a long one with length $\mu+T_1$. If the agent has not made a breakthrough by T_1, then the project is terminated. On the other hand, if he makes the first breakthrough at $\tau_1 < T_1$, then he gets the remaining time on the long clock, $\mu + T_1 - \tau_1$, to make the second breakthrough. Thus he is rewarded for an early first-stage breakthrough with more time to make a second-stage one, and he is rewarded by making an early second-stage breakthrough with a higher cash payment. The details are summarized in the following result.

Proposition 3.2. *Suppose that* (12) *holds. Then, the optimal contract can be implemented with a hard deadline in the first stage* T_1^**, a conditional (hard) deadline in the second stage* $T_2^*(\tau_1) = \mu + T_1^* - \tau_1$ *and the reward schedule* $w(\tau_1, \tau_2) = \phi c(\mu + T_2^*(\tau_1) - \tau_2)$ *such that:*

- *If the agent is not successful in the first stage prior to* T_1*, then the project is terminated.*
- *If the agent succeeds in stage 1 at* $\tau_1 < T_1$*, he is given additional time* $T_2^*(\tau_1)$ *to complete the second stage. If the agent does not complete the second stage prior to prior to* $T_2(\tau_1)$*, then the project is terminated.*
- *The agent is rewarded only if the project ultimately succeeds and in the amount* $w(\tau_1, \tau_2)$*.*

3.1 Soft Deadlines

If (12) does not hold, then there exists a $\bar{u}$ such that V_1^c is convex on $[0, \bar{u}]$. To characterize the principal's value function for this case, we find the steepest possible line through the origin that intersects V_1^c. That is, let

$$u_S = \arg\max_u \frac{V_1^c(u)}{u}.$$

Proposition 3.3. *If* (12) *does not hold, then the principal's value function under the optimal contract is given by* $V_1(u) = \frac{V_1^c(u_S)}{u_S} u$ *for* $u \in [0, u^S]$ *and* $V_1(u) = V_1^c(u)$ *and for* $u \geq u_S$*.*

In this case, randomizing over termination dates is needed to optimally provide incentives. That is, the principal gives the agent a *soft deadline*, after which the project may be terminated, or it may not. Note that a breakthrough in the first stage has no value to the principal without its counterpart in the second stage. Thus, the soft deadline can be useful because it allows the principal to provide incentives in the first stage without further reducing the second stage deadline. Conditional on not terminating the project, the principal effectively "stops the clock", thus preserving the the probability of ultimate success conditional on reaching the second stage.

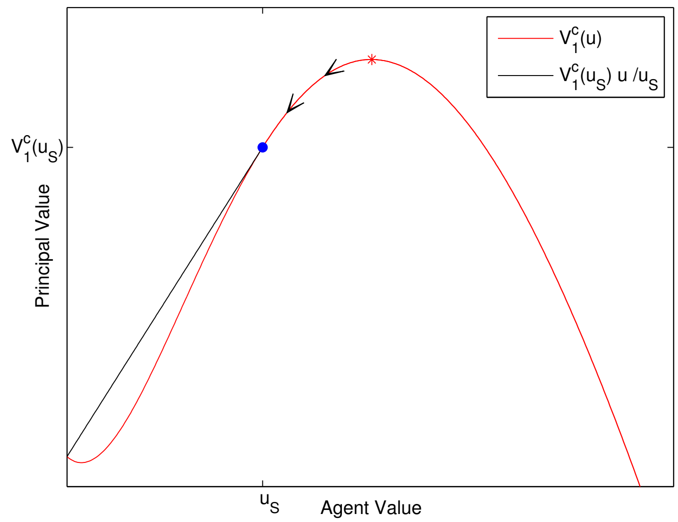

Figure 2: This figure illustrates the value function and dynamics in the first stage when (12) does not hold. Starting from the red asterisk, the continuation values follow the left pointing arrows on the red line until either (1) a breakthrough occurs at which point they jump upward as in Figure 1 (not pictured here) or (2) the blue dot is reached at which point the principal randomizes between termination (a jump to the origin) and remaining at the blue dot. Conditional on making a breakthrough from the blue dot, continuation values jump to the same point regardless of the time at which the breakthrough occurs.

