Contracting for Experimentation and the Value of Bad News*

Sinem Hidir†

September 24, 2019

This paper studies the optimal provision of incentives to acquire information about a project's quality. News arrives in the form of a conclusive good signal or nonconclusive bad signals while the agent invests in information and both effort and signals are private information. The optimal contract consists of history dependant payments upon terminal signals and a deadline that is extended upon the disclosure of bad signals, as long as information acquisition remains efficient. The extensions of the deadline help back-load the payments and mitigate the inefficiency caused by early stopping while keeping the agent's rent constant. When there is stopping before a deadline, it happens at the first best stopping belief.

Keywords: Dynamic moral hazard, continuous time principal-agent model, experimentation, Poisson arrivals, private signals, Bayesian learning.

JEL Codes: D82, D83, D86, O32.

*This paper is based on the first chapter of my PhD thesis presented at the Toulouse School of Economics, June 2015. I am grateful to Jacques Crémer, Christian Hellwig, and Thomas Mariotti for their advice and I also benefited from discussions with thank Alessandro Bonatti, George Georgiadis, Renato Gomes, Qingmin Liu, George Lukyanov, Lucas Maestri, Andrea Prat, Harry Di Pei, Yuliy Sannikov, Teck Yong Tan and various seminar and conference participants.

†University of Warwick. email: s.hidir@warwick.ac.uk.

1 Introduction

Firms often consult experts in order to acquire information about new and untested technologies. Consider a pharmaceutical company who provides funds to a scientist for performing tests on a new drug. The scientist has the duty of testing different compounds to find out about their efficiency or side effects. In such a setting, both positive and negative information arise only as a result of costly research. Similarly, during the development of a new product, market research may reveal a low potential demand which is interpreted as bad news making it less profitable to continue.

In these examples, there is uncertainty about the time required to obtain conclusive information. In addition, the incentives of the principal who wants to take the decision and the agent who is hired to acquire information are not aligned. The principal would like to obtain information, whether both good or bad, as soon as possible and avoid spending resources inefficiently, potentially diverting them onto other projects. However, the agent is keen on getting paid for as long as possible hence would like to prolong the contract as much as possible. This paper solves for the optimal way through which incentives must be provided for the agent to acquire information and report the intermediary as well as terminal findings so that the principal will either decide to implement the project or stop spending resources on it.

The problem considered is one of optimal contracting for information acquisition with moral hazard and endogenously arising private information. The agent has a choice between investing in information or shirking to keep the resources provided by the principal for his private benefit. While the agent invests in information, signals arrive via a Poisson arrival process at a rate that doesn't depend on the project type. Conditional on the project being good, a good or a bad signal can realize while for a bad project, only a bad signal can ever realize. Hence, one good signal is sufficient to conclude that it is optimal to invest in the project. As a good signal ends the contract, the agent must be given sufficient incentives to disclose it. A bad signal decreases the belief about the project being good and after sufficiently many bad signals, it is optimal to stop investing in information, even without an agency problem. Again, the agent should be incentivized to disclose bad signals. The principal's problem is then one of finding the optimal contract which incentivizes the agent to both acquire information and disclose it.

We find that the optimal contract consists of history dependant payments upon terminal signal disclosures, an initial deadline and a rule for extension after bad signal disclosures. These extensions are optimal as long as the belief is above that belief at which the principal would stop in the first best benchmark. Stopping before reaching a deadline happens either upon a *terminal* bad signal disclosure[1] or upon a good signal disclosure. Hence, even though many times contract may stop inefficiently early due to reaching a deadline without learning, stopping before a deadline only happens *efficiently*. In addition, the only payments to the agent are made upon first-best stopping.

To understand the intuition, it must be seen that the optimal contract needs to satisfy two types of incentive constraints: first, the *ex-ante* incentive constraint to work and second the *ex-post* incentive constraint to disclose the signals. It is a simplifying feature that the agent's shirking does not lead to a belief divergence between him and the principal, as no signals arrive when the agent shirks. Hence, there is no informational rent from postponing effort of the type called *procrastination rents* in the literature on experimentation with agency.[2] In addition, although there are many possible deviations available to the agent once having acquired a signal, the agent's payoff is bounded by the one in which he shirks until the deadline. The agent's working constraint binds at any moment in the contract and the total agency rent is only a function of the time initially allocated to him, while the incentive constraints define how much of this rent can be provided through continuation values as opposed to direct payments.

Payments upon termination decrease as getting closer to the deadline for a given belief, with upward jumps after the disclosure of each additional bad signal until the stopping belief is reached. In order to induce the agent to reveal a good signal or a terminal bad signal, the principal has to compensate him for the possibility of hiding it and shirking until the deadline. A payment is made to the agent only upon the disclosure of a good signal at any point in time or upon a bad signal which pushes the belief below the first best threshold and leads to stopping.[3] Although the agent is willing to reveal non terminal bad signals without any reward as long as his continuation value does not decrease, it is optimal to promise more contract time as this relaxes his incentive constraint to work as well as providing longer time for experimentation which still has positive value.

---

[1] A terminal signal is used to denote either a good signal or the first bad signal upon which the belief falls below the first best experimentation threshold. This can be thought of as a certain number of research methods that lead to bad signal that it is no longer optimal to try another one.

