Experimentation, Private Observability, and the Timing of Monitoring

Alexander Rodivilov[ζ]

May 2016

Abstract

We consider a principal that must hire a financially constrained agent to execute a project of uncertain feasibility. When the principal allows the agent to hold several trials privately, the agent may not announce when the experiment results in success. Under the optimal contract, the agent's private observations are inconsequential. However, private observability of success plays a role in the optimal monitoring period. When the agent publicly observes success, the principal monitors the agent from the start of their relationship. This contrasts with the Bergemann and Hege (1998) result, where optimal monitoring occurs toward the end of the project, when the agent is raising funds in a competitive market. However, when the agent observes success privately and both parties remain patient, monitoring is most useful at the end of the relationship.

[ζ] University of Washington, arodivil@uw.edu. I am grateful to Fahad Khalil and Jacques Lawarrée for their valuable suggestions and insightful conversations. I also am thankful for the helpful comments of Quan Wen, Philip Bond, and participants at Midwest Economic Theory Meetings 2016.

Introduction

The research and development process for new drugs and medicines is complex, lengthy and expensive.[1] In the United States, a new medicine completes its journey from initial discovery to the marketplace in 12 years on average, with clinical trials alone last between five and seven years. Research and development (R&D) of a successful drug costs an average of $2.6 billion. This amount incorporates the costs of all possible failures: Among the thousands and sometimes millions of compounds that have to be screened and assessed early in the R&D process, only a few will receive approval. The overall probability of clinical success, which is simply the likelihood that a drug entering clinical testing eventually will be approved, is estimated to be less than 12%.[2]

Consider a pharmaceutical company (principal) that is attempting to develop a new drug. The firm initially is unsure about the drug's efficacy; it only understands that combining certain microelements will produce the intended results. For example, promising results from laboratory tests on animals may cause the firm to believe the drug can cure certain human diseases. To verify this, the firm employs a physician researcher (agent) who will carry out protocol in performing tests on a group of volunteers, as required by law.

Before the company can sell its drug to patients, the firm must submit for U.S. Food and Drug Administration (FDA) approval. In preparation for the approval from the FDA, a candidate drug must demonstrate its effectiveness through extensive human trials. In later studies, clinical trials enroll patients who have the condition that the medicine is designed to address. Their participation is crucial for researchers to evaluate the medicine's effectiveness and whether it requires adjustments in dosing and timing.

A series of clinical trials demonstrate whether the drug functions as expected. For each trial that does not result in desirable effects, the firm will become more pessimistic about the drug's quality and may abandon the project. If, however, the drug affects research volunteers positively, the agent can prove success by providing verifiable test results for FDA approval.

---

[1] For example, industry-wide research and development spending in the United States was estimated to cost $45 billion in 2011. See http://www.nytimes.com/2011/03/07/business/07drug.html.
[2] See Dimasi et al. (2003) for a detailed analysis.

If the company does not to observe the agent testing the drug directly, it may doubt whether the agent is exerting all efforts and resources properly. If the agent chooses to shirk his responsibilities, he simply can report that the drug was tested unsuccessfully. If the company remains unaware of this falsehood, it will adjust its beliefs about the drug's quality accordingly, becoming more pessimistic. Even if the agent discovers the drug is performing well, he might delay announcing success in favor of personal gain. The company can avoid this by introducing incentives that encourage the agent to announce success immediately. Finally, since the agent lacks his own finances, the principal can pay the agent at certain intervals, but penalizing him is not feasible.

Under certain circumstances, the principal may be able to detect that the agent is shirking his responsibilities by reviewing volunteers' testimonies. The principal also can hire a monitor to supervise the agent. The optimal contract would then be contingent on the monitor's reports. Since the relationship between the company and the agent is dynamic, a monitor might be useful during certain periods only, such as during early or later trials.

This paper studies an optimal contract that will ensure that the agent works properly during every period of the repeated clinical trials, prevent the agent from postponing announcement or hiding evidence of a successful trial and state that the financially constrained agent will not be asked to make payments to the pharmaceutical company in the event that time and resources are wasted on an unsuccessful project. We explore specific questions that arise in such an environment: How many unsuccessful iterations of the project will the principal tolerate? Can the principal benefit from hiring a monitor?[3] In a dynamic relationship, collecting information during every period is prohibitively costly, raising the question of when monitoring is most effective.

The optimal contract has to take into account four crucial features of the relationship between the principal and the agent: First, the results of each trial affects the relationship; with each failure the agent reports, the principal becomes more pessimistic. Second, during each period, the agent chooses privately whether to exert the effort necessary for the project to

[3] Contract theory literature suggests that optimal incentive schemes should use all available information related to the agent's performance, and under certain circumstances, the agency relationship creates a demand for monitoring. See the literature on the "informativeness principle," such as Holmstrom (1979) and Shavell (1979).

succeed. Third, the agent is financially constrained, so the principal cannot sell the project to the agent. Finally, the agent observes successful project implementation privately. As a result, the payment structure must ensure that the agent neither postpones nor hides the announcement of a successful implementation.

Our model for optimal contracting features a project with uncertain feasibility; it could yield a positive return or nothing at all. To illustrate this, we use a simple two-armed bandit model:[4] The agent can "pull the risky arm" by exerting all his efforts toward implementing the project, or he can "pull the safe arm" and shirk his responsibilities. While pulling the risky arm is costly, it allows the project to be implemented successfully if it is feasible. Pulling the safe arm, however, yields zero return, regardless of project quality. If the project fails, both the principal and the agent update their beliefs regarding the project's quality. As time without success increases, the principal becomes more pessimistic and may even terminate the relationship with the agent and discard the project.

We will demonstrate that without a monitor, a high-powered contract is optimal, meaning both the principal and the agent benefit only if the project is successful. If the agent postpones or hides a successful trial, the value of his reward decreases, although the nominal value of the reward increases to account for rising pessimism. When monitoring is included, private observability of success factors into the optimal time period for monitoring. If success is publicly observed, the principal should monitor the agent at the beginning of their relationship. However, if the agent observes success privately, the optimal timing for monitoring is affected by patience. If the contract includes a high discount factor, monitoring should be performed during the final trial.

To better understand the reasoning behind these results, consider the first-best scenario, in which effort is observed publicly. Since each trial is costly, and the principal becomes more pessimistic when success is not announced, the first-best solution is characterized by a stopping rule. The agent is allowed to attempt to successfully implement the project for several periods only. If the agent does not achieve success despite exerting all his efforts, both parties become more pessimistic about the project's quality. As a result, the expected value of the project

[4] See Keller et al. (2005) for more details.

diminishes every subsequent period. If the agent reports success, the project is released to market immediately, and the principal benefits. On the other hand, if the agent never succeeds, the project is abandoned, and the principal gains nothing.

When the agent chooses his effort level privately, he receives a strictly positive rent. The agent could shirk his duties and report that the project failed. The principal can motivate the agent to exert effort by paying a higher reward for success and a lower one for failure. The gap between these payments must be wide enough for the agent to believe it is in his best interest to exert all his efforts after taking into account current beliefs of the project's quality and probability of success. If the agent and principal share the same beliefs about the project's quality, a standard moral hazard problem takes place within each trial period. To minimize risk, the principal ideally would sell the project to the agent; however, this is not feasible because the agent is financially constrained. Instead, the agent receives a positive rent. Since the principal benefits only if the project is released to market, it gains advantage by awarding little to the agent if failure, and, given the limited liability constraint, the agent is paid nothing if he fails overall.

Moreover, if the agent and the principal do not hold common beliefs, the former receives additional reward. If the agent deviates from project goals at one period, his chance to succeed in all future periods remains. Although the agent will not receive anything if he deviates from his duties during a particular period, he becomes relatively more optimistic than the principal for all future periods. That means a deviation at one period carries into all future periods by creating asymmetric beliefs among parties. In some sense, the agent is relatively more patient than the principal in all periods except the last. During this final period, the agent cannot benefit from shirking his responsibilities, since he will not gain from this, and his rent is contingent on the combination of moral hazard and limited liability only. Because of the positive rent the agent receives, the project could be terminated inefficiently early.

The agent's promised reward consists takes moral hazard and the learning rent into account. The first component is always increasing, since the agent becomes more pessimistic as time proceeds without success, and his motivation to exert himself becomes costlier. The second component, however, is non-monotonic and depends on the discount factor. Consider a case in which both the agent and the principal are patient (the discount factor equals one). Under these circumstances, the principal can wait for success indefinitely. Without loss of generality, the principal can offer a contract with constant nominal reward and a deterministic deadline that will ensure the agent will exert effort in every period. Since the moral hazard rent is increasing strictly and the nominal reward is constant, the learning rent decreases strictly. If the principal is patient, the agent benefits less from deviating from project goals since the fixed deadline gives him a smaller horizon to benefit from asymmetric beliefs.

However, if parties to a contract are impatient (the discount factor is less than one), the learning rent becomes non-monotonic. As in the previous case, this means that except for the final period, if the agent shirks and does not incur the cost-of-effort, he will have an additional attempt to successfully implement the project and receive a reward. This weakens the agent's incentives to work during each period. The principal, however, benefits if the project is successfully implemented only, and it cannot wait indefinitely, as later success means discounting. This allows the agent to be relatively more patient than the principal. By this logic, the learning rent is increasing. However, the later the agent deviates, the fewer periods remain to exploit the difference in beliefs and, as a result, the learning rent eventually decreases before vanishing completely during the final trial. Since the agent has more incentives to deviate at the beginning of the relationship, the optimal contract makes the discounted reward strictly decreasing. This would prevent private observability of success from exacerbating the contracting environment.

In some settings, the principal can observe and verify experimental success easily. For example, the agent may have difficulty hiding a revolutionary drug's success from the public.[5] However, success might be much more difficult to ascertain when information gathering does not involve extreme outcomes.[6] In this scenario considered, private observability of success does not worsen the problem; however, we will demonstrate that it becomes a crucial factor in defining the optimal timing for monitoring.

[5] During the drug explosion in the 1950s, some drugs were found to serve purposes for which they were not originally intended. For example, Thorazine was developed as an antihistamine before it became the first antipsychotic.

[6] FDA provides companies with guidance during the phases of human clinical trials. Even so, the number of people in a clinical trial of a new drug is usually small in comparison to the number of people who may take the drug if it reaches the market. This can make detecting rare effects difficult.

Consider first the case in which success is observed commonly. The principal would benefit from monitoring if during the monitoring period, it does not have to reward the agent in case he succeeds—we refer to this as the *static effect*. Recall that except for the final testing period, the agent can shirk his duties and attempt to implement the project during later trials. As a result, his incentives to work during each period depend not only on the payment determined by the contract and success of that particular trial, but also on the payments for all subsequent periods.[7] Thus, the *dynamic effect* emerges: The principal can diminish all of the agent's rewards in all periods before he is monitored. While static effect dominates, and monitoring is optimal during the first period, this would not be that case if the agent was raising funds in a competitive market.[8] In this case, the agent's reward for success is non-monotonic and can either increase or decrease for earlier periods. This optimal payment scheme disrupts the dynamic agency effect and makes monitoring most optimal during the later trials.