Proposition 3.4. *Suppose that (12) does not hold. Then, the optimal contract can be implemented with a soft deadline T_1^{**} in the first stage, a conditional (hard) deadline in the second stage $T_2^{**}(\tau_1) = \mu + \frac{1}{\sigma} + (T_1^{**} - \tau_1)^+$, a termination rate σ and the reward schedule $w(\tau_1, \tau_2) = \phi c(\mu + T_2^{**}(\tau_1) - \tau_2)$ such that:*

- *If the agent does not have a breakthrough in the first stage prior the soft deadline, then the principal (randomly) terminates the project at rate σ.*
- *If the agent has a breakthrough in the first stage at τ_1, he is given additional time $T_2^{**}(\tau_1)$ to complete the second stage. If the agent does not succeed in the second stage prior to $T_2^{**}(\tau_1)$, then the project is terminated.*
- *The agent is rewarded only if the project ultimately succeeds prior to being terminated and in the amount $w(\tau_1, \tau_2)$.*

3.2 Deadlines: First stage versus Second stage

One natural question is whether the agent should have more time to make a breakthrough in the first stage or in the second stage.

Proposition 3.5. *Conditional on making a breakthrough in the first stage, the expected amount of time the agent has to complete the second stage is strictly larger than the (expected) amount of time the agent has to complete the first stage.*

In other words, the principal keeps the agent on a relatively "short leash" until he makes the initial breakthrough. After that, he is granted a longer horizon to bring the project to fruition. This result seems intuitive but it depends critically on the tangibility of progress. When the principal uses only hard deadlines, conditional on an ultimate success, the total payment to the agent depends only on the total time to completion (i.e., $\tau_1 + \tau_2$). This same feature is not true when the principal uses a soft deadline. Specifically, conditional on not being terminated prior to the first breakthrough, the reward is constant over any $\tau_1 > T_1^{**}$ beyond when the principal starts randomizing. Hence, for any such τ_1 the agent is compensated *as if* the total time to completion was $T_1^{**} + \tau_2$ regardless of when the first breakthrough actually occurred.

4 Intangible Progress

We now suppose that progress is intangible. That is, the first breakthrough is not observed by the principal and hence the time at which it occurs, τ_1, cannot be directly contracted upon. We maintain that the second breakthrough is tangible (i.e., the agent can neither falsify nor hide the ultimate success of the project). First, we consider a limited space of contracts in which communication between the principal and agent is not possible. This specification seems quite natural for certain situations in which communicating may be prohibitively difficult or costly. Furthermore, it is not obvious that being able to condition on self reports from the agent has any value to the principal. After all, the only reason for the principal to use this information is to strengthen the agents incentives and reduce the rents he can extract. And given that the agent must be provided with incentives to report truthfully, whether this information can be used in a meaningful way is one of the key questions we wish to answer. In Section 4.2, we present some preliminary analysis on this point.

4.1 No Communication

Suppose that there is no communication between the principal and the agent and consider any reward scheme w_t (for ultimate success) and deterministic termination rule, T. When termination is near and the agent has not yet progressed to the second stage (i.e., for t_1 close to T), it becomes extremely costly for the principal to satisfy incentive compatibility.

Lemma 4.1. *For any bounded reward scheme w and deterministic termination rule, T, there exists a $\bar{t} < T$ such that if the agent has not progressed to the second stage by $\bar{t}$, he will divert the entire investment (i.e., $a_t = 0$) for all $t \in (\bar{t}, T)$.*

The intuition is actually quite straightforward. A simple calculation yields that

$$\Pr(\tau_2 \leq T | \tau_1 > t_1) = 1 - e^{\lambda(T-t_1)} \left(1 + \lambda(T - t_1)\right).$$

As $(T-t_1) \to 0$, the above expression (and hence the benefit of investing cash flows) converges to zero at a rate proportional to $(T-t_1)^2$, whereas the benefit of cash flow diversion is proportional to $(T-t_1)$, implying the agent will eventually prefer diversion as long as W is bounded.

Lemma 4.1 suggests several possibilities. Either, the optimal contract without communication involves positive probability of cash flow diversion, or, it involves a more complex termination rule in which the principal randomizes over the time at which termination occurs. This suggests there is a potential for value from communication.