[2] The procrastination rent is present in other papers on dynamic agency where the belief goes down as long as no success is observed. For example, in Horner and Samuelson (2013), the principal finds it optimal to downsize the project or take his outside option for some periods in order to decrease the informational rent of the agent. In our setting, the principal does not find it optimal to take the outside option and continue with the project later on. Similarly, in Halac et al. (2016) and Moroni (2016) among others, an informational rent is born because the agent's deviation leads him to hold a different prior than the principal.

[3] This threshold is the belief at which experimentation ends upon a bad signal disclosure and will be solved for in the paper, equivalently the number of methods that should be tried until stopping. It will correspond to the same stopping belief as in first best.

The intuition for the extensions is as follows: a longer time horizon provides more time to potentially discover a good project but also makes it more costly to induce the agent to disclose terminal signals. Hence, committing to a deadline keeps the moral hazard rents under control. As time passes without a signal being disclosed, it becomes more likely that the agent has been shirking and termination serves as a punishment. At the same time, the deadline sometimes causes experimentation to end inefficiently early without the belief having decreased. The disclosure of a bad signal reveals that the agent has been exerting effort and the extension of the deadline decreases the distortion due to the deadline. As the agent's incentive constraint binds at any moment, these extensions do not increase the agency rent compared to a contract with a fixed deadline. Indeed, there are other incentive compatible contracts with identical cost in which the deadline doesn't change upon the disclosure of intermediary bad signals (or gets extended by a lower amount) but the disclosure of terminal signals requires higher payments. The optimal contract with extended deadlines improves on this contract by providing longer time for learning at the same cost. Finally, when the belief falls below the first-best threshold before a deadline, the contract ends without reaching the deadline. At this point, as the value of information is negative without agency problem, it is optimal to stop immediately.

1.1 Related Literature

This paper relates closely to the literature on dynamic agency with learning. A closely related paper is by [Gerardi and Maestri (2012)] which studies contracting with moral hazard to acquire soft information about a project quality in each period. As the agent is sure to get a signal in each period while he incurs cost, there is a latest date at which learning will be complete. As termination payment can be conditioned both on the recommendation and the outcome, the agent's rent comes from the possibility of guessing a good state without acquiring information. In our setting, in addition to being hard information, signals arrive at random times while effort is exerted. Hence, there is no predetermined date by which sufficient information should be acquired.

Klein (2016) considers a setup where the agent has the choice between experimenting on a risky arm to prove a hypothesis or cheating to create fake successes. In the optimal contract, to make cheating unattractive, the principal rewards only when the success is repeated several times, given that if the risky arm is good, successes arrive faster than on the cheating option. In our setup, one good signal is conclusive but several bad signals may be disclosed before stopping, and we are interested in how much time to allocate to the agent and incentivize termination via disclosing signals.

There is a wide literature on learning in dynamic agency with a cash diversion problem, such as Bergemann and Hege (1998, 2005) and Horner and Samuelson (2013). The learning structure in our setup is different. In the mentioned works, learning happens through good news conditional on a good project and the lack of good news is interpreted as a bad news or a failure[4]. Hence, even without an agency problem, as time goes by the belief would be pessimistic enough that the project should be abandoned. It is not possible to distinguish an agent who is shirking from an agent who is working on a bad project. In our setting, news arrives only while the agent experiments on a research path in form of a good or a bad signal while in the absence of signal arrivals beliefs remain constant. Hence, without an agency problem, there is no stopping until obtaining a good signal or sufficiently many bad signals that the benefit of experimentation falls below its cost. In the presence of agency, although a bad signal leads to a lower belief about the project quality, it is a result of the agent's effort. This information structure leads to new contractual features, mainly the use of extended deadlines upon bad signal disclosures.

It is common in the dynamic contracting literature with moral hazard to provide incentives through affecting the agent's continuation value. Sannikov (2008) considered a continuous time principal agent model and finds in the optimal contract the agent's continuation value has an upward drift as wages are back-loaded. In my setting, payments are back-loaded by increasing the agent's continuation value after a bad signal disclosure through extended deadline and payments are only made upon first-best stopping.

Green and Taylor (2016) study contracting for a two stage project without learning and find that the optimal contract has an endogenously determined deadline. However, the extensions happen for different reasons than in our setup: the first stage has to be completed in order to complete the second stage which is final, and as the deadline approaches, the probability of completing both stages in time decreases. As the progress report is through communication and not verifiable, a hard deadline upon no progress would mean the agent would have incentives to lie and report progress. Hence, randomization after a soft deadline that is not met is optimal due to the convex nature of the probability of two successive breakthroughs. In my setup, one good signal arrival is sufficient to conclude that the project is good, and the bad signals only make an eventual success less likely. Nonetheless, an extension happens upon their disclosure as they are a signal of agent's effort.