When the agent observes success privately, the principal remains paying a positive reward for success during every monitoring period except for the last. For example, if a reward in the current period is smaller than the discounted value of a reward in one of the concurrent periods, the agent optimally will postpone an announcement of success. Given that the optimal contract without monitoring includes decreasing discounted reward value, the principal can decrease a reward of one period up to the discounted reward of the following period only. In the final period, however, the principal optimally can promise paying nothing. As discount factor increases, the benefits of monitoring all periods except the last become smaller, and, as a result, optimal monitoring takes place at the end.

We argue that these findings are not only theoretical, but also empirically significant. Business literature[9] on venture capital for innovation emphasizes the importance of relationship financing and monitoring. For example, Gorman and Sahlman (1989) found in their survey that venture capitalists visit their companies frequently and devote significant amount of time to participating in decision-making. According to the survey, a typical venture capitalist spends

---

[7] Halac et al. (2016) discuss this dynamic agency effect with respect to moral hazard. See Mason and Valimaki (2011), Bhaskar (2013), Horner and Samuleson (2013) and Kwon (2013) for dynamic agency effects in similar settings.
[8] See Bergemann and Hege (1998) for an example of such an environment.
[9] See Gorman and Sahlman (1989), Sahlman (1990) and references therein.

about 80 hours a year onsite with a company whose board he or she serves. In addition, he or she holds frequent telephone conversations and works on the company's behalf to attract new investors and management candidates.

With respect to optimal monitoring time, Lerner (1995) demonstrated that venture capitalists' representation on the board of directors increases over certain periods, especially during chief executive officer turnover. While a venture capitalist may have reasons to participate in the firm's recruitment of a management team and strategic planning, this paper, does not examine such motives and considers monitoring that serves the purpose of eliminating the moral hazard problem only.

The rest of the paper is organized as follows: Section 1 explores related literature; Section 2 explains the model and the contract space with payoffs, and provides a solution for the first-best benchmark; Section 3 provides a description of the optimal contract with moral hazard and limited liability; Section 4 extends results for the case in which the principal can hire a monitor with private as well as common observability of success; and Section 5 concludes the paper.

1. Related Literature

The only researchers we are aware of that studied optimal monitoring time set in optimal contracts for experimentation are Bergemann and Hege (1998), who considered the provision of venture capital in a dynamic agency model. The optimal share contract operates on the provision that if the entrepreneur succeeds, he conveys a part of the project to the investor. The authors demonstrated that long-term contracts are optimal, the project is terminated inefficiently early and the expected share of the entrepreneur decreases. Here, monitoring is optimal toward the end of the project. The authors consider an entrepreneur who is raising funds in a competitive market, whereas in our research, the principal is making the high-stakes offer. We will show that this difference is crucial for the timing of monitoring. In Bergemann and Hege's environment, the share for the earlier periods can rise or fall but has at most one extremum in time, and then it is a maximum, whereas in our paper, a nominal reward for success is always increasing. Our paper complements Bergemann and Hege's (1998) result in that it highlights the pivotal roles of market structure and private observability of project success on the optimal timing of monitoring.

When studying monitoring, Bonatti and Horner (2011) considered a team of agents who work together on a project of uncertain feasibility. Positive probability leads to a positive reward for each agent, and this depends on the agents' combined efforts and the project being good. However, success will never be observed, regardless of effort level, if the project is bad. As time proceeds without success, agents become more pessimistic in the quality of the project. Crucially, agents only observe their own effort levels and only form beliefs regarding the effort of other teammates. Because the relationship is dynamic, a team member can postpone exerting effort since there always is a chance that other agents' efforts will suffice for early success, which all agents benefit from, regardless of effort. An intriguing result is that monitoring (knowing effort level of all the other agents) does not necessarily improve the outcome. A trade-off arises: Though observing other teammates' effort choices prevents unreasonable pessimism, it might lead to early high-level effort faster learning and later low-level effort. This demonstrates that setting binding deadlines can be more efficient than monitoring. However, in a setting with one principal and one agent, only the principal would benefit from observing effort choices.

Literature on experimentation includes Keller et al.'s (2005) study, which analyzed a strategic experiment that used a two-armed bandit model with a risky arm that might yield payoffs after exponentially distributed at random times and a safe arm that offers a safe payoff. In particular, we model projects in which success follows geometrically (a discrete version of exponentially) distributed at random times.

Bergemann and Hege's 2005 study built on their 1998 study, with one crucially distinct feature—the time horizon is infinite and the funding decision is renegotiated period by period. Two methods of funding are considered. First, venture capitalist, a form of relationship financing, actively involves decision-making within portfolio companies. The capitalist and the agent receive symmetric information, so if the agent diverts funds, his future option will depend on current beliefs since the capitalist understands the failure was uninformative. Second, angel investors provide funds and do not observe the decision made by the entrepreneur directly in a form of arm's-length relationship. A simple trade-off exists: Relationship financing does not require paying information rent, whereas arm's-length funding implicitly commits the investor to a shorter funding horizon. The authors concluded that arm's-length relationships are preferred always.

Halac, Kartik, and Liu (2016) considered the problems of creating a contract for a project of uncertain feasibility with adverse selection and moral hazard. The agent's probability of implementing the task was conditional on the project's ability to succeed and be profitable. The optimal contract involves paying the agent initially and penalizing him progressively if success has not been observed. Such penalty contracts are appropriate since they mitigate the agent's incentives to hide success. Our research differs from Halac, Kartik, and Liu's (2016) in that we assume the probability of success is known, but the agent is financially constrained. The agent thus cannot be penalized for failure; instead, we consider how bonus contracts and optimal monitoring might discourage the agent from hiding success.

Bonatti and Horner (2012) studied a model in which the state of the world (or the quality of the project) is a worker's skill that is revealed through output, and wage is based on the expected output, or assessed ability. This model features a continuous effort level. However, the authors do not assume limited liability and allow punishment for achieving certain deadlines.

Manso (2011) applied a similar model to study a manager who must be incentivized to perform an innovative task. The optimal incentive scheme exhibits significant tolerance and even reward both for early failure and long-term success. Moreover, commitment to a long-term compensation plan is essential to motivating innovation. An important difference is that in Manso's paper, the relationship lasts exogenously for two periods, whereas in our paper, the principal chooses how many failures to tolerate.

Gerardi and Maestri (2012) analyzed how an agent can be incentivized to acquire and truthfully report information over time. The authors assumed that the principal observes the state of nature with some time lag and, as a result, the optimal contract can reward or punish the agent after comparing the agent's report with the revealed state. An optimal contract specifies the duration of the testing phase and rewards based on the credibility of reports. The authors consider extreme (the agent is rewarded only if his report matches the state), evidence-based (the principal pays the agent only when the relationship ends and the true state is observed) contracts without loss of generality. Finally, Mason and Valimaki (2015) considered in a similar model an infinitely lasting relationship with continuous effort and solve for a stationary wage, but they did not study monitoring.

Most of these papers assumed project success would be publicly observed.[10] We, in contrast, assume that the agent could observe success privately. We also assume that information is hard—that is, the agent can either postpone announcing a successful implementation or hide it completely by destroying evidence.[11]

2. The Model

The project

A principal owns a valuable idea that could result in a lucrative project, but he lacks the decisive skills needed to implement it. He is considering hiring a financially constrained agent with relevant competency. Both parties initially are uncertain about the project's quality; that is, the common prior on the project being "good" is $\beta_0 \in (0,1)$[12]. If the project is good, then it can be implemented successfully with a known positive probability, in which case it will yield a fixed return of $V > 0$, which is commonly known at the beginning of the relationship. To implement the good project, the agent must exert effort that is assumed to be subject to a binary choice: $e \in \{0,1\}$. If the project is bad, then it will yield zero, regardless of effort.[13] Exerting effort costs $c > 0$ per period.

Probability of success		*Effort* $e = 1$	*Effort* $e = 0$
Quality of the project	*Good*	$\lambda > 0$	0
	Bad	0	0

Table 1. Probability of success.

[10] Halac et al. (2016) discuss the robustness of an optimal contract to project success being privately observed by the agent.
[11] Gerardi and Maestri (2012) consider a case of soft information, i.e., when the agent can make up the evidence of success.
[12] It is important that β_0 is strictly positive and strictly less than one. Otherwise, no additional information arrives as the relationship proceeds; in this case, there is no learning regarding the quality of the project, and the problem simplifies to standard dynamic moral hazard.
[13] We refer to an implementation of the project as "success" and to lack of success as "failure."

The agent's ability, λ, which is the probability of achieving success given that the project is good, conditional on exerting effort, is common knowledge during contracting. Finally, we assume that the effort choice is not observable and that the agent can postpone announcement of or hide a successful implementation.

Learning the quality of the project

An important feature of this model is learning project quality. When the agent does not succeed, despite exerting effort, he updates[14] his beliefs regarding the quality of the project using Bayes' rule and becomes more pessimistic. Denoting by $\tilde{\beta}_t$, the updated belief of the agent that the project is good at the beginning of period t after $t-1$ failures, we present:

$$\tilde{\beta}_t = \frac{\tilde{\beta}_{t-1}(1-\lambda^{\theta})}{\tilde{\beta}_{t-1}(1-\lambda^{\theta})+(1-\tilde{\beta}_{t-1})}\text{, which simplifies to } \tilde{\beta}_t = \frac{\beta_0(1-\lambda^{\theta})^{t-1}}{\beta_0(1-\lambda^{\theta})^{t-1}+(1-\beta_0)}.$$

Since the agent chooses effort level privately, both parties do not share the same beliefs necessarily as their relationship evolves. The principal becomes more pessimistic every period the agent does not announce success. However, the agent secretly can shirk and become relatively more optimistic from that period on. Consider a hypothetical scenario in Figure 1 below:

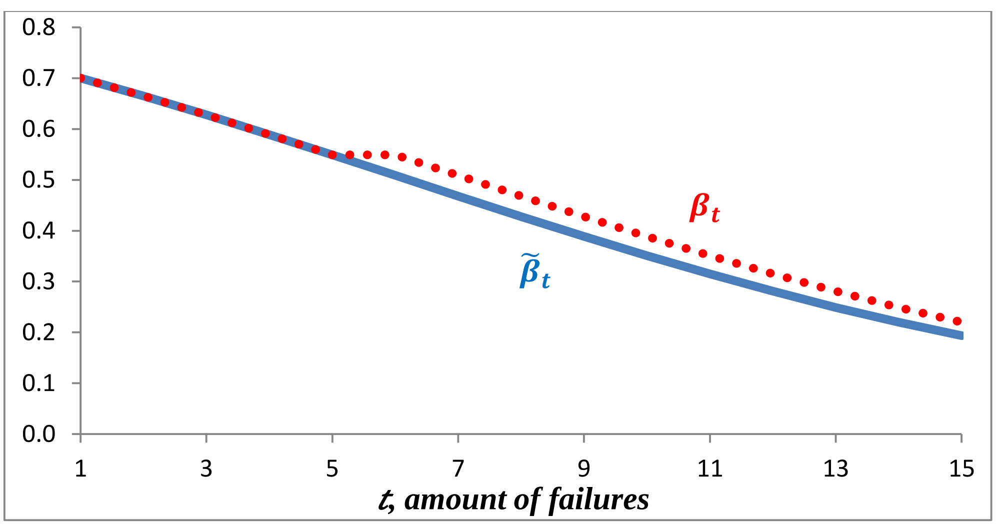

Figure 1. Learning the quality of the project with $\lambda = 0.15$ and $\beta_0 = 0.7$.