To see why, consider the case of cash flow diversion and suppose that upon reaching $\bar{t}$, the principal asks the agent "Have you made a breakthrough yet?" To induce the agent to report truthfully, the principal promises to make a "severance" payment to the agent of $P = \phi c(T-\bar{t})$ if he answers "no" and continue funding the project until T with the same reward scheme if he answers "yes". This induces truthful reporting and does not change the overall probability of ultimate success because the agent who accepts the severance payment would have otherwise been diverting funds for private benefit. Further, it lowers the principal's expenditure by $(1-\phi)c(T-\bar{t})$ conditional on the agent reporting "no". Thus, this simple form of communication with the agent strictly improves the principal's expected payoff. From this argument, it appears there is in fact a role for communication.[10]

4.2 With Communication

Let us now introduce the possibility of communication more generally. By the Revelation Principle, it is without loss to look at contracts in which the agent truthfully reports progress if and when it occurs. We let $\hat{\tau}_1$ denote the time at which the agent self-reports his first breakthrough. To induce truthful reporting requires two sets of constraints, which we refer to as the *no-false-progress* (NFP) and *no-hidden-progress* (NHP). Let $U_1(t_1)$ denote the agent's continuation utility in the first stage of the project and $U_2(\tau_1, t_2)$ denote his continuation utility after truthfully reporting a breakthough at τ_1 and after a period of length t_2 has elapsed in second stage since the report.

To get a flavor for the implications of the truth-telling constraints, we note that the solution to the setting with tangible progress (characterized in Section 3) will *not* induce truthful reporting. Specifically, the agent would falsely report progress at T_1^* because this would raise his payoff by at least $\phi c\mu$ (in violation of NFP).

The class of contracts under consideration consists of a first stage deadline, T_1, a schedule of second-stage deadlines conditional on the agent's report, $T_2(\hat{\tau}_1)$, a severance payment if the

[10] We conjecture, though it remains to be shown, that randomized termination rules can also be improved upon with communication.

first-stage deadline is reached without a reported breakthrough, P, and a reward schedule $W(\hat{\tau}_1, \tau_2)$. Making a severance payment at the second deadline is not optimal because it undermines incentives for investment and truthful reporting. We conjecture (and intend to prove) that this class is without loss of generality with respect to the set of feasible payoffs.

Diverting cash flows is an inefficient form of compensation. That is, the principal can compensate the agent more efficiently through direct payments. Hence, it is again without loss to focus on contracts in which the agent does not divert cash flows. The contract must satisfy the analogous conditions to those in the case of tangible progress. The main difference is that at the deadline T_1, the agent receives the severance payment, resulting in the boundary condition

$$U_1(T_1) = P. \tag{BC1}$$

In addition and as mentioned earlier, the contract must satisfy truth-telling constraints.

No False Progress

Suppose the agent falsely reports a stage-one breakthrough at some $t < T_1$, and let $\tilde{U}_2(t, t_2)$ be his *off-path* value function for falsely reporting progress at t and has yet to make a breakthrough after an additional period of length t_2. The *no-false-progress* constraint is written

$$U_1(\tau_1) \geq \tilde{U}_2(\tau_1, 0), \quad \forall \tau_1 \in [0, T_1]. \tag{NFP}$$

That is, the agent must prefer not to falsely report a breakthrough at any point in the first phase of the project. A few key observations are as follows. First, combining (IC_1) and (NFP) yields

$$U_2(\tau_1, 0) \geq \phi c\mu + \tilde{U}_2(\tau_1, 0).$$

Next, note that one option available to the agent after falsely reporting progress is to divert funds for the remainder of the contract, so

$$\tilde{U}_2(\tau_1, 0) \geq \phi c T_2(\tau_1). \tag{13}$$

Combining this with the previous inequality yields

$$U_2(\tau_1, 0) \geq \phi c(\mu + T_2(\tau_1)). \tag{14}$$

This says that the agent receives information rents in the second stage of the project; unlike in the observable-progress setting, IC_2 does not always bind.

The agent's off-path value function satisfies the HJB equation

$$\frac{d\tilde{U}_2(\tau_1, t)}{dt} = -\phi c - \max_{a_t} \; a_t \left(\lambda \left(U_2(\tau_1, t) - \tilde{U}_2(\tau_1, t)\right) - \phi c\right). \tag{15}$$

Lemma 4.2 (Off-Path Behavior)**.** *If (14) binds, then*

$$\tilde{U}_2(\tau_1, t_2) = \phi c(T_2(\tau_1) - t_2). \tag{16}$$

That is, $a_{t_2} = 0$ *for all* $t_2 \in [0, T_2(\tau_1))$.