Manso (2010) shows in a two period model that motivating an agent to innovate may require

[4]Other settings where this is the assumption include Keller and Rady (2010), Halac, Kartik and Liu (2016), Bonatti and Horner (2011), Moroni (2016), Guo (2016), Klein (2016)

tolerating or even rewarding early failure in a setting with no moral hazard and a tradeoff between exploration and exploitation. In our setup, moral hazard is present and the agent is rewarded through experimentation time. In addition, the problem is to choose a stopping rule for experimentation as opposed to a defined time horizon. Guo (2016) studies dynamic delegation of experimentation without monetary transfers in which learning happens only through good news or the lack of it. In an extension which allows for non-conclusive good news, there is a result on sliding deadlines. The reason for the extended deadlines is due to experimentation becoming more profitable after an intermediary good signal is realised, which is the opposite case in my setting as extensions come after bad signals which make an eventual success less likely. Moroni (2016) considers experimentation in an organization where there are two milestones the feasibility of which are initially unknown and shows that the optimal contract rewards agents for earlier completion of the first stage by giving more time on the second stage rather than making a monetary payment. This is also an example of back-loading of payments through extended time, but there is no link between the two stages whereas in my setting the extension happens when the profitability of the project goes down due to a bad signal. Ulbricht (2016) considers a model of delegated search where the agent privately samples outcomes by incurring a cost and decides when to adopt one. There is adverse selection as the agent privately observes in the beginning the distribution of outcomes. The optimal contract pays a fixed bonus if and only if a certain target is reached by the adopted outcome. As opposed to his setup, in our paper there is a commonly known project success probability and uncertainty is about how long it may take to acquire conclusive news.

Akcigit and Liu (2016) consider two firms competing for an innovation on a risky research line where when one firm reaches a dead end and switches to the safe arm, the other firm which is uninformed may keep experimenting inefficiently on the risky arm. Learning happens only upon arrival of news, and a good or bad signal are both conclusive about the research line. Indeed, our setting could be considered as one with a similar learning structure in that learning only happens upon signal arrivals. However, after a bad signal, experimentation continues on a new research line but eventual success becomes less likely.

Outside the experimentation literature, there is setups with moral hazard where bad news is rewarded. For example, Levitt and Snyder (1997) consider a static setting in which an agent's effort creates a private signal about the project's potential success. Unlike in our setting, rewarding for bad signal is necessary in order to make sure the agent is truthful so that the principal can avoid investing in the project, in other words the agent would not reveal bad news without a sufficient reward. Lastly, Chade and Kovrijnykh (2016) consider a setting in which an agent is hired to evaluate different options over time and her private effort determines the precision of a signal that is publicly observed. Rewarding for bad news occurs when the prior is low and the agent's effort generates a bad signal, which is due to the signal outcome suggesting that effort was exerted. Here, the bad signal does not affect the value of the next option hence there is no learning over time. In my setting, the extended time upon bad news is not required to induce the agent to disclose the non-terminal bad signals, but it is optimal as it provides more experimentation time without increasing the agency rent.

2 Model

A principal (she) hires an agent (he) to learn about the quality of a project in order to decide whether to adopt it or not. They share a common prior ρ_0 that the project is good (equivalently the state of the world is good). A good project is profitable and has a net value v for the principal while a bad project has negative value. The principal provides resources c per unit time in order to be used by the agent for information acquisition who could also shirk and keep it for his private use.[5] Both parties are risk neutral and share a discount factor r. Outside options are zero and the agent has limited liability.

When the agent exerts effort over an interval $[t, t + dt]$ by incurring cost $a_t c dt$, a signal arrives with probability $a_t \lambda dt$ where the Poisson arrival rate λ is independent of the project type. If the project is good, with probability θ a good signal is realized (i.e. success), while for a bad project, only a bad signal (i.e. failure) can be realized. A good signal is conclusive and upon its disclosure, it is optimal for the principal to invest in the project. A bad signal decreases the belief about the project being a profitable one.[6] Signals are verifiable, they can be hidden but not constructed.

Using the Bayes' rule, upon a bad signal disclosure, the belief about project quality goes down as follows:

$$\rho_{k+1} = \frac{(1-\theta)\rho_k}{1-\theta\rho_k}$$

where ρ_k is the belief after the disclosure of k bad signals. It is important for our analysis that

[5] I have described the setting as one in which the principal provides resources for experimentation and the agent has the possibility to divert benefits to his own use. However, the setting could also be chosen as one in which the agent enjoys the benefit c from leisure when putting in low effort and remaining in the contract, without the assumption of investment by the principal. The main results would carry on.

[6] Another equivalent interpretation is to consider multiple research lines in an innovation process. This type of learning structure is present in Akcigit and Liu (2016) who consider two firms competing on 1 research line in which both good and bad signal are conclusive about the project type.

belief is not a function of the time but of the number of bad signals already realized (for the agent) and that disclosed (for the principal). These will be equivalent in the optimal incentive compatible contract.

Assumption 1. *Learning is initially optimal in the absence of an agency problem:*

$$v\lambda\theta\rho_0 > c$$

where the left hand side is the probability of a good signal arrival multiplied by the value of a good project. This assumption ensures that without the agency problem the principal would be willing to invest in information at least initially.

First-best Benchmark: Consider the principal's problem in case she could carry out information acquisition herself. Given assumption 1, the first-best consists of a stopping belief level where n^* is the minimum n which satisfies:

$$v\theta\lambda\rho_{n+1} - c < 0$$

After $n^* + 1$ bad signals, it is optimal to stop. Due to the stochastic arrival of signals while effort is exerted, it is not known ex-ante the future date when sufficient learning will have taken place. Hence, without the agency problem, there is not a stopping time but a stopping level of belief.