[14] A failure is more informative if beliefs are close to $\frac{1}{2}$, while beliefs change slowly when parties are relatively certain about the quality of the project. See Bergemann and Hege (1998) for more details.

Given the parameters, the bold line reflects the evolution of beliefs if the agent continues exerting effort for 15 periods. Suppose the agent secretly deviates at $t = 5$, but reports that success has not been achieved, despite exerting effort. The principal would use this report to update his beliefs and become relatively more pessimistic from period $t = 6$ on. The agent, in contrast, would understand that the reported failure was uninformative, and at the beginning of period $t = 6$ would have the same beliefs as in the previous period. Importantly, this difference in beliefs following one deviation at $t = 6$ would carry into all future periods until the relationship ends.

Contracts and payoffs

Both parties are risk neutral and share a common discount factor $\delta \in (0,1]$. An optimal contract must specify how many failures the principal will tolerate and a sequence of transfers as a function of the agent's reports,[15] which in this case is whether or not the agent succeeded. All transfers are from the principal to the agent.

The contract is given by $\varpi = (T, \{b_t\}_{t=1}^{T}, \{w_t\}_{t=1}^{T})$, where $T \in \mathcal{N}$ is the duration of the relationship, b_t is the payment to the agent in case he reports success at period $1 \leq t \leq T$ and w_t is the payment to the agent conditional on reporting failures from the beginning of the relationships up to period $1 \leq t \leq T$.

The optimal contract will have to satisfy the following incentive compatibility constraint for every period $1 \leq t \leq T$:

(IC) $\quad b_t \geq \delta b_{t+1}$ for $t = 1, \ldots, T-1$,

$\quad b_t \geq w_t$ for $t = 1, \ldots, T$,

Given the optimal contract and effort levels the agent chooses, we can specify the agent's expected utility and the principal's expected profit. The agent's expected utility from accepting contract ϖ at time zero while exerting an effort profile $\vec{e}$ and reporting each project truthfully as a failure or success is:

$$U(\varpi, \vec{e}) = (1 - \beta_0) \sum_{t=1}^{T} \delta^t (w_t - e_t c)$$

[15] Because success is observed privately, both parties do not necessarily share the same history as the relationship evolves.

$$+\beta_0 \sum_{t=1}^{T} \delta^t (\prod_{s=1}^{t-1}(1-\lambda e_s))\big((e_t(\lambda b_t - c) + (1-\lambda e_t)w_t)\big),$$

where $\vec{e} = (e_1, \dots, e_T)$ is an effort profile with $e_t \in \{0,1\}$[16] for $1 \leq t \leq T$.

First, the agent has a chance to succeed during the relationship; this occurs only if both the project is good, which is true with probability β_0, and if the agent is exerting effort. Conditional on the project being good, the relationship lasts for an arbitrary period, $t \leq T$, with probability $\prod_{s=1}^{t-1}(1-\lambda e_s)$. If the agent exerts effort at period t, his expected payoff at this period is:

$$\lambda b_t + (1-\lambda)w_t - c,$$

whereas in case he shirks, the agent receives only w_t, as defined by the contract. Second, if the project is bad, which happens with probability $1-\beta_0$, the agent never succeeds, regardless of effort profile.

The principal's expected profit from offering contract ϖ at time zero if the agent exerts an effort profile $\vec{e}$ and reports failures and project success truthfully is:

$$\pi(\varpi, \vec{e}) = -(1-\beta_0)\sum_{t=1}^{T} \delta^t w_t$$

$$+\beta_0 \sum_{t=1}^{T} \delta^t (\prod_{s=1}^{t-1}(1-\lambda e_s))\big((e_t \lambda (V - b_t) - (1-\lambda e_t)w_t)\big).$$

The optimal contract will have to satisfy the following moral hazard constraint at each period for all possible histories and all possible effort paths in the future:

$$(MH) \quad \vec{1} \in arg \max_{\vec{e}} U(\varpi, \vec{e}).$$

Given that the MH constraint is satisfied, the principal's expected profit from offering contract ϖ at time zero becomes:

$$\pi\big(\varpi, \vec{1}\big) = -(1-\beta_0)\sum_{t=1}^{T} \delta^t w_t + \beta_0 \sum_{t=1}^{T} \delta^t (1-\lambda)^{t-1}(\lambda(V - b_t) - (1-\lambda)w_t).$$

[16] We refer to $e_t = 1$ as "work" and to $e_t = 0$ as "shirk".

Finally, since the agent is constrained financially, all the transfers from the principal to the agent must be nonnegative:

$(LL)\quad w_t, b_t \geq 0$ for $t = 1, \dots, T$.

The first-best benchmark

Consider the first-best case: The principal observes the effort choice and project outcome. As the relationship proceeds, if success is not being achieved, then every period the marginal benefit, $\lambda\tilde{\beta}_t V$, is the expected value of the project and takes into account both probability of success and current beliefs. Since beliefs are declining as time goes on without success, the marginal benefit decreases strictly. The marginal cost-of-effort, c, is constant. As a result, the first-best solution is characterized by stopping time $T \in \mathcal{N}$, such that the agent is allowed to exert effort up until that date only, as follows:

$$T^{FB} = arg \max_{t \in \mathcal{N}} \left\{\lambda\tilde{\beta}_t V \geq c\right\}^{17}.$$

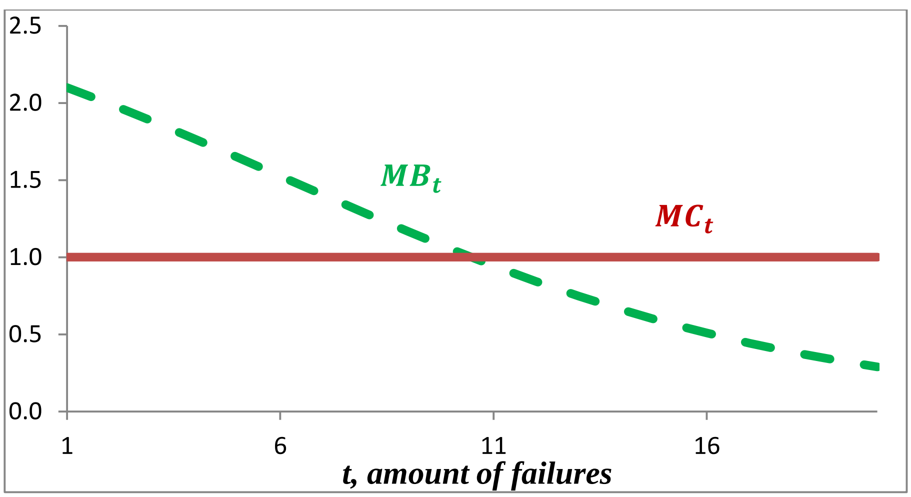

Figure 2. The firs best benchmark with $\lambda = 0.15$, $\beta_0 = 0.7$, $V = 20$ and $c = 1$.

Consider the example in Figure 2, where $\lambda = 0.15$, $\beta_0 = 0.7$, $V = 20$ and $c = 1$, and where the agent starts with $MB_1 = 2.1$ and continues experimenting with the project for ten periods at most.

[17] Recall that $\tilde{\beta}_t$ are beliefs evolved as a result of the agent exerting effort in all periods until t.

3. The Second-Best Contract

When the agent chooses effort level privately, the optimal contract must ensure the agent works in every period, which is guaranteed by the MH constraint. Since the agent is constrained financially, he cannot pay the principal, as the LL constraint reflects. Under certain circumstances, the agent can postpone or even hide success. To understand how this fact matters and if it affects the structure of the optimal contract, consider the agent's incentive to announce that he successfully completed the project at $t < T$. This decision is affected not only by the payment tied to success or failure in this particular period, as determined by the optimal contract, but, in addition, by payments in all subsequent periods of the planning stage. For example, if the discounted value of the promised reward for success in the future exceeds the current value, then the agent will postpone an announcement; if the agent is rewarded for consecutive failures,[18] he would benefit from hiding success completely. The IC constraint ensures the agent neither postpones success nor hides it. The principal's optimization problem in this case becomes the following[19]:

$$[\mathbb{P}^{SB}]\ \max_{\varpi} \pi\left(\varpi, \vec{1}\right) \text{ subject to}$$

(MH) $\vec{1} \in arg \max_{\vec{e}} U(\varpi, \vec{e})$,

(IC) $b_t \geq \delta b_{t+1}$ for $t = 1, \ldots, T-1$,

$b_t \geq w_t$ for $t = 1, \ldots, T$,

(LL) $w_t, b_t \geq 0$ for $t = 1, \ldots, T$.

Before we present a detailed solution to the principal's optimization problem, consider the agent's incentives to deviate at period $t \leq T$, assuming that the agent was behaving $1 \leq s < t$ without success in all prior periods and will work $T \geq s > t$ in all subsequent periods. In case the agent decides to shirk at the beginning of period $t \leq T$, his continuation value from the relationship is:

[18] For example, Manso (2011) and Chade and Kovrijnykh (2016) explored models where the agent is rewarded for delivering bad news in a different setting.
[19] We assume that V is high enough, and it is optimal for the principal when the agent exerts effort in every period.

$$U_t(\varpi, e_t = 0, e_s = 1 \text{ for } s \neq t) = w_t + \left(1 - \tilde{\beta}_t\right) \sum_{s=t+1}^{T} \delta^{s-t}(w_t - c)$$

$$+\tilde{\beta}_t \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s + (1-\lambda)w_s - c).$$

Note that if the agent follows this one-period deviation, he gets only w_t at period t, since he fails for sure. If the project is good, which, based on history, is true with current beliefs $\tilde{\beta}_t$, then the agent has a chance to succeed in all future periods $s > t$ until the relationship is terminated. If the project is bad, which is true with probability $1 - \tilde{\beta}_t$, the agent will receive w_t in all future periods despite exerting effort.

In contrast, if the agent decides to work at period t, his continuation value from the relationship is:

$$U_t\left(\varpi, \vec{e} = \vec{1}\right) = -c + \lambda\tilde{\beta}_t b_t + \left(1 - \tilde{\beta}_t\lambda\right)w_t + \left(1 - \tilde{\beta}_t\right) \sum_{s=t+1}^{T} \delta^{s-t}(w_t - c)$$

$$+\tilde{\beta}_t \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t}(\lambda b_s + (1-\lambda)w_s - c).$$

Notice that at period t, the agent has a chance to succeed and receive b_t. This occurs either if the project is good or with probability $\lambda\tilde{\beta}_t$. In case the agent is unlucky, with probability $1 - \tilde{\beta}_t\lambda$, he gets w_t, despite exerting high effort. As in the case where the agent deviates, if the project is bad, the agent will receive w_t for all future periods.