Because there is no reason to pay the agent more than necessary, we conjecture that (14) will bind. Then, Lemma 4.2 says that following a false report, the optimal strategy for the agent will be to divert funds for the rest of the contract, implying that (NFP) can be recast as

$$U_1(\tau_1) = \phi c T_2(\tau_1), \quad \forall \tau_1 \in [0, T_1]. \tag{17}$$

Moreover, using a similar logic, we conjecture (IC_1) will bind, yielding

$$U_1(t_1) = P + \phi c(T_1 - t_1), \quad \forall t_1 \in [0, T_1].$$

Combining this with the previous equality gives

$$T_2(\tau_1) = P/\phi c + T_1 - \tau_1. \tag{18}$$

That is, putting additional time on the second clock requires a larger severance payment in order to keep the agent from falsely reporting progress and then diverting funds until the second deadline.

No Hidden Progress

The *no-hidden-progress* constraint says that the agent must always prefer reporting a first-stage breakthrough to hiding it for any length of time while diverting funds

$$U_2(\tau_1, 0) \geq \phi c(\hat{\tau}_1 - \tau_1) + U_2(\hat{\tau}_1, 0), \quad \forall \tau_1 \in [0, T_1), \text{ and } \hat{\tau}_1 \in (\tau_1, T_1]. \tag{NHP}$$

Note that if (IC1) and (14) bind, then we have

$$U_2(\tau_1, 0) = P + \phi c(\mu + T_1 - \tau_1). \tag{19}$$

This automatically satisfies (NHP) – so the no-hidden-progress constraint will not be active provided that (IC1) binds and (14) holds with equality.

Preliminary Analysis of the Optimal Contract with Communication

The principal's contract design problem is to maximize her expected payoff by choice of: T_1, $T_2(\tau_1)$, P, and $W(\tau_1, \tau_2)$ subject to the constraints identified above. As noted above (IC_1) and (NFP) imply that (IC_2) must be slack at least some of the time, hence we will relax this constraint and then check that it is satisfied in the relaxed problem.

We now proceed by assuming (IC_1), (NFP), and (14) bind with equality, in which case, as shown above, (NHP) is redundant. At the beginning of the second stage the principal's value function is

$$
\begin{aligned}
V_2(\tau_1, 0) &= \int_0^{T_2(\tau_1)} \lambda(b - c\mu - W(\tau_1, \tau_2))e^{-\lambda\tau_2}\, d\tau_2 \\
&= (b - c\mu)\left(1 - e^{-\lambda T_2(\tau_1)}\right) - U_2(\tau_1, 0) \\
&= b - c\mu)\left(1 - e^{-\lambda T_2(\tau_1)}\right) - \phi c(\mu + T_2(\tau_1)),
\end{aligned}
$$

where the last line follows from (14).

Similarly, at the beginning of the first stage the principal's value function satisfies

$$
V_1(0) = \int_0^{T_1} \lambda(V_2(\tau_1, 0) - c\mu)e^{-\lambda\tau_1)}\, d\tau_1 - Pe^{-\lambda T_1)}.
$$

Substituting for $V_2(\tau_1, 0)$ from the previous expression and for $T_2(\tau_1)$ from (18), and performing the integration yields

$$
V_1(0) = (b - 2c\mu)\left(1 - e^{-\lambda T1}\right) - \lambda T_1(b - c\mu)e^{-\lambda(P+T_1)} - (P + \phi c T_1), \tag{20}
$$

which must be maximized with respect to T_1 and P. Note that W does not appear in the problem. Hence, we must also find a reward schedule that implements the solution; i.e., one that satisfies (IC_2) and binding (14)

$$
u_2(\tau_1, 0) = \int_0^{T_2(\tau_1)} \lambda W(\tau_1, \tau_2)e^{-\lambda\tau_2}\, d\tau_2 = \phi c(\mu + T_2(\tau_1)). \tag{21}
$$

Lemma 4.3 (Reward Schedule). *Suppose that (IC_1), (NFP), and (14) bind with equality. Then, the following reward schedule implements any T_1 and $T_2(\tau_1) = P/\phi c + T_1 - \tau_1$*

$$
W(\tau_1, \tau_2) = \phi c\left(\frac{\mu^2}{T_2(\tau_1)}e^{\lambda\tau_2} + \mu + T_2(\tau_1) - \tau_2\right). \tag{22}
$$