On the other hand, in the presence of agency, as the principal commits to a deadline in order to control the moral hazard rent, the contract may end at a higher belief than in the first best. However, an important result emphasized in the paper is that the belief at which stopping before a deadline happens in the presence of agency is also the first best stopping belief.

First, let's consider the case in which the agent's decision to experiment or shirk is perfectly observed by the principal, but not the arrival of the signals. Although the agent has the option to hide a signal, there is no gain in doing so as his expected surplus is always zero when the principal can monitor his effort. In addition, as the agent cannot lie about the realization of the signals, he does not get informational rents either and the principal's problem is identical to the benchmark case in which she experiments alone. This means that the private observation of signals alone does not lead to any distortion in the principal's problem compared to first best.

For learning to be profitable in presence of agency at least in the beginning, we should have:

Assumption 2. $v\lambda\rho_0 > 2c$

which we will assume is satisfied. The term $\lambda\rho_0\theta$ is the expected benefit of an instant of effort, and $2c$ is the total cost of an instant for the principal: the first c is the *flow cost* provided upfront by the principal to the agent and the second c is the *agency cost*, in other words the minimum continuation value per unit time the principal has to promise the agent in order to deter shirking.

3 Contracts

We consider contracts of the form $\mathcal{C} = \{W, T, a\}$ where W represents the history dependent payments to the agent upon termination and T represents the termination date. As signals are the only indicators of effort, no payment to the agent is required as long as no signal is realised. In addition, we find that no payment is necessary upon bad signal realizations which are not terminal, hence any payment will only be due to a terminal signal disclosure. Committing to a deadline ex-ante is necessary in order to control the moral hazard rent of the agent. The final term a is the effort induced in the contract, and it will always be optimal for the principal to induce effort.

As the principal wants to find the optimal contract which makes the agent work and disclose the signals without delay, her present expected value at time zero from a contract C is:

$$F_0(\mathcal{C}) = E^a\left[\int_{t=0}^{T\Lambda\gamma} e^{-rt}(a_t v\lambda\theta\rho_t\chi_t - c)dt - W_{T\Lambda\gamma}\right] \tag{1}$$

where ρ_t is the belief at time t which depends on how many bad signals have already been disclosed, and $\chi_t = 1$ only if no good signal is realised until t and $W_{T\Lambda\gamma}$ is the payment at the end of the contract. The term γ is the time when stopping happens before a deadline. Once a good signal is realized and disclosed, the profit of the project will be realized by the principal and the contract should end. The belief only evolves at certain points in time when signals are realized, hence we will denote it as a function of the number of bad signals already disclosed which will be ρ_k.

The agent's utility is a function of the payments he receives and the resources he keeps from shirking:

$$V_0(\mathcal{C}) = E^a\left[W_{T\Lambda\gamma} + \int_{t=0}^{T\Lambda\gamma} e^{-rt}[(1-a_t)c]dt\right] \tag{2}$$

A contract $\mathcal{C}$ is incentive compatible if it induces the agent to work and disclose the signals without delay.

First, in an optimal contract no payment should be made as long as no signal is disclosed, as the signals are the only hard information about the agent's effort.[7] Second, as long as the principal keeps investing, it is never optimal to induce the agent to shirk as this implies a waste of resources. While the agent shirks, no signal can arrive and the beliefs of the principal and the agent cannot diverge. A final immediate feature of optimal contracts is that if the principal stops providing resources to the agent, she will not start again, as in the absence of new information, the beliefs and hence the value of information remain constant. This leads us to conclude that the contract will have a deadline such that, if reached, experimentation will stop once and for all.[8]

Definition 1. *A termination rule is denoted by $T^k(h^t)$ which specifies the current deadline at time t when the public belief is ρ_k (k bad signals have already been revealed) and the history is h^t.*

We will call k, the number of bad signals already disclosed, as the **state**. The public history at any moment can be summarized by the times at which bad signals have been disclosed, $h^t = \{t_1, t_2, ..., t_k\}$, as these are the only elements of the history that are relevant for the contract. We find that in the optimal contract the updated deadline is indeed a function of the previous deadline and the time remaining once a bad signal is disclosed.

From now on, we will simplify the notation for the deadline as T^k, the agent's continuation value as $V_{t,k}$ and the principal's continuation value as $F_{t,k}$ at time t and state k. The payment to the agent due to the disclosure of a good or a bad signal at state k are denoted respectively by $\{w_{t,k}(G), w_{t,k}(B)\}$. As we find that the only payments will be due to terminal signals, these will equivalently be the payments at termination. While doing this, the public history h^t is not omitted, which will be important in determining the payments and the deadline. Now we rewrite the principal's continuation value in a recursive way at time t and state k:

$$(3) \qquad F_{t,k} = \int_{s=t}^{T^k} e^{-(s-t)(\lambda+r)}[\lambda(\theta\rho_k(1 - w_{s,k}(G)) + (1 - \theta\rho_k)(-w_{s,k}(B) + F_{s,k+1})) - c]ds^9$$

[7]The principal has to pay a positive rent in order to make the agent work and reveal the signals, as the agent gets a positive benefit from shirking. Then, given that the signals arrive only while the agent is working, the flow payment should be set to zero in any optimal contract and payments must be only upon the revelation of signals.