When the agent deviates at period t, he knows that failure at this period should not change beliefs and make parties more pessimistic regarding the project's quality. However, if this deviation is not observed by the principal, she will consider a failure reported at period t as a signal that the project is more likely to be bad. Importantly, this difference in beliefs reverberates into all future periods. Thus, in this model, the moral hazard problem in each period translates into asymmetrical information regarding beliefs about the project's quality in all consecutive periods.

Combining the two continuation values, the moral hazard constraint at period t (assuming that the agent was behaving in all prior periods $s < t$ and will work in all subsequent periods $s > t$) becomes the following:

$$(MH_t)\ b_t - w_t \geq \frac{c}{\lambda\tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s + (1-\lambda)w_s - c) \text{ for } 1 \leq t \leq T.$$

A gap between b_t and w_t can be divided into two components explicitly. First is the moral hazard rent, which arises because the agent chooses effort level privately. Second is the learning rent, which emerges if the agent deviates at some period and is relatively more optimistic than the principal until the end of the relationship. Note that the learning rent at the final period T is zero since the agent cannot exploit the fact that he is more optimistic. In addition, if the principal increases a reward for success at an arbitrary period t, b_t, he has to scale up all the rewards in all the previous periods automatically, as demonstrated in the MH_t constraint.

Proposition 1. The agent receives a positive reward if the project is implemented successfully and nothing otherwise. In particular,

$w_t = 0$ and $b_t = \frac{c}{\lambda\tilde{\beta}_t} + c\sum_{s=1}^{T-t}\delta^s \frac{(1-\beta_0)}{\beta_0(1-\lambda)^{t+s-1}}$ for $1 \leq t \leq T^{SB}$ with the following properties:

i) if $\delta = 1$ b_t is constant[20];
ii) if $0 < \delta < 1$ b_t is strictly increasing, whereas δb_t is strictly decreasing.

Moreover, the project is terminated inefficiently early; that is, $T^{SB} < T^{FB}$[21].

Proof: See Appendix A.

As demonstrated by the proposition above, the optimal contract makes the value of the discounted reward strictly decreasing. An immediate and perhaps fascinating conclusion from this observation is that the agent's private observability of success does not exacerbate the problem; that is, the agent will never postpone an announcement of success even if the IC constraint was not taken into account directly when solving the principal's optimization problem. This result plays a key role in our analysis, as we will demonstrate that private observability of success becomes critical when it comes to optimal monitoring timing.

[20] Halac, Kartik, and Liu (2016) argue that in the case of no discounting, the principal can be restricted to use constant bonus contracts. Bonatti and Horner (2011) have a similar result in their model with one agent only.
[21] Note that in this setting the optimal stopping time is deterministic. See Gerardi and Maestri (2012) for a formal proof and Green and Taylor (2015) for an example with a stochastic deadline.

To demonstrate the intuition behind Proposition 1, we clarify the dynamics of the moral hazard and learning rents. The moral hazard rent, $\frac{c}{\lambda\tilde{\beta}_t}$, is increasing strictly to account for the agent's increasing pessimism. The learning rent, $c\sum_{s=1}^{T-t}\delta^s\frac{(1-\beta_0)}{\beta_0(1-\lambda)^{t+s-1}}$, however is non-monotonic. Since $c\sum_{s=1}^{T-t}\delta^s\frac{(1-\beta_0)}{\beta_0(1-\lambda)^{t+s-1}} = \frac{c(1-\beta_0)}{\beta_0(1-\lambda)^{t-1}}\frac{\delta}{(1-\lambda-\delta)}\left(1-\left(\frac{\delta}{1-\lambda}\right)^{T-t}\right) = \frac{c\delta(1-\beta_0)}{\beta_0(1-\lambda-\delta)}\frac{(1-\lambda)^{T-t}-\delta^{T-t}}{(1-\lambda)^{T-1}}$, the learning rent is increasing at period $1 \le t \le T$ if and only if

$$\text{either } \delta < 1-\lambda \text{ and } \left(\frac{1-\lambda}{\delta}\right)^{T-t}\ln\delta > \ln(1-\lambda) \text{ or}$$

$$\delta > 1-\lambda \text{ and } \left(\frac{1-\lambda}{\delta}\right)^{T-t}\ln\delta < \ln(1-\lambda).$$

We present a particular example to better demonstrate the decomposition of the two rents in Figure 3 below. Suppose $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 1$, and $c = 1$. Note that when there is no discounting, the nominal value of b_t is constant, whereas the discounted value of the reward is decreasing, as suggested by Proposition 1. In this case, $\left(\frac{1-\lambda}{\delta}\right)^{T-t}\ln\delta > \ln(1-\lambda)$ and $\delta > 1-\lambda$, and the learning rent is decreasing for all periods.

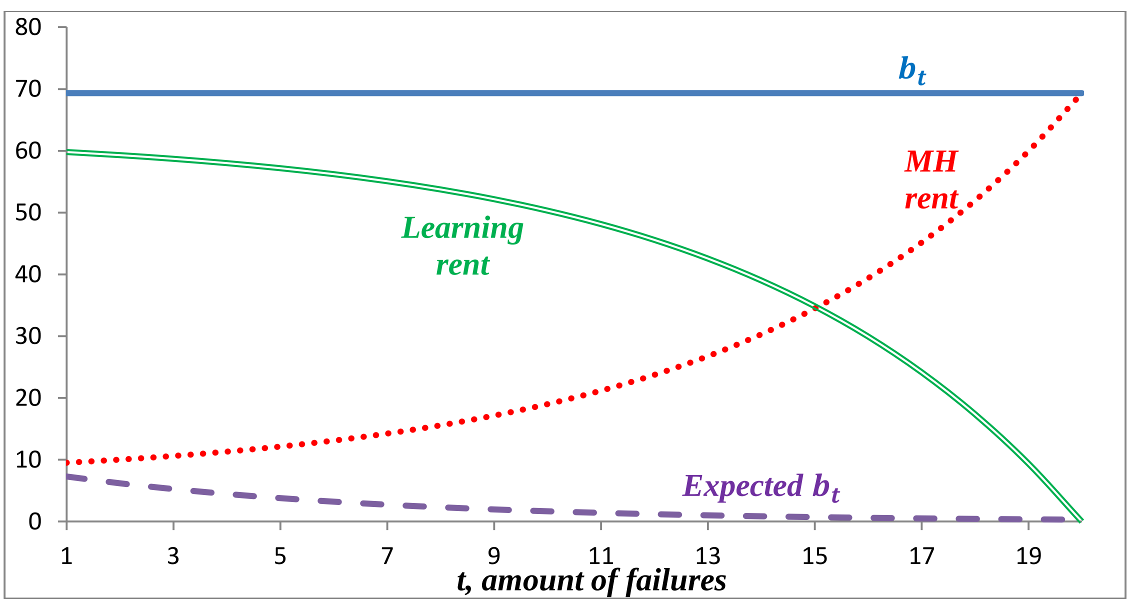

Figure 3. The optimal contract with $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 1$ and $c = 1$.

Now consider an example with discounting, as depicted in Figure 4 below. Suppose that $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 0.9$, and $c = 1$. In this case, $\delta > 1-\lambda$ and $\left(\frac{1-\lambda}{\delta}\right)^{T-t}\ln\delta < \ln(1-\lambda)$

for $t \leq 12$, and the learning rent is increasing for these periods and decreasing thereafter. As in the previous case, the moral hazard rent is increasing strictly to account for the agent's increasing pessimism.

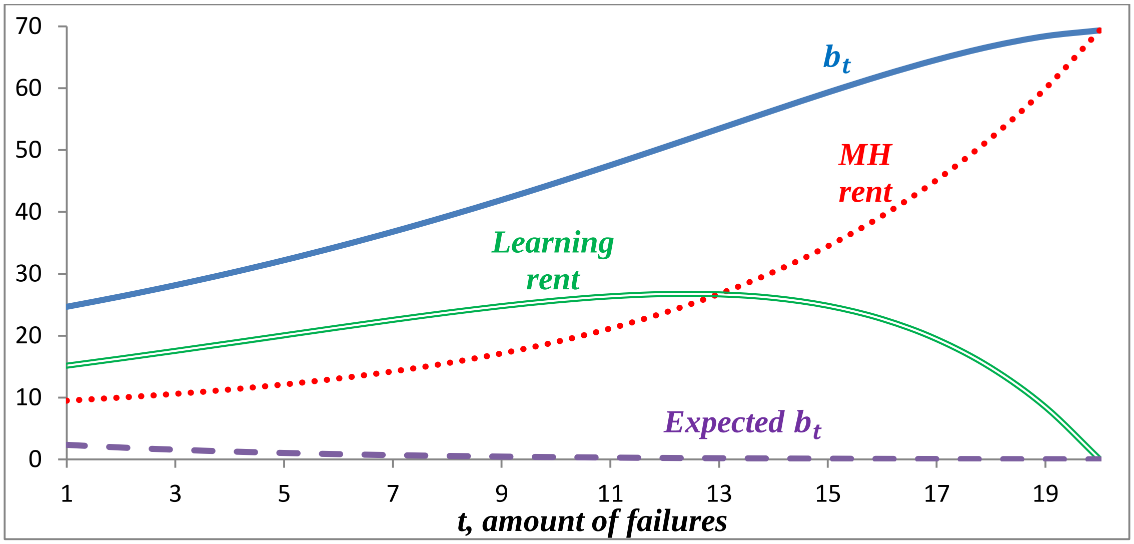

Figure 4. The optimal contract with $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 0.9$ and $c = 1$.

4. Monitoring

Given the optimal contract described in the previous section, the agent receives a strictly positive rent, and the project is terminated inefficiently early. One way the principal can alleviate this inefficiency is by hiring a monitor who, by assumption, can observe perfectly the effort level the agent chooses. The obvious benefit of hiring a monitor is that now, for some periods, the principal can promise to pay less since moral hazard problem is alleviated. It turns out the exact reduction of the reward depends on whether success might be hidden by the agent or not. This will be important for future analysis, since without monitoring, private observability does not play any role, whereas, as we will show, private observability is crucial for determining the optimal monitoring timing.

To keep exposition simple, we assume that monitoring allows a perfect assessment of the agent's effort; however, our results could be extended easily to account for noisy monitoring. In addition, we assume that monitoring costs $\gamma > 0$ per period, which might represent the cost of installing equipment like video cameras or recording devices.

We begin the analysis of the optimal timing of monitoring with an example, where the relationship lasts exogenously for two periods ($T = 2$), and the principal can perfectly observe the effort level during one period only when success is publicly observable. The principal's optimization problem (assuming it is optimal when the agent exerts effort in every period) is:

$$\max_{\varpi} \pi\left(\varpi, \vec{1}\right) \text{ subject to}$$

(MH) $(1{,}1) \in arg \max_{(e_1, e_2)} U(\varpi, \vec{e})$,

(LL) $w_1, w_2, b_1, b_2 \geq 0$.