Proposition 4.4 (Optimal Two-Stage contract with Self-Reported Progress)**.** *Suppose that (IC_1), (NFP), and (14) bind with equality. A necessary condition for a strictly positive first period deadline is*

$$b - 2c\mu > \phi c\mu e. \tag{23}$$

If the optimal first period deadline is strictly positive then it is strictly bigger than μ and the optimal severance payment is given by

$$P^{*} = -\mu \ln\left(\frac{b - 2c\mu}{b - c\mu}\right). \tag{24}$$

The above analysis suggests that a positive severance payment will generally be part of the solution when communication is possible. Providing the agent with incentives not to falsely report progress via a severance payment is, however, costly to the principal. For this reason, it is probable that the optimal schedule of second-stage deadlines will be relatively short as compared with the first-stage deadline. In other words, once the agent reports the first breakthrough, he will not have much time before he must deliver the final project. This would contrast sharply with the case of observable progress where the first deadline is always shorter than the second (i.e., Proposition 3.5).

Except for a couple of technical proofs, the analysis of the setting with observable progress explored in Sections 2 and 3 is largely complete. The analysis of unobservable progress presented in Sections 4.1 and 4.2 is somewhat preliminary. In particular, we aim to answer the following set of questions.

1. *Intangible Progress.* Thus far, we have restricted the class of contracts we consider. Lemmas 2.1 as well as some other preliminary analysis suggest that a general optimal contract for the principal belongs to this class, but this needs to be rigorously established. Further, we need to verify the active constraints and derive both necessary and sufficient conditions for a strictly positive first period deadline and severance payment.

   Without communication, we have yet to determine the optimal contract. This is a critical step in understanding the value of communication and may help shed light on the design of real-world incentive schemes as well as the optimal schedule for reporting when doing so is costly (see below).

2. *Tangible vs Intangible.* We have conjectured that the second deadline will be short when progress is intangible relative to when it is tangible, but this (of course) must be formally demonstrated. More generally, we are interested in how the nature of progress affects the incentive scheme. Presumably, the agent will extract additional rents when

the progress is intangible, but the extent to which total surplus is affected remains an open question.

3. *Value of Communication.* An interesting set of questions involves comparing the optimal contract without reports to the optimal one with self-reported progress. Specifically, how much does the principal gain by inducing the agent to truthfully report his progress and in which environments does this matter the most?

5 Discussion and Extensions

After developing a deeper understanding on the role of progress in two-stage projects and the value of communication, there are a number of promising avenues to explore. A non-exhaustive list of these potential directions is discussed below.

N-Stage Projects. An obvious and important extension is to consider projects with more than two stages. For example, in the case of observable progress it would be edifying to know whether the result given in Proposition 3.5 that the agent is granted more time to complete each successive phase of the project is robust.

With intangible progress, each stage (save the last) will require a severance payment. Taking the limit as $N \to \infty$ can answer another interesting set of questions.[11] As N becomes arbitrarily large, the importance of an individual stage becomes less relevant. Nevertheless, progress will remain a useful monitoring device for the principal and we expect the tangibility of progress to continue to play an important role in determining the optimal incentive scheme. We are curious to see how the respective limits compare to the (stationary) Brownian cash flow diversion models (e.g., DeMarzo and Sannikov, 2006).

Asymmetric Stages. In reality, different stages are unlikely to be symmetric. In some environments, progress may be more rapid but more expensive in later stages ($\lambda_2 > \lambda_1$ and $c_2 > c_1$). In others, it may be easier for the agent to divert funds in early stages ($\phi_1 > \phi_2$). More generally, the parameters λ, c, and ϕ, should be allowed to vary across stages of the project. Performing comparative statics with respect to these parameters will provide valuable insights regarding the design of optimal incentive schemes across a broader class of applications. Also, when progress is intangible and stages are asymmetric (e.g., $\phi_1 > \phi_2$), (NHP) will likely bind, so that the agent must be given incentives not to hide an early breakthrough.