[8]This is under the assumption that it is too costly to provide the agent an infinite amount of time to experiment.

[9]where $e^{-(s-t)\lambda}$ is the probability of time s being reached without any signal arrival conditional on the agent working, hence that the belief remains at ρ_k. The term $e^{-(s-t)r}$ is the discount factor that applies when a signal arrives and is disclosed at time s. During an infinitesimal time period of dt, with probability λdt a signal arrives and is disclosed. With probability $\theta\rho_k$ the signal is a good one and upon its disclosure the contract ends with the principal making a payment $w_{t,k}(G)$, or the signal is a bad one and the state moves to $k+1$ providing the principal with the continuation value $F_{t,k+1}$, given that $k < n^*$. The detailed derivation of

Then, the problem of the principal is to choose an optimal deadline subject to satisfying the incentive compatibility conditions of the agent.

Agent's incentive constraints: The agent chooses effort $a_t \in \{0, 1\}$ where 1 indicates exerting effort and 0 the decision to shirk and keep the investment c for himself. The agent also chooses a disclosure plan $x_t \in \{G, B, 0\}$ as a function of his private history. The agent could possibly keep and disclose a signal later on or never but cannot construct a fake signal. Hence, the private history of the agent can possibly be very complicated. However, attention is restricted to contracts which induce truthful reporting in which the agent's private history h_A^t coincides with the public history, h^t. The agent has no possibility to privately learn by shirking, as no signals arrive while he shirks and the beliefs of the agent and the principal do not diverge.[10] Then, the law of motion of the agent's continuation value at time t and belief ρ_k is as follows:

$$\begin{aligned} V_{t,k} &= a_t \lambda dt[\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)(w_{t,k}(B) + (1-rdt)V_{t+dt,k+1})] \\ &\quad + (1-a_t)cdt + (1-\lambda a_t dt)(1-rdt)V_{t+dt,k}{}^{11} \end{aligned}$$

Letting dt go to 0 leads to the following equation for the agent:

$$(4)\ \ 0 = V'_{t,k} + \max_{a_t}\{-(\lambda a_t + r)V_{t,k} + a_t\lambda[\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)(w_{t,k}(B) + V_{t,k+1})] + (1-a_t)c\}$$

The first type of incentive constraint is the *no shirking constraint* which makes sure that the agent prefers experimenting to shirking at any moment, which we get by making use of equation 14.

Lemma 1. *Given a contract C the agent will choose $a_t = 1$ if and only if:*

$$\Pi_t \geq \frac{c}{\lambda} + V_{t,k} \tag{5}$$

where $\Pi_t = \theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)(w_{t,k}(B) + V_{t,k+1})$ denotes the value for the agent of a signal arrival and $V_{t,k+1}$ the continuation value after an additional bad signal.

Lemma 1 provides the local incentive constraint to work. Let us now explain briefly why this constraint is enough for global incentive compatibility. When the agent deviates to shirk, he does not get any signals and hence the beliefs of the agent and the principal cannot differ. Then,

equation (3) is provided in the Appendix.

[10] This is different from the literature on experimentation in which the agent's shirking leads to private learning and a more optimistic belief than the principal.

the future incentives to work of the agent are not modified by the deviation at time t. This means if equation (15) is satisfied for any t, it will be satisfied globally as well.

Lemma 2. *The continuation value of the agent at any t in the optimal contract is:*

$$V_{t,k} = \frac{c}{r}(1 - e^{-r(T^k - t)})$$

The optimal contract leaves the agent indifferent at any moment between exerting effort or not as the working constraint binds at any moment. Then, for a project with an initial deadline T, the total agency rent at time 0 is given by $V_{0,k} = \frac{c}{r}(1 - e^{-rT})$. Hence, the optimal contract provides the agent exactly the same amount of rent.

The second type of constraints are the *disclosure constraints* which make sure that the agent is willing to disclose the signals that he acquires without delay. The ability to keep the signals and disclose later on adds the complication of many possible histories after deviations, such as hiding a signal and shirking or experimenting in order to possibly get another one. For now I will restrict attention to local deviations which consist of delaying the disclosure of a signal and shirking for an instant.

Lemma 3. *If the local constraints are sufficient conditions for global incentive compatibility, then the disclosure constraint for a good signal is:*

$$w_{t,k}(G) \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)} w_{T^k,k}(G) \tag{6}$$

The disclosure constraint for a bad signal is:

$$w_{t,k}(B) + V_{t,k+1} \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)}(w_{T^k,k}(B) + V_{T^k,k+1}) \tag{7}$$

which mean the agent should prefer disclosing the signal now rather than hiding it and shirking until the deadline.

3.1 Benchmark case when one signal is conclusive

First, consider the case in which after one bad signal, it is optimal for the principal to stop acquiring information and to decide not to invest even in first best scenario without an agency problem. This is the case when the following holds:

Assumption 3. $v\lambda\theta\rho_1 < c$

We will show that under this assumption, the optimal contract with agency should also terminate after only one bad signal or upon reaching the deadline.