First, suppose the principal chooses to monitor the agent at the beginning of the relationship. In this case, the MH constraint could be replaced by:

$$(MH_2)\ \lambda\tilde{\beta}b_2 + \left(1 - \lambda\tilde{\beta}\right)w_2 - c \geq w_2,$$

which ensures that the agent behaves at $t = 2$, given that success has not been achieved at period $t = 1$. Then, a solution to the optimization problem with monitoring involves $b_2 = \frac{c}{\lambda\tilde{\beta}}$ and $w_2 = 0$, where $\tilde{\beta} = \frac{\beta_0(1-\lambda)}{\beta_0(1-\lambda)+1-\beta_0}$. The principal's expected profit in this case is the following:

$$\pi_{m=1} = \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1} \lambda V - \delta^2 \beta_0 (1-\lambda)\lambda b_2 - \gamma$$

$$= \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1} \lambda V - \delta^2 c(1 - \beta_0\lambda) - \gamma.$$

Second, suppose the principal hires the monitor at $t = 2$. In this case, the MH constraint could be replaced by:

$$(MH_1)\ \lambda\beta_0 b_1 + (1 - \lambda\beta_0)w_1 - c \geq w_1,$$

which ensures that the agent behaves at $t = 1$, given that he will exert effort at $t = 2$. Then, the solution to the optimization problem involves $w_1 = b_2 = w_2 = 0$ and $b_1 = \frac{c}{\lambda\beta_0}$. The principal's expected profit is the following:

$$\pi_{m=2} = \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1} \lambda V - \delta\beta_0\lambda b_1 - \gamma = \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1} \lambda V - \delta c - \gamma.$$

Since $-\delta^2 c(1-\beta_0\lambda) > -\delta c$, it is optimal to monitor at $t = 1$. The intuition is that if monitoring occurs at $t = 1$, the principal expects to pay a reward at $t = 2$, conditional on the agent failing at $t = 1$, despite exerting effort as reflected by $(1-\beta_0\lambda)$ in the principal's expected profit $\pi_{m=1}$. The example above is straightforward but does not capture the main intuition fully, as when the relationship lasts for two periods and the monitor is hired, the agent does not receive any learning rent. We will demonstrate, however, that the result of this example extends to a general setting where the duration of the contract is long enough that the agent is granted a strictly positive learning rent.

Suppose now the principal can monitor the agent perfectly at any period $1 \leq m \leq T^M_{Public}$, where T^M_{Public} is the duration of the contract with monitoring when success is publicly observed. We would like to understand all the benefits from monitoring in this case. First, the principal can avoid paying b_m since the moral hazard problem at period m vanishes. This static effect is at the heart of our analysis, as it will be playing an important role when success is observed privately. Since the static effect alleviates the moral hazard problem at the period of monitoring, the principal benefits from it only if the agent succeeds at period m, which in turn is possible only if the project is good. Thus, the static effect is:

$$SE_m = \delta^m \Pr(\text{success at } t = m)\, b_m = \delta^m \beta_0 (1-\lambda)^{m-1} \lambda b_m.$$

Second, recall from the MH_m constraint that if the principal decreases a reward for success, b_m, he can scale down all the rewards in all the preceding periods, $1 \leq s < m$. This effect, which we call the *dynamic effect*, will be shown to play an auxiliary role in the environment we consider. Importantly, unlike with the static effect, the principal benefits from the dynamic effect even if the agent does not succeed at some period m. Since the agent always, except for the final period, has a chance to succeed in the later periods, the future rewards make it costlier to ensure the agent behaves at the beginning of the relationship. Intuitively, the promise of future monitoring echoes into the earlier periods, as it acts as a threat and it changes the agent's options if he decides to shirk in the earlier periods.

The dynamic effect is defined as:

$$DE_m = \sum_{i=1}^{m-1} \delta^i \Pr(\text{success at period } i < m)\, [\text{nominal decrease in } b_i].$$

Thus, the *total effect* of monitoring at period m, $TE_m = DE_m + SE_m$ combines the benefit of paying less at period m, which is decreasing in time, and the benefit of scaling down all the rewards in previous periods, which is non-monotonic. It turns out that the former effect is dominant, as stated in Proposition 2.

Proposition 2. When success cannot be hidden, monitoring is optimal at the beginning of the relationship. Moreover, the project is terminated inefficiently early: $T^{SB} \leq T^{M}_{Public} < T^{FB}$.

Proof: *See Appendix B.*

Let me finish this section with a numerical example. Suppose $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 0.9$, and $c = 1$. The static effect, $SE_t = \delta^t \beta_0 (1-\lambda)^{t-1} \lambda b_t$, is strictly decreasing, as reflected by a dashed line in Figure 5 below. The dynamic effect, $DE_t = \delta^t \beta_0 \lambda^2 (1-\lambda)^{t-2} (t-1) b_t$, is non-monotonic. For early periods $t \leq 6$, the agent has to be paid more to behave, since if he deviates once, he can leverage the fact that he is relatively more optimistic until the deadline. However, as time goes by without success, both parties become more pessimistic, and since the expected value of b_t goes down, the dynamic effect diminishes, as well.

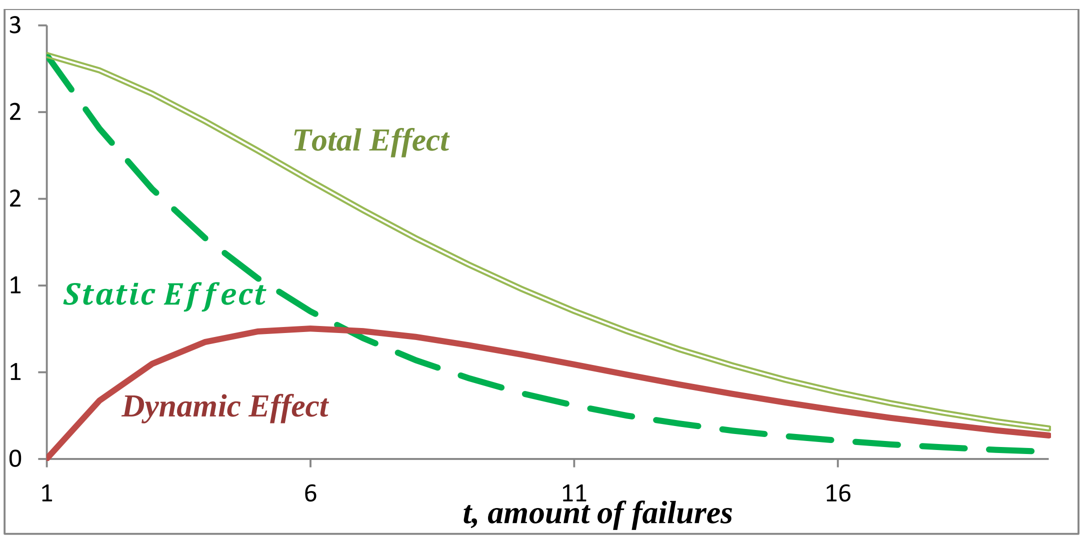

Figure 5. Effect from monitoring when success is publicly observable with $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 0.9$ and $c = 1$.

Suppose the agent can postpone an announcement of a successful implementation. We will denote $T^{M}_{Private}$ as the duration of the contract with monitoring when success is observed privately. The principal can still benefit from hiring a monitor; however, now the modified reward structure must ensure the agent does not have incentives to postpone or hide success. This is where the additional IC constraints:

$(IC)\quad b_t \geq \delta b_{t+1}$ for $t = 1, \dots, T-1$,

$b_t \geq w_t$ for $t = 1, \dots, T$,

become relevant and, as we will demonstrate, will be binding for the period when monitoring is implemented.

How can the principal benefit from monitoring when the agent observes success privately? First, as with public observability, the principal can pay less during the monitoring period because the moral hazard problem at period m vanishes. The static effect, however, is different. If the principal sets $b_m = 0$, then in the case the agent succeeds at this exact period, he will postpone an announcement until the later period. Thus, the static effect now has to be modified and becomes:

$$SE_m = \begin{cases} \delta^m \beta_0 (1-\lambda)^{m-1} \lambda (b_m - \delta b_{m+1}) \text{ for } 1 \leq m < T^M_{Private} \\ \delta^m \beta_0 (1-\lambda)^{m-1} \lambda b_m \text{ for } m = T^M_{Private} \end{cases}.$$

In addition, the dynamic effect, has to be modified to take into account the IC constraint, as well.

As in the previous case, we will first consider an example where the relationship lasts for two periods ($T = 2$), and the principal can perfectly observe the effort level during one period only when the agent privately observes success. The principal's optimization problem is:

$$\max_{\varpi} \pi(\varpi, \vec{1}) \text{ subject to}$$

$(MH)\ \vec{1} \in arg \max_{\vec{e}} U(\varpi, \vec{e})$,

$(LL)\ \{w_t\}_{t=1}^2, \{b_t\}_{t=1}^2 \geq 0$.

First, suppose the principal monitors the agent at $t = 1$. In this case, the MH constraint could be replaced by:

$(IC)\ b_1 \geq \delta b_2$,

$(MH_2)\ \lambda \tilde{\beta} b_2 + (1 - \lambda \tilde{\beta}) w_2 - c \geq w_2$, that ensures that the agent behaves at $t = 2$, given that success has not been achieved at period $t = 1$ and, in addition, that he does not postpone an announcement of success from the first period. Then, a solution to the optimization problem involves $w_1 = w_2 = 0$, $b_2 = \frac{c}{\lambda\tilde{\beta}}$, and $b_1 = \delta b_2 = \delta\frac{c}{\lambda\tilde{\beta}}$, where $\tilde{\beta} = \frac{\beta_0(1-\lambda)}{\beta_0(1-\lambda)+1-\beta_0}$. The principal's expected profit in this case becomes:

$$\pi_{m=1} = \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1}\lambda V - \delta^2 c(1-\beta_0\lambda) - \delta^2\beta_0\lambda\frac{c}{\lambda\tilde{\beta}} - \gamma$$

$$= \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1}\lambda V - \delta^2 c\left(1-\beta_0\lambda+\frac{\beta_0}{\tilde{\beta}}\right) - \gamma.$$

Second, suppose the principal hires a monitor at the second period. In this case, the MH constraint could be replaced solely by:

(MH_1) $\lambda\beta_0 b_1 + (1-\lambda\beta_0)w_1 - c \geq w_1,$

which ensures that the agent behaves at $t = 1$, given that he will exert effort at $t = 2$. Note that in this case, the principal does not have to pay anything at the final period, since the agent cannot benefit from hiding his early success.[22] The solution to the optimization problem involves $w_2 = 0$ and $b_1 = \frac{c}{\lambda\beta_0}$. The principal's expected profit is the following:

$$\pi_{m=2} = \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1}\lambda V - \delta\beta_0\lambda b_1 - \gamma = \beta_0 \sum_{i=1}^{2} \delta^i (1-\lambda)^{i-1}\lambda V - \delta c - \gamma.$$

It is optimal to monitor at $t = 2$ if $\delta^2 c\left(1-\beta_0\lambda+\frac{\beta_0}{\tilde{\beta}}\right) > \delta c$ or, equivalently, when:

$$\delta > \frac{(1-\lambda)}{(2-\lambda)(1-\beta_0\lambda)},$$

whereas if $\delta < \frac{(1-\lambda)}{(2-\lambda)(1-\beta_0\lambda)}$, monitoring is performed optimally at the beginning of the relationship.