Projects in Teams. Consider the two-stage project with tangible progress and now suppose the principal has the ability to split the individual stages across two different agents. Is this a

[11] To maintain positive and finite project value, $\lambda \to \infty$ at a rate proportional to N.

preferable arrangement for the principal to using a single agent for both stages (as in Section 3). One observation is that even when using two agents, the principal will give less time to the first agent and more time to the second agent conditional on a breakthrough in the first stage. This is because a success in the first stage is worth less to the principal than the ultimate success. More generally, we plan to explore under what conditions is it a good idea to break up different stages across different agents and under what conditions should they be combined.

Heterogeneous Agents. Agents might differ in a number of dimensions not readily observed by the principal. For instance, some might be more talented (higher λ) and/or more honest (lower ϕ). Consider, for example, a setting with observable progress where there are two possible types of agent, $\lambda_H > \lambda_L$. Here, the frequency of successes can be used by the principal as a screening instrument in order to help separate types and mitigate payment of information rents. In a setting of unobservable progress it would be especially interesting to study a situation where agents receive intrinsic private benefits from completion of the project, which they also privately observe. Calibrating the optimal incentive scheme to account for such unobserved heterogeneity is a challenging but potentially fruitful avenue for study.

Monitoring and Verification. How frequently should the principal monitor the progress of the agent? How much should a principal invest in a technology that can verify progress? How should it be optimally used? Should the principal use experts in order to verify reports from the agent. Should the experts be "inside" the organizations or should they be independent contractors? An obvious advantage of such technology is to relax the no-hidden-progress and no-false-progress constraints. The trade-off is that monitoring and verification is costly: indeed the principal may still have to provide incentives to the verifying experts to induce them to exert effort and report truthfully. Further exploration of these questions may provide insights as to what type of projects should be conducted inside the organization and which ones should be outsourced.

Costly Reporting. Our model assumes that there are no costs associated in sending messages from the agent to the principal. But for many relevant applications, communicating with the principal (e.g., submitting progress reports) may be a costly and time consuming exercise. In this case, the principal will also need to decide on a *schedule* for progress reports and in doing so, trade off the ultimate success of the project for more information about its current progress.

Limited Commitment. Thus far, we have allowed the principal to commit to a long-term contract that can be conditioned on the entire past performance (or self reports). In some applications, especially when breakthroughs occur infrequently (e.g., basic research), it is difficult to justify this assumption. It might be more reasonable to think that the principal has only a limited form of commitment. For example, the principal can commit to a length of time for which she agrees to fund the project (e.g., a three year grant), but cannot commit to whether (or how much) additional funding she will provide at the end of the grant. This is an important consideration that we believe can be addressed within our framework.

6 Conclusion

We study the optimal provision of incentives in a dynamic multistage agency setting. When breakthroughs are tangible, they serve as a useful monitoring device, and the principal controls the agents incentives through a sequence of increasing deadlines and a reward scheme that decreases over time. When breakthroughs are intangible, communication with the agent (i.e., self-reported progress) can have value to the principal. However, the principal must provide incentives for truthful reports, which requires the use of payments upon project termination, which we have interpreted as severance pay.

References

Baker, G. (2000): "The Use of Performance Measures in Incentive Contracting," *American Economic Review*, 90, 415–420.

Baker, G. P., R. Gibbons, and K. J. Murphy (1994): "Subjective Performance Measures in Optimal Incentive Contracts," *The Quarterly Journal of Economics*, 109, 1125–1156.

Battaglini, M. (2005): "Long-term contracting with markovian consumers," *American Economic Review*, 95, 637–658.

Bergemann, D. and J. Valimaki (2010): "The Dynamic Pivot Mechanism," *Econometrica*, 78, 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.

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. Tchistyi (2014): "Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives," .

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.

Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012): "Dynamic CEO Compensation," *Journal of Finance*, 67, 1603–1647.

Eso, P. and B. Szentes (2007): "Optimal information disclosure in auctions and the handicap auction," *Review of Economic Studies*, 74, 705–731.

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.

Green, E. J. (1987): "Lending and the Smoothing of Uninsurable Income," in *Minnesota Studies in Macroeconomic, Volume I*, ed. by E. C. P. Wallace and Neil, University of

Minnesota Press.

Halac, M., N. Kartik, and Q. Liu (2013): "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.*

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.

Levin, J. (2003): "Relational Incentive Contracts," *American Economic Review*, 93, 835–857.

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.

Macleod, W. B. (2003): "Optimal Contracting with Subjective Evaluation," *American Economic Review*, 93, 216–240.