Proposition 1. *When it is optimal to stop after one bad signal in the first best, it is also optimal to do so in the second best contract with agency.*

From the above, we know that when $(\lambda\theta\rho_1 v - c) > 0$, it is not optimal to stop once the first bad signal is realised. Furthermore, we will now show that we can improve upon the contract with fixed deadline by a contract in which the deadline will be extended. To show this, we will simply focus on the case in which $(\lambda\theta\rho_2 v - c) < 0$, in other words the case in which upon a second bad signal it is optimal to stop. Later, we generalize this to cases when it is optimal to stop upon $n'th$ bad signal in the first best.

Proposition 2. *When it is not optimal in the first best to stop after the first bad signal, then in the optimal contract the deadline is extended upon the first bad signal.*

3.2 Optimal contracts in the general case

Now, let's consider the general case for a given first best stopping belief ρ_{n^*+1}. First, at any deadline T^k, it should be that $V_{T^k,k} = 0$ as the contract ends once the agent reaches the deadline without an additional signal. In addition, the disclosure constraints are irrelevant at this moment.[12] The incentive constraint to work is the only condition that should be satisfied at the deadline T^k for any k hence it binds in the optimal contract:

$$\lambda\theta\rho_k w_{t,\rho_k}(G) + (1 - \lambda\theta\rho_k) w_{t,\rho_k}(B) = \frac{c}{\lambda} - (1 - \theta\rho_k) V_{T^k,k+1} \tag{8}$$

which is found by replacing $V_{T^k,k} = 0$ in the expression in lemma (1) and making it bind.

Before we go on to state the main results, let us explain briefly why contracts having any $T^{k+1}(t_{k+1}) < T^k$, in other words whose horizon shortens upon the disclosure of bad signals, are not optimal. This is easy to see: the agent is willing to disclose non terminal bad signals without receiving any payment as long as his continuation value doesn't decrease. If the horizon of the contract shortens due to a bad signal disclosure, it must be that some payment was made to the agent. However, this payment could have instead been provided as extra time to the agent given that learning is still efficient. Or keeping the remaining time constant without any payment would be sufficient for the agent to disclose the bad signal. Now, I proceed to solve

[12] Once the agent acquires a signal, he is indifferent to disclose it or not given that the deadline is already reached. The assumption is that the agent chooses to disclose the signals whenever indifferent.

for the optimal contract under the condition $T^{k+1}(t_{k+1}) \geq T^k$ for all k and the next lemma provides the main step.

Lemma 4. *In any nonterminal state, it is optimal to set* $w_{T^k,k}(G) = w_{T^k,k}(B) = 0$ *so that the continuation value after a bad signal disclosure at the deadline is:*

$$V_{T^k,k+1} = \frac{c}{\lambda(1-\theta\rho_k)} \tag{9}$$

from the condition at the deadline in equation (30) *for* $k < n^*$*.*

In the optimal contract the payments upon signals at the deadline are set to zero and the continuation value after the disclosure of a bad signal is strictly positive. As the incentive constraint binds, the agent's rent doesn't increase due to this extension. Now, using Lemma 4, at any $t < T^k$, $w_{t,k}(G)$ and $w_{t,k}(B)$ are found by making the disclosure constraints bind, and the the remaining rent in equation need to make the incentive constraint to work bind is provided through extra time after the disclosure of bad signals.

Lemma 5. *In any non terminal state* $k < n^*$*, the payment upon the disclosure of a bad signal is zero and the continuation value increases as follows:*

$$V_{t,k+1} = V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)}$$

Now we can see that the continuation value of the agent (or the current deadline) and the public belief are sufficient to account for history dependency. The agent's effort affects the history through the signals which determine belief. At a given time, what matters for the agent's incentives are the belief and his continuation value. As the disclosure of signals already determine the agent's continuation value by determining the deadline, the only relevant part of the history at time t is the current belief ρ_k and the remaining time. We can now find the rule for the extension of the deadline. At the deadline, replacing $\frac{c}{\lambda(1-\theta\rho_k)} = \frac{c}{r}(1 - e^{-r(T^{k+1}(T^k))})$ by using the condition (9) and taking logs:

$$T^{k+1}(T^k) - T^k = \frac{ln(\frac{\lambda(1-\theta\rho_k)}{\lambda(1-\theta\rho_k)-r})}{r}$$

Doing the same for $t < T^k$ yields:

$$\frac{c}{r}(e^{-r(T^k(t)-t)} - e^{-r(T^{k+1}(t)-t)}) = \frac{c}{\lambda(1-\theta\rho_k)} \tag{10}$$

It is easy to see that $T^{k+1}(t) - T^k$ is decreasing in k. As the belief becomes more pessimistic, there is less extended time after a bad signal. This is because acquiring a bad signal is more likely hence a lower extension in the time is enough to provide the same incentives for effort. Also, in a given state k, the extended time is decreasing in t: the further t is from T^k, less costly is the extension from a time t point of view, due to discounting. Now let's summarize the different ways through which a contract ends:

- reaching the deadline T^k in any state k
- whenever a good signal is disclosed
- upon the disclosure of a bad signal in state n^* [13]

Finally, the next proposition summarises the optimal contract.