The intuition is straightforward: If monitoring occurs at $t = 1$, the principal has to pay a reward to ensure the agent does not postpone announcing success, whereas if monitoring occurs at the final period, no positive reward is needed. We will show that it is always optimal to monitor at the end of the relationship when the discount factor is large enough.

[22] We assume that if the agent is indifferent between announcing success and postponing this announcement, he would choose the former always.

Proposition 3. When agent observes success privately and the discount factor is high enough, monitoring is used optimally at the end of the relationship. The project is terminated inefficiently early: $T^{SB} \leq T^{M}_{Private} \leq T^{M}_{Public} < T^{FB}$.

Proof: See Appendix C.

Consider an example in Figure 6 where the discount factor is not high enough for the monitoring to occur optimally at the end of the relationship.

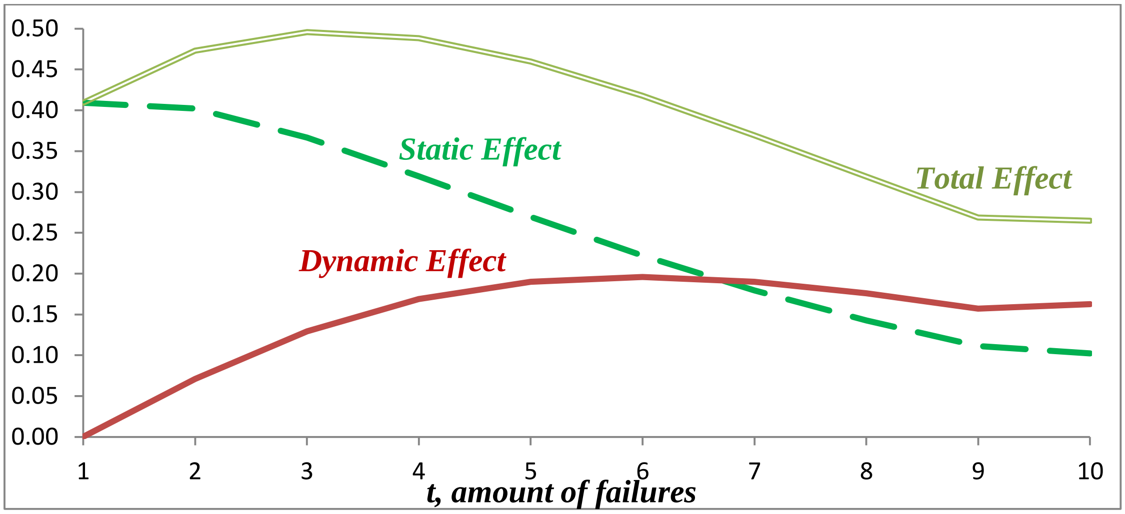

Figure 6. Effect from monitoring when success is private with $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 0.86$ and $c = 1$.

When $\delta = 0.86$, monitoring is employed optimally at the third period. However, if we increase the discount factor up to $\delta = 0.91$, monitoring will be employed optimally toward the end of the relationship. In this case, both the static and dynamic effects are non-monotonic, as Figure 7 demonstrates below.

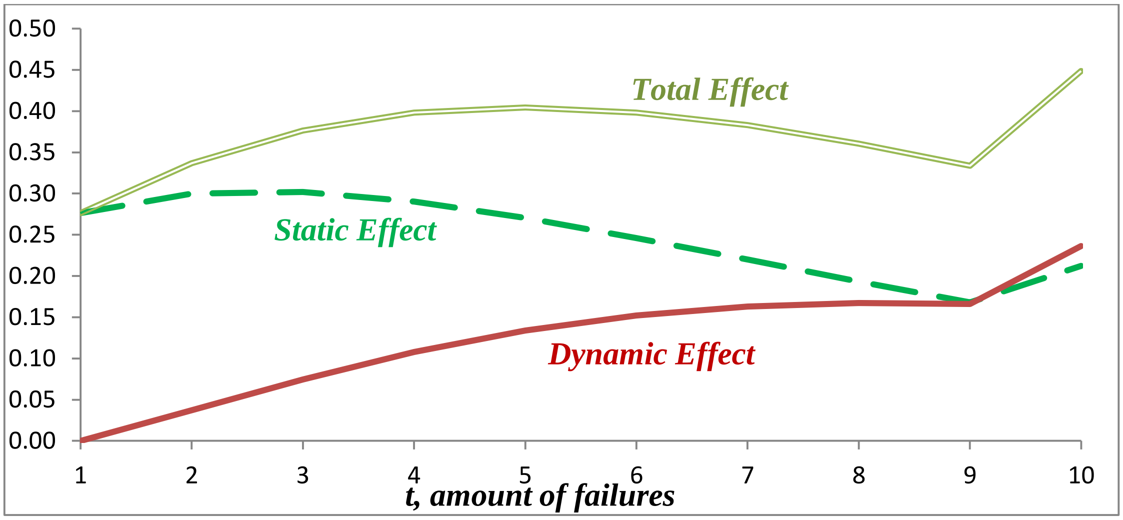

Figure 7. Effect from monitoring when success is private with $\lambda = 0.15$, $\beta_0 = 0.7$, $\delta = 0.9$ and $c = 1$.

5. Concluding Remarks

After studying optimal contracts for experimentation, we have found that the duration of the relationship between the principal and the agent is distorted compared to the first-best level; this leads to terminating the project early. The agent receives a twofold positive rent: He is compensated for exerting unobservable effort as well as for downgrading his beliefs regarding the project quality. Given the optimal contract, the agent is rewarded only if he succeeds and receives nothing otherwise. Observing success is irrelevant since the value of discounted reward is decreasing, though the nominal reward is increasing.

Monitoring at the beginning of the relationship improves the efficiency of financial contracting. This is in contrast with Bergemann and Hege's (1998) results, which demonstrated that monitoring is optimal toward the end of the project, when the agent is raising funds in a competitive market. In their settings, the agent benefits from monitoring the most and, as a result, monitoring is optimal toward the end of the project to extend the duration of the relationship. However, when the agent observes success privately and the parties to a contract are patient enough, it is again optimal to monitor at the end of the relationship, and the results of Bergemann and Hege (1998) prove true.

Our research demonstrates that without monitoring, the contract is in danger of being terminated inefficiently early. When the agent observes success privately, the duration of the relationship is extended when the monitor is hired. If success is observed commonly, the project is extended even further with monitoring. Thus, we emphasize a pivotal role of private observability and market structure on the optimal monitoring timing.

The results of this research are clear and testable. In the case that hiding success from the principal is prohibitively costly, the agent formed the original idea for the project, and investors are competing to fund the project, monitoring should be performed at the end of the trial periods. However, if the agent is hired by the owner of the project, monitoring is employed optimally at the beginning of the relationship. If success is enormously costly to observe, as is the case in this research, monitoring is performed toward the end if parties to a contract are patient enough. Although we did not explore this case, we suggest that the same result would hold in the environment considered by Bergemann and Hege (1998).

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.

Bhaskar V. (2012). Dynamic Moral Hazard, Learning and Belief Manipulation. *Working Paper*.

Bhaskar V. (2013). The Ratchet Effect Re-examined: A Learning Perspective. *Working Paper*.

Bonatti A. and Horner J. (2011). Collaborating. *American Economic Review*, 101, 632–663.

Bonatti A. and Horner J. (2012). Career Concerns with Coarse Information, *Working Paper*.

Cremer J. (1995). Arm's-Length Relationships. *Quarterly Journal of Economics* 110, 275-295.

DiMasi J., Hansen R., and Grabowski H. (2003). The Price of Innovation: New Estimates of Drug Development Costs. *Journal of Health Economics* 22: 151–85.

Gerardi D. and Maestri L. (2012). A Principal-Agent Model of Sequential Testing. *Theoretical Economics* 7, 425–463.

Gomes R., Gottlieb D. and Maestri L. (2015). Experimentation and Project Selection: Screening and Learning. *Working Paper*.

Gorman M. and Sahlman W. A. (1989). What do venture capitalists do? *Journal of Business Venturing* 4(4), 231–248

Halac M., N. Kartik and Q. Liu (2016). Optimal Contracts for Experimentation. *Review of Economic Studies*, *forthcoming*.

Hoffmann F. and Pfeil S. (2010). Reward for luck in a dynamic agency model. *Review of Financial Studies* 23, 3329-3345.

Holmstrom, Bengt (1979). Moral hazard and observability. *Bell Journal of Economics* 10, 74-91.

Horner J. and L. Samuelson (2014). Incentives for Experimenting Agents. *RAND Journal of Economics* 44(4), 632-663.

Inderst R. and Ottaviani M. (2012). Competition through commissions and kickbacks. *American Economic Review* 102, 780–809.

Keller G., Rady S. and Cripps M. (2005). Strategic Experimentation with Exponential Bandits. *Econometrica* 73, 39–68.

Kwon S. (2013). Dynamic Moral Hazard with Persistent States. *Working Paper*.

Lerner J. and U. Malmendier (2010). Contractibility and the Design of Research Agreements. *American Economic Review* 100, 214-246.

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

Mason R. and J. Valimaki (2015). Getting it done: dynamic incentives to complete a project. *Journal of the European Economic Association* 13 (1), 62-97.

Myerson R. B. (1986). Multistage Games with Communication. *Econometrica* 54 (2), 323-358.

Pakes, Ariel, and Shmuel Nitzan. (1983). Optimum contracts for research personnel, research employment, and the establishment of "rival" enterprises. *Journal of Labor Economics* 1(4), 345-365.

Pavan A., Segal I. and Toikka J. (2014). Dynamic Mechanism Design: A Myersonian Approach. *Econometrica* 82(2), 601-653.

Rothschild M. (1974). A Two-Armed Bandit Theory of Market Pricing. *Journal of Economic Theory* 9, 185–202.

Sahlman W. A. (1990). The Structure and Governance of Venture-Capital Organizations. *Journal of Financial Economics* 27, 473-521.

Shavell S. (1979). Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics* 10, 55-73.

Appendix A (The Second-Best Contract)

Proof of Proposition 1

The principal's optimization problem is the following:

$$[\mathbb{P}^{SB}]\ \max_{\varpi} \pi(\varpi, \vec{1}) \text{ subject to}$$

(MH) $\vec{1} \in arg \max_{\vec{e}} U(\varpi, \vec{e})$,

(IC) $b_t \geq \delta b_{t+1}$ for $t = 1, \dots, T-1$,

$b_t \geq w_t$ for $t = 1, \dots, T$,

(LL) $w_t, b_t \geq 0$ for $t = 1, \dots, T$.

We first solve an auxiliary problem $\mathbb{P}1$, where the global MH constraint is replaced by a sequence of MH_t constraints for $t = 1, \dots, T$ that ensure the agent does not want to deviate at period t, given that he was behaving in all prior periods $s < t$ and will work in all subsequent periods $s > t$. In addition, $\mathbb{P}1$ ignores the AS constraint, which will demonstrate automatic satisfaction.