Maestri, L. (2012): "Bonus Payments versus Efficiency Wages in the Repeated Principal-Agent Model with Subjective Evaluations," *American Economic Journal: Microeconomics*, 4, 34–56.

Manso, G. (2011): "Motivating Innovation," *Journal of Finance*, 66, 1823–1860.

Mason, R. and J. Välimäki (2011): "Dynamic moral hazard and stopping," *... of Exeter* *and ....*

Pavan, A., I. Segal, and J. Toikka (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica (forthcoming).*

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.

Sannikov, Y. (2008): "A continuous-time version of the principal-agent problem," *Review of Economic Studies*, 75, 957–984.

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.*

Williams, N. (2011): "Persistent Private Information," *Econometrica*, 79, 1233–1275.

Zhang, Y. (2009): "Dynamic contracting with persistent shocks," *Journal of Economic Theory*, 144, 635–675.

Appendix

Several of the proofs below are incomplete and/or contain only a sketch of the argument.

Proof of Lemma 2.1. Cash flow diversion is an inefficient form of compensation. Hence, any contract in which the $a_s < 1$ over some interval can be improved upon by increasing the promised utility for success such that the agent will be induced to and/or direct monetary payments. Since both principal and agent are risk nuetral and do not discount cash flows, it is without loss to make all payments at the end of the project. Further, any payment upon termination distorts the agents incentives. □

Proof of Proposition 2.2. Noting that $V(u)$ is concave in u we can differentiate and get that $u_0^* = \frac{\phi c \log\left(\frac{\lambda b - c}{\phi c}\right)}{\lambda}$. Since $U_t' = -\phi c$, this implies that the time of project termination is $T^* = \frac{1}{\lambda} \log\left(\frac{\lambda b - c}{\phi c}\right)$. That w is of the stated form follows from noting that (4) binds and substituting in $U_t = \phi c(T^* - t)$. □

Proof of Lemma 3.1. Twice differentiating $V_1^c(u)$ gives

$$\frac{d^2}{du^2} V_1^c(u) = \frac{\lambda e^{-\frac{\lambda u}{c\phi} - 1} \left(c\lambda(u - (e-2)b\phi) - b\lambda^2 u + 2(e-1)c^2\phi\right)}{c^3\phi^3} \tag{25}$$

The term inside the parentheses is strictly decreasing in u given (1). Therefore, the function is concave for all $u \geq 0$ if and only if $\frac{d}{du^2} V_1^c(0) \leq 0$. Evaluating (25) at $u = 0$, we get that

$$sign\left(\frac{d}{du^2} V_1^c(0)\right) \leq 0 \iff sign\left(\frac{\lambda(2(e-1)c - (e-2)b\lambda)}{ec^2\phi^2}\right) \leq 0.$$

The second inequality is equivalent to (12). □

Proof of Proposition 3.2. That the optimal contract can be implemented with the two deadlines follows from $U_t' = -\phi c$, that continuation utility jumps up by $\phi c \mu$ conditional on a breakthrough and that the project must be terminated when $U_t = 0$. That w is of the stated form follows from similar arguments to those given in the proof of Proposition 2.2. □

Proof of Proposition 3.3. By Lemma 3.1, when (12) does not hold, there exists a $\bar{u} > 0$ such that V_1^c is convex on $[0, \bar{u})$. Standard arguments (e.g., DeMarzo and Fishman (2007)) imply that under these conditions it is optimal to randomize. □

Proof of Proposition 3.4. To be completed. □

Proof of Proposition 3.5. We break the proof into two separate cases depending on whether (12) holds.

1. Suppose that (12) holds. From Proposition 3.2, we have that

$$\mathbb{E}[T_2^*|\tau_1 \leq T_1^*] = \mu + T_1^* - \mathbb{E}[\tau_1|\tau_1 \leq T_1^*]$$

   Noting that $\mathbb{E}[\tau_1|\tau_1 \leq T_1^*] < \mathbb{E}[\tau_1] = \mu$, yields, $\mathbb{E}[T_2^*|\tau_1 \leq T_1^*] > T_1^*$ and completes the argument.