Proposition 3. *The optimal contract has the following features:*

- $V_{t,k} = \int_t^{T^k} ce^{-r(s-t)}ds$. *The continuation value of the agent at any moment and belief is equal to the payoff he would obtain instead by shirking until the deadline* T^k.
- $w_{t,k}(B) = 0$ *for* $k < n^*$. *No payment is made upon a bad signal disclosure as long as it is not terminal.*
- $T^{k+1}(t_{k+1}) > T^k$ *for* $k \leq n^*$*: the current deadline is extended upon non terminal bad signal disclosures.*
- $w_{t,k}(G) = V_{t,k}$. *The termination payment upon a good signal in any state is equal to the continuation payoff of the agent at that moment.*

In state n^* both types of signals are terminal hence payment should be made upon the disclosure of either signal, and in expectation it should be sufficient to incentivize the agent's effort.

Lemma 6. *The payments in state* n^* *should satisfy the following:*

$$\theta\rho_{n^*}w_{t,n^*}(G) + (1 - \theta\rho_n^*)w_{t,n^*}(B) = \frac{c}{\lambda} + V_{t,n^*} \tag{11}$$

$$w_{t,n^*}(S) \geq V_{t,n^*} + e^{-r(T^{n^*}-t)}w_{T,n^*}(S) \tag{12}$$

for $S \in \{G, B\}$.

[13]this is stopping before a deadline due to a terminal bad signal disclosure.

The first equation is the incentive constraint to work and the second one is the incentive constraint to disclose the signals. As the continuation value is zero after the disclosure of either type of signal in this terminal state, the payments need to be strictly positive upon the disclosure of either signal. In addition, both should be at least equal to the continuation value of the agent at that instant.

There are an infinite number of payment pairs that may satisfy the agent's incentive conditions in state n^* in which a good or a bad signal are both terminal. However, the optimal payments in states $k < n^*$ are unique and only positive in case of a good signal disclosure. The principal does not have to promise payments upon the disclosure of bad signals which are not terminal: as long as the agent's continuation value, which is a function of the time remaining in the contract, does not go down upon disclosing a bad signal, he is willing to disclose it without any payment. On the other hand, as a good signal terminates the contract, a payment is required to incentivize its disclosure. The payment upon good signal is an increasing function of the remaining time in the contract.

The principal minimizes the payments promised upon a good signal and increases the agent's continuation value after bad signal disclosures. The agent's expected payoff from effort has three components: the payment upon the disclosure of a good signal, payment upon a bad signal and the continuation value in the more pessimistic state after the disclosure of a bad signal. The disclosure constraints give the minimum payments which make sure the agent discloses the terminal signals and the remaining rent is provided as extended contract horizon while always making the incentive constraint to work bind.

Figure 1 demonstrates the updating of deadlines, which are a function of the time of the disclosure and the current deadline. The principal prefers to reward the agent through extended time until the belief falls to ρ_{n^*+1} which is the threshold at which stopping before a deadline happens and coincides with the first best benchmark without agency. Above this belief level the value of information is positive while it is negative below this belief hence the principal prefers to stop and pay the agent rather than continue. An increase in the continuation value is not necessary in order to induce the agent to disclose a bad signal which is not terminal: as long as his continuation value doesn't decrease, the agent is willing to disclose a bad signal without a reward. The extension is optimal as it leads to more time while keeping the total agency rent constant. Indeed, there are other incentive compatible contracts with the same agency rent in which agent's continuation value doesn't increase by the same amount after bad signals and the payments promised upon good signals are higher, which implies less information acquisition for

Figure 1: Updating of T^k

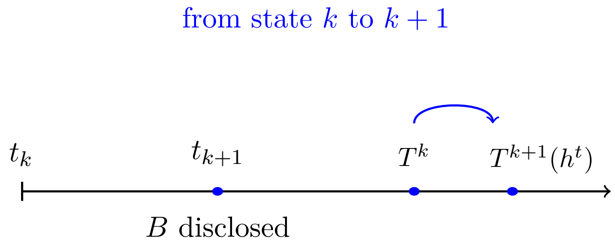

the same cost.

In a given state, the extension in the time horizon is higher the earlier a bad signal is disclosed: due to the discount factor, for the same increase in the agent's continuation value, the extension is higher the more distant T^k is from t. In addition, the extension is smaller in more pessimistic states: when a bad signal is more likely, a smaller extension provides the same expected increase in the continuation value.

Due to the deadlines, a contract may end inefficiently early at a higher belief than in the first best. This is why the extensions after the disclosure of bad signals are profitable.

The agent is promised an increased continuation value upon the disclosure of a bad signal in any state $k < n^*$ which together with the payment upon good signal disclosure makes the incentive constraint to work bind. This means the agency rent is kept constant. The term $\frac{c}{\lambda(1-\theta\rho_k)}$, which is the flow opportunity cost of effort for the agent divided by the probability of acquiring a bad signal, is the increase in the agent's continuation value provided upon disclosing a bad signal.

The disclosure of a bad signal causes the belief about the project quality to go down, but it only happens while the agent exerts effort. From the agent's point of view, the expected value of acquiring a signal is what matters for his incentives to work and how it is decomposed among the three is irrelevant. The decomposition of this rent matters for the disclosure constraints. Hence, it is optimal for the principal to minimize the payments and increasing the continuation value of the agent in the more pessimistic state reached after the disclosure of a bad signal. The principal is better off rewarding the agent through extended time as long as the intrinsic value of experimentation is positive.