The optimization problem $\mathbb{P}1$ is:

$$[\mathbb{P}1]\ \max_{\varpi} \pi(\varpi, \vec{1}) \text{ subject to}$$

(MH_t) $b_t - w_t \geq \frac{c}{\lambda \tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s + (1-\lambda)w_s - c)$ for $t = 1, \dots, T$,

(LL) $w_t,\ b_t \geq 0$ for $t = 1, \dots, T$.

Lemma 1. The following payment sequence solves $\mathbb{P}1$:

$$w_t = 0 \text{ and } b_t = \frac{c}{\lambda \tilde{\beta}_t} + c \sum_{s=1}^{T-t} \delta^s \frac{(1-\beta_0)}{\beta_0 (1-\lambda)^{t+s-1}} \text{ for } 1 \leq t \leq T.$$

Note that increasing w_t makes it more difficult to satisfy the MH_t constraints and lessens the objective function. As a result, the optimal solution must have $w_t = 0$ for $1 \leq t \leq T$, and the problem can be rewritten as:

$$[\mathbb{P}1]\ \max_{\varpi} \pi(\varpi, \vec{1}) \text{ subject to}$$

(MH_t) $b_t \geq \frac{c}{\lambda\tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c)$ for $t = 1, \dots, T$,

(LL) $b_t \geq 0$ for $t = 1, \dots, T$.

The auxiliary problem $\mathbb{P}1$ has the Lagrangian:

$$\mathcal{L} = \beta_0 \sum_{t=1}^{T} \delta^t (1-\lambda)^{t-1} \lambda (V - b_t)$$

$$+ \sum_{t=1}^{T} \mu_t \left(b_t - \frac{c}{\lambda\tilde{\beta}_t} - \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c) \right) + \sum_{t=1}^{T} \xi_t\, b_t.$$

The Kuhn-Tucker conditions for the optimization problem are:

$[b_t]$ $-\beta_0 \delta^t (1-\lambda)^{t-1}\lambda + \mu_t - \sum_{j=1}^{t-1} \mu_j \delta^{t-j}(1-\lambda)^{t-j-1}\lambda + \xi_t = 0$ for $1 \leq t \leq T$,

complemented by the constraints of the problem and the corresponding complementary slackness conditions.

If $\xi_t > 0$ for some $1 \leq t \leq T$, then MH_t would be violated and, as a result, we must have $\xi_t = 0$ for $t = 1, \dots, T$.

Consider the first-order conditions with respect to b_t:

$t = 1$: $-\beta_0 \delta \lambda + \mu_1 = 0 \Longrightarrow \mu_1 = \beta_0 \delta \lambda > 0$;

$t = 2$: $-\beta_0 \delta^2 \lambda (1-\lambda) + \mu_2 - \mu_1 \delta \lambda = 0 \Longrightarrow \mu_2 = \beta_0 \delta^2 \lambda > 0$;

repeating this procedure until the final period T, we have:

$t = T$: $\mu_T = \beta_0 \delta^T \lambda > 0$.

Thus, all MH_t constraints must be binding.

First, we will prove that with $b_t = \frac{c}{\lambda\tilde{\beta}_t} + c \sum_{s=1}^{T-t} \delta^s \frac{(1-\beta_0)}{\beta_0 (1-\lambda)^{t+s-1}}$, it is the case that:

$$b_t = \frac{c}{\lambda\tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c) \text{ for } 1 \leq t \leq T.$$

For $t = T$, the equality $b_T = \frac{c}{\lambda\tilde{\beta}_T}$ trivially follows from the proposed formula itself.

For any $t < T$, assume that MH_s holds for $s = t + 1$; that is:

$$b_{t+1} = \frac{c}{\lambda\tilde{\beta}_{t+1}} + \sum_{s=t+2}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c).$$

We need to show that $b_t = \frac{c}{\lambda\tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c)$, or, using the line above:

$$b_t = \frac{c}{\lambda\tilde{\beta}_t} + \delta(\lambda b_{t+1} - c) + \sum_{s=t+2}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c)$$

$$= \frac{c}{\lambda\tilde{\beta}_t} + \delta(\lambda b_{t+1} - c) + \left(b_{t+1} - \frac{c}{\lambda\tilde{\beta}_{t+1}}\right) = \frac{c}{\lambda\tilde{\beta}_t} - \frac{c}{\lambda\tilde{\beta}_{t+1}} - \delta c - (1+\delta\lambda)b_{t+1}.$$

Since $b_{t+1} = \frac{c}{\lambda\tilde{\beta}_{t+1}} + \frac{c(1-\beta_0)}{\beta_0(1-\lambda)^t}\frac{\delta}{(1-\lambda-\delta)}\left(1 - \left(\frac{\delta}{1-\lambda}\right)^{T-t-1}\right)$ and

$b_t = \frac{c}{\lambda\tilde{\beta}_t} + \frac{c(1-\beta_0)}{\beta_0(1-\lambda)^{t-1}}\frac{\delta}{(1-\lambda-\delta)}\left(1 - \left(\frac{\delta}{1-\lambda}\right)^{T-t}\right)$, it suffices to show that:

$$\frac{c}{\lambda\tilde{\beta}_t} + \frac{c(1-\beta_0)}{\beta_0(1-\lambda)^{t-1}}\frac{\delta}{(1-\lambda-\delta)}\left(1 - \left(\frac{\delta}{1-\lambda}\right)^{T-t}\right)$$

$$= \frac{c}{\lambda\tilde{\beta}_t} - \frac{c}{\lambda\tilde{\beta}_{t+1}} - \delta c - (1+\delta\lambda)\left(\frac{c}{\lambda\tilde{\beta}_{t+1}} + \frac{c(1-\beta_0)}{\beta_0(1-\lambda)^t}\frac{\delta}{(1-\lambda-\delta)}\left(1 - \left(\frac{\delta}{1-\lambda}\right)^{T-t-1}\right)\right),$$

which is easily verified for any $1 \le t \le T$. *Q.E.D.*

We will demonstrate that with the proposed solution, any period is optimal for the agent to work in, regardless of previous effort history profile. Consider the final period T. Note that if the agent deviates and shirks his duties at some arbitrary period $t < T$, he only can be more optimistic at period T. Thus, for any history of prior effort, the current belief β_T can be higher only than $\tilde{\beta}_t$. Now (MH_T) $\lambda\tilde{\beta}_T b_T = c$ is satisfied since $\beta_T \ge \tilde{\beta}_T$ and $\lambda\beta_T b_T \ge c$. Next, assume that working in any period is optimal for the agent, regardless of the previous effort history profile at period $t + 1 \le T$. Consider period t as any history of prior effort with current beliefs β_t. Since we already showed that $b_t = \frac{c}{\lambda\tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c)$, for any $\beta_t \ge \tilde{\beta}_t$ it is apparent that $b_t \ge \frac{c}{\lambda\beta_t} + \sum_{s=t+1}^{T} \delta^{s-t}(1-\lambda)^{s-t-1}(\lambda b_s - c)$, and working is optimal for the agent.

Finally, it can be shown by induction that any reward profile that makes every MH_t constraint binding must coincide with $b_t = \frac{c}{\lambda \tilde{\beta}_t} + c \sum_{s=1}^{T-t} \delta^s \frac{(1-\beta_0)}{\beta_0 (1-\lambda)^{t+s-1}}$ for $1 \leq t \leq T$. Moreover, since we already proved that working for the agent is optimal in any period, regardless of the previous effort history profile, this also ensures that the agent would find it optimal to work in period t for any possible effort profile before t.

Recall that when solving $\mathbb{P}1$, we ignored the IC constraint. Since $w_t = 0$ and $b_t > 0$, the proposed solution obviously satisfies $b_t \geq w_t$ for $t = 1, \dots, T$. Thus the final condition we must check is $b_t \geq \delta b_{t+1}$ for $t = 1, \dots, T-1$.

Given that $\sum_{s=1}^{T-t} \frac{\delta^{s-1}}{(1-\lambda)^{s-1}} = \frac{1-\left(\frac{\delta}{1-\lambda}\right)^{T-t}}{1-\frac{\delta}{1-\lambda}}$, by performing some algebra, one could verify that:

$$b_t = \frac{c}{\lambda \tilde{\beta}_t} + c\delta \frac{(1-\beta_0)}{\beta_0} \frac{1-\left(\frac{\delta}{1-\lambda}\right)^{T-t}}{(1-\lambda)^{t-1}(1-\lambda-\delta)} = \frac{c}{\lambda \tilde{\beta}_t} + c\delta \frac{(1-\beta_0)}{\beta_0} \frac{(1-\lambda)^{T-t} - \delta^{T-t}}{(1-\lambda)^{T-1}(1-\lambda-\delta)}.$$

Then, it follows that $b_t \geq \delta b_{t+1}$ if and only if:

$$\frac{c}{\lambda}\left(\frac{\beta_0(1-\lambda)^{t-1}+1-\beta_0}{\beta_0(1-\lambda)^{t-1}} - \delta \frac{\beta_0(1-\lambda)^t+1-\beta_0}{\beta_0(1-\lambda)^t}\right) \geq \frac{\delta c(1-\beta_0)\left(\delta(1-\lambda)^{T-t-1} - \delta^{1+T-t-1} - (1-\lambda)^{T-t} + \delta^{T-t}\right)}{\beta_0(1-\lambda)^{T-1}(1-\lambda-\delta)},$$

$$\frac{c\left(\beta_0(1-\lambda)^t(1-\delta) + (1-\beta_0)(1-\lambda-\delta)\right)}{\lambda \beta_0 (1-\lambda)^t} \geq -\frac{\delta c(1-\beta_0)}{\beta_0(1-\lambda)^t},$$

or, alternatively, $(1-\delta)\left(\beta_0(1-\lambda)^t + (1-\beta_0)(1-\lambda)\right) \geq 0,$

which holds as an equality if $\delta = 1$ (and as a strict inequality as long as $\delta < 1$ for any t).

Thus, $w_t = 0$ and $b_t = \frac{c}{\lambda \tilde{\beta}_t} + c \sum_{s=1}^{T-t} \delta^s \frac{(1-\beta_0)}{\beta_0(1-\lambda)^{t+s-1}}$ for $1 \leq t \leq T$ is a solution to the principal's optimization problem.

Finally, since $\lambda \tilde{\beta}_t b_t > c$, for all $1 \leq t \leq T-1$, the project is terminated inefficiently early, $T^{SB} < T^{FB}$.

Q.E.D

Appendix B (The Optimal Timing of Monitoring When Success Cannot Be Hidden)

Proof of Proposition 2

We will call T_{Public}^{M} the duration of the contract when the principal performs monitoring and success is observed publicly. The total benefit, TE_m, from monitoring is a sum of the static and dynamic effects. The static effect from monitoring at period m, SE_m is:

$$SE_m = \delta^m \beta_0 (1-\lambda)^{m-1} \lambda b_m,$$

where $b_m = \frac{c}{\lambda \tilde{\beta}_m} + c \sum_{s=1}^{T-m} \delta^s \frac{(1-\beta_0)}{\beta_0 (1-\lambda)^{m+s-1}}$.