2. Suppose that (12) does not hold. Let $\tilde{T}_1$ denote the random variable represnting the first period deadline. From Proposition 3.3, we have that

$$\begin{aligned}
\mathbb{E}[T_2^*|\tau_1 \leq \tilde{T}_1] &= \mu + \frac{1}{\sigma} + \mathbb{E}[T_1^{**} - \tau_1)^+|\tau_1 \leq \tilde{T}_1] \\
&\geq \mu + \frac{1}{\sigma} + \mathbb{E}[T_1^{**} - \tau_1|\tau_1 \leq \tilde{T}_1] \\
&> \mu + \frac{1}{\sigma} + T_1^{**} - \mathbb{E}[\tau_1|\tau_1 \leq \tilde{T}_1] \\
&> \frac{1}{\sigma} + T_1^{**} = \mathbb{E}[\tilde{T}_1].
\end{aligned}$$

□

Proof of Lemma 4.1. A simple calculation yields that

$$\Pr(\tau_2 \leq T|\tau_1 > t_1) = 1 - e^{\lambda(T-t_1)}\left(1 + \lambda(T - t_1)\right).$$

As $(T - t_1) \to 0$, the above expression (and hence the benefit of investing cash flows) converges to zero at a rate proportional to $(T - t_1)^2$, whereas the benefit of cash flow diversion is proportional to $(T - t_1)$, implying the agent will prefer diversion as long as W is bounded. □

Proof of Lemma 4.2. Because $U_2(\tau_1, 0) = \phi c(\mu + T_2(\tau_1))$ and $U_2(\tau_1, 0 \geq \phi c\mu + \tilde{U}_2(\tau_1, 0)$, (16) is actually the only feasible solution. We verify that it indeed satisfies (15. Note that $u_2(\tau_1, 0) = \phi c(\mu + T_2(\tau_1))$ and $\frac{dU_2(\tau_1,t_2)}{dt_2} \leq -\phi c$ together imply $U_2(\tau_1, t_2) \leq \phi c(\mu + T_2(\tau_1) - t_2)$. Thus, for all $t_2 \in [0, T_2(\tau_1))$

$$\lambda\left(U_2(\tau_1, t_2) - \tilde{U}_2(\tau_1, t_2)\right) - \phi c \leq \lambda\left(\phi c(\mu + T_2(\tau_1) - t_2) - \phi c(T_2(\tau_1) - t_2)\right) - \phi c = 0.$$

Hence, $a(t_2) = 0$ for all $t_2 \in [0, T_2(\tau_1)]$. Substituting this into (15) gives $\frac{d\tilde{U}(\tau_1,t_2)}{dt_2} = -\phi c$. Integration and the boundary condition $\tilde{U}_2(\tau_1, T_2(\tau_1)) = 0$ establish the claim. □

Proof of Lemma 4.3. We have

$$\begin{aligned} U_2(\tau_1, t_2) &= \int_{t_2}^{T_2(\tau_1)} \lambda W(\tau_1, \tau_2) e^{\lambda(t_2 - \tau_2)}\, d\tau_2 \\ &= \phi c \left(\frac{\mu}{T_2(\tau_1)} + 1 \right) (T_2(\tau_1) - t_2). \end{aligned}$$

This yields (21) for $t_2 = 0$ and evidently satisfies (IC_2). □

Proof of Proposition 4.4. First we show that a non-trivial optimum must involve the value of P given in (24). To see this, compute the first-order conditions

$$\frac{\partial V_1(0)}{\partial T_1} = \left((\lambda b - 2c) + (\lambda T_1 - 1)(\lambda b - c)e^{-\lambda P}\right) e^{-\lambda T_1} - \phi c = 0 \tag{26}$$

and

$$\frac{\partial V_1(0)}{\partial P} = \lambda T_1 (\lambda b - c) e^{-\lambda(P + T_1)} - \phi c = 0. \tag{27}$$

Solving these for P yields (24). Moreover, substituting(24) into (26) yields the necessary condition for a non-trivial optimum

$$H(T_1^{**}) \equiv (\lambda b - 2c)\lambda T_1^{**} e^{-\lambda T_1^{**}} - \phi c = 0.$$

Note that

$$H(0) = \lim_{T_1 \to \infty} H(T_1) = -\phi c.$$

Thus a necessary condition for the existence of a non-trivial optimum is $\max_{T_1} H(T_1) > 0$. It is straightforward to verify that $H(\cdot)$ is maximized at $T_1 = \mu$, and $H(\mu) > 0$ iff (23) holds. It remains to check the second-order condition and decide about whether low outside option for principal works. □