Optimal stopping before a deadline: state n^*

A crucial result of the paper is that n^* is the last state in which experimentation takes place which is identical to the first best stopping.

Proposition 4. *Whenever the contract ends before a deadline, it ends efficiently either due to the disclosure of a good signal or due to the belief reaching* ρ_{n^*+1} *which is the stopping belief in the benchmark setting without agency:*

$$\lambda\theta\rho_{n^*+1} < c$$

where n^* *is the lowest value which satisfies this condition.*

Let us provide the intuition for why this belief coincides with the first best benchmark. The total cost of the contract to the principal per unit time is $2c$, where the first c is the investment per unit time promised to the agent over the remaining time horizon and the second c is the reward per unit time required to prevent the agent from shirking. By initially committing to a deadline, the principal promises the agent the equivalent of utility c over the contract horizon. Then, in order to be induced to disclose a signal that will terminate the relationship, the agent should be compensated for the value of remaining in the project until the deadline and shirking, which is equal to $\frac{c}{r}(1 - e^{-r(T^k - t)})$. The second c is due to the ex-ante flow opportunity cost of experimentation for the agent which should be promised in order to make him work at any instant. The optimal way to provide this rent is through experimentation time after bad signal disclosures while keeping constant the total agency rent as long as $\lambda\theta\rho_k \geq c$, which corresponds to the benchmark case without agency. As the probability of getting a bad signal is $\lambda(1 - \theta\rho_k)$, this leads to an increase of $\frac{c}{\lambda(1-\theta\rho_k)}$ in the agent's continuation value. Once the belief is such that $\lambda\theta\rho_k < c$, an extension is no longer profitable and the principal promises a payment to the agent upon a bad signal disclosure in order to end the contract.

3.3 The optimal T^0

Proposition 5. T^0 *which is the time initially allocated to the agent to experiment is given by:*

$$T^0 = \frac{ln(\frac{\lambda(v\theta\rho_0 + (1-\theta\rho_0)F_{T^0,1}) - (1-\theta\rho_0)c}{\theta\rho_0 c})}{\lambda} \tag{13}$$

The value $F_{T^0,1}$ is positive and less than 1 (given that the value of a project discovered as good is equal to 1), in addition it depends only on $T^1(T^0) - T^0$ which can be found from the agent's continuation value $V_{T^0,1} = \frac{c}{\lambda(1-\theta\rho_0)}$.

It is easy to see that T^0 is decreasing in c: as the cost of information increases, the principal allocates a shorter time to the agent. The sign of the derivative with respect to λ is positive, hence higher rate of information arrival leads to the allocation of more time. The derivatives with respect to θ and ρ_0 could be either positive or negative depending on the sign of $(c - \lambda F_{T^0,1})$. An increase in these parameters have 2 opposing effects: first, a higher belief about project quality implies that longer time in contract is more profitable. On the other hand, given that the possibility of good signal is higher, the agent is likely to receive a signal earlier. Hence, the sign of the derivative with respect to these parameters depend on which of the two effects dominate. The common tradeoff in these comparative statics is between the value of experimentation due to faster learning versus more agency rents that should be provided to the agent.

4 Extensions

In this section, we consider two modifications to the original model.

4.1 Public Signals

We consider the case in which the signals are publicly observed both by the agent and the principal. In this setting, the agency problem is only due to the agent's possibility to shirk. When the signals are publicly observed, the disclosure constraints are no longer relevant. This means the agent's rent is a pure moral hazard rent. As the agent cannot choose whether to disclose the signals or not, the only constraint that should be satisfied is the one which makes sure that he works.

Proposition 6. *An optimal contract in the presence of publicly observed signals has the following features:*

- $w_{t,k}(G) = w_{t,k}(B) = 0$ *for all* $k < n^*$. *The payments promised upon the realizations of either type of signal are zero as long as it doesn't happen in the terminal state* n^*.
- $V_{t,k+1} = \frac{c}{\lambda(1-\theta\rho_k)} + \frac{V_{t,k}}{(1-\theta\rho_k)}$ *for* $k < n^*$. *Incentives to the agent to exert effort are provided through increased continuation values upon the disclosure of bad signals.*
- $\theta\rho_n^* w_{t,n^*}(G) + (1-\theta\rho_n^*)w_{t,n^*}(B) = \frac{c}{\lambda} + V_{t,n^*}$. *In state* n^*, *as experimentation will end with one additional signal, the payment upon disclosure of any signal must be positive.*

In the presence of public signals, the only positive payments are made when experimentation ends in state n^* due to the arrival of a good or a bad signal. In all the previous states, due to the public observability of signals, the principal is able to set the payment even upon a success equal to zero, and incentivize the agent solely through the disclosure of bad signals and hence an extended experimentation time. Indeed, realization of a good signal at a state $k < n^*$ is unlucky for the agent as it ends the contract without providing any positive payment, but in overall his expected benefit from working is high enough that he is willing to work. The agent gets a positive payment only if a signal is realized in state n^*. The reason is that the principal no