The dynamic effect is:

$$DE_m = \sum_{i=1}^{m-1} \delta^i \, \text{Pr(success at period } i < m) \, [\text{nominal decrease in } b_i].$$

First, we need to define a nominal decrease in b_i for $i < m$ that is possible because of monitoring that will occur at period m. Recall that the optimal payment structure makes all MH_t constraint binding or, equivalently,

$$b_t = \frac{c}{\lambda \tilde{\beta}_t} + \sum_{s=t+1}^{T} \delta^{s-t} (1-\lambda)^{s-t-1} (\lambda b_s - c),$$

and by decreasing a reward in a certain period m (in the right-hand side), the principal can decrease a reward in the left-hand side. Thus, a nominal decrease in b_i for $i < m$ is:

$$\delta^{m-i} (1-\lambda)^{m-i-1} \lambda b_m.$$

As a result, the dynamic effect becomes:

$$DE_m = \sum_{i=1}^{m-1} \delta^i \big(\beta_0 (1-\lambda)^{i-1} \lambda\big) \, \delta^{m-i} (1-\lambda)^{m-i-1} \lambda b_m = \delta^m \beta_0 \lambda^2 (1-\lambda)^{m-2} (m-1) b_m.$$

We can then calculate the total effect from monitoring at period m:

$$TE_m = \delta^m \beta_0 (1-\lambda)^{m-1} \lambda b_m + \delta^m \beta_0 \lambda^2 (1-\lambda)^{m-2} (m-1) b_m$$

$$= \delta^m \beta_0 (1-\lambda)^{m-2} \lambda \big(1 - \lambda + \lambda (m-1)\big) b_m.$$

We will prove that TE_m is strictly decreasing in m. First, recall form Proposition 1 that b_t was chosen optimally such that $b_t \geq \delta b_{t+1}$ for $t = 1, \dots, T-1$, and, as a result, $\delta^m b_m$ is decreasing in m.

It suffices to show that $\varphi(m) = (1-\lambda)^{m-2}\big(1-\lambda+\lambda(m-1)\big)$ is decreasing in m as well. Notice that $\varphi(1) = \varphi(2) = 1$.

Because $\frac{d\varphi(m)}{dm} = (1-\lambda)^{m-2}\Big(\lambda + \big(1-\lambda+\lambda(m-1)\big)\ln(1-\lambda)\Big)$, it is sufficient to show that $f(\lambda, m) = \lambda + \big(1-\lambda+\lambda(m-1)\big)\ln(1-\lambda)$ is negative for $m > 1$. Note that $\frac{\partial f}{\partial m} = \lambda\ln(1-\lambda) < 0$ for any m and $f(\lambda, 2) = \lambda + \ln(1-\lambda)$. Consider $g(\lambda) = f(\lambda, 2) = \lambda + \ln(1-\lambda)$, with $g'(\lambda) = -\frac{\lambda}{1-\lambda}$ being negative for all $0 < \lambda < 1$ and $\lim_{\lambda\to+0} g(\lambda) = 0$. Thus, $f(\lambda, m)$ is negative, and, as a result, $\frac{d\varphi(m)}{dm}$ is decreasing in m.

Since TE_m is decreasing in m, monitoring is implemented optimally at the very first period. Since the principal can promise paying less, he can use these funds to extend the duration of the relationship up to T^M_{Public} and, as a result, $T^{SB} < T^M_{Public}$. Since the agent still receives a positive rent, the project is still terminated inefficiently early: $T^M_{Public} < T^{FB}$.

Q.E.D

Appendix C (The Optimal Timing of Monitoring When Success Might Be Hidden)

Proof of Proposition 3

We will call $T^M_{Private}$ the duration of the contract when the agent observes success privately and the principal performs monitoring. As in the case when success is publicly observed, the total effect, TE_m, of monitoring is a sum of the static and dynamic effects. However, for this case, we redefine both effects to account for the possibility of hiding success. In particular, the additional IC constraints become relevant:

$$b_t \geq \delta b_{t+1} \text{ for } t = 1, \dots, T^M_{Private} - 1,$$

$$b_t \geq w_t \text{ for } t = 1, \dots, T^M_{Private}.$$

Since we proved that all $w_t = 0$, the second constraint will not affect either of the two effects, whereas the first constraint will limit the amount of money the principal will save by monitoring, except during the final period of the relationship. The static effect of monitoring at period m, SE_m is:

$$SE_m = \delta^m \beta_0 (1-\lambda)^{m-1} \lambda (b_m - \delta b_{m+1}),$$

where $b_m = \frac{c}{\lambda \tilde{\beta}_m} + c \sum_{s=1}^{T-m} \delta^s \frac{(1-\beta_0)}{\beta_0 (1-\lambda)^{m+s-1}}$[23].

To simplify notation, we will define function η_m as follows:

$$\eta_m = \begin{cases} b_m - \delta b_{m+1} \text{ for } 1 \leq m < T^M_{Private} \\ b_m \text{ for } m = T^M_{Private} \end{cases}$$

Thus, $SE_m = \delta^m \beta_0 (1-\lambda)^{m-1} \lambda \eta_m$ for $t = 1, \dots, T^M_{Private}$.

The dynamic effect is then:

$$DE_m = \sum_{i=1}^{m-1} \delta^i \left(\beta_0 (1-\lambda)^{i-1} \lambda\right) \delta^{m-i} (1-\lambda)^{m-i-1} \lambda \eta_m = \delta^m \beta_0 \lambda^2 (1-\lambda)^{m-2} (m-1) \eta_m.$$

Thus, the total effect becomes

[23] For convenience, we will say that $b_{T^M_{Private}+1} = 0$.

$$TE_m = \delta^m \beta_0 (1-\lambda)^{m-2} \lambda \big(1-\lambda+\lambda(m-1)\big) \eta_m = \delta^m \beta_0 \lambda \varphi(m) \eta_m.$$

First, consider the case when $\delta = 1$. From Proposition 1, it follows that optimal b_m is constant over time, and, as a result, $\eta_m = 0$ for $t = 1, \dots, T^M_{Private} - 1$. For the final period, however, $\eta_{T^M_{Private}} = b_{T^M_{Private}} > 0$. Cleary, monitoring at the final period is optimal.

Consider now the case for $\delta < 1$. Recall that from Proposition 1, it follows that b_m is increasing strictly, whereas δb_{m+1} is decreasing strictly. In particular, for $t = 1, \dots, T^M_{Private} - 1$:

$$\eta_m = b_m - \delta b_{m+1} = \frac{c}{\lambda}\left(\frac{\beta_0(1-\lambda)^{t-1}+1-\beta_0}{\beta_0(1-\lambda)^{t-1}} - \delta \frac{\beta_0(1-\lambda)^{t}+1-\beta_0}{\beta_0(1-\lambda)^{t}}\right)$$

$$-\frac{\delta c(1-\beta_0)\big(\delta(1-\lambda)^{T-t-1}-\delta^{1+T-t-1}-(1-\lambda)^{T-t}+\delta^{T-t}\big)}{\beta_0(1-\lambda)^{T-1}(1-\lambda-\delta)}$$

$$= \frac{c}{\lambda}\left(\frac{\beta_0(1-\lambda)^t(1-\delta)+(1-\lambda-\delta)(1-\beta_0)}{\beta_0(1-\lambda)^t}\right) + \frac{\delta c(1-\beta_0)}{\beta_0(1-\lambda)^t} = \frac{c(1-\delta)\big(\beta_0(1-\lambda)^{t-1}+1-\beta_0\big)}{\lambda\beta_0(1-\lambda)^{t-1}}.$$

As a result, η_m evolves as follows:

$$\eta_m = \begin{cases} \frac{c(1-\delta)\big(\beta_0(1-\lambda)^{m-1}+1-\beta_0\big)}{\lambda\beta_0(1-\lambda)^{m-1}} & \text{for } 1 \le m < T^M_{Private} \\ \frac{c\big(\beta_0(1-\lambda)^{m-1}+(1-\beta_0)\big)}{\lambda\beta_0(1-\lambda)^{m-1}} & \text{for } m = T^M_{Private} \end{cases}.$$

$$TE_m = \delta^m \beta_0 (1-\lambda)^{m-2} \lambda \big(1-\lambda+\lambda(m-1)\big) \eta_m$$

$$= \begin{cases} (1-\delta)c\,\frac{\delta^m\big(1-\lambda+\lambda(m-1)\big)\big(\beta_0(1-\lambda)^{m-1}+1-\beta_0\big)}{(1-\lambda)} & \text{for } 1 \le m < T^M_{Private} \\ c\,\frac{\delta^m\big(1-\lambda+\lambda(m-1)\big)\big(\beta_0(1-\lambda)^{m-1}+(1-\beta_0)\big)}{(1-\lambda)} & \text{for } m = T^M_{Private} \end{cases}.$$

We define time period $1 \le j < T^M_{Private}$ as follows:

$$j \in \arg\max_{1 \le s < T^M_{Private}} \left\{ c\,\frac{\delta^s\big(1-\lambda+\lambda(s-1)\big)\big(\beta_0(1-\lambda)^{s-1}+1-\beta_0\big)}{(1-\lambda)} \right\}^{24}.$$

[24] Since $T^M_{Private} < \infty$, a time period j is well defined, although it might be not be unique due to the non-monotonicity of $\delta^m\big(1-\lambda+\lambda(m-1)\big)\big(\beta_0(1-\lambda)^{m-1}+1-\beta_0\big)$ in m.

We will show that with a high enough discount factor, monitoring at the final period is always optimal. First, there is a (unique) $\tilde{\delta}$, such that:

$$(1-\tilde{\delta})c\frac{\delta^{j}\big(1-\lambda+\lambda(j-1)\big)\big(\beta_0(1-\lambda)^{j-1}+1-\beta_0\big)}{(1-\lambda)} = c\frac{\delta^{T^M_{Private}}\Big(1-\lambda+\lambda\big(T^M_{Private}-1\big)\Big)\Big(\beta_0(1-\lambda)^{T^M_{Private}-1}+(1-\beta_0)\Big)}{(1-\lambda)}$$

or, equivalently, $$\tilde{\delta} = 1-\frac{c\frac{\delta^{T^M_{Private}}\left(1-\lambda+\lambda\left(T^M_{Private}-1\right)\right)\left(\beta_0(1-\lambda)^{T^M_{Private}-1}+(1-\beta_0)\right)}{(1-\lambda)}}{c\frac{\delta^{j}\left(1-\lambda+\lambda(j-1)\right)\left(\beta_0(1-\lambda)^{j-1}+1-\beta_0\right)}{(1-\lambda)}} \leq 1.$$

For any $\delta > \tilde{\delta}$, the total effect of monitoring at period m, TE_m achieves its highest value at the final period $T^M_{Private}$, whereas when $\delta < \tilde{\delta}$ monitoring is implemented, it is optimal at period j, which, in general, is not the final period. As the discount factor increases, both the static and dynamic effects diminish, and monitoring is implemented optimally at the final period. Since the principal promises paying less, the money that he saves could be used to extend the duration of the relationship beyond $T^M_{Private}$. In comparison with the case in which success is observed publicly, the principal saves less money; however, the second-best outcome is improved marginally. As a result, $T^{SB} \leq T^M_{Private} \leq T^M_{Public} < T^{FB}$.

Q.E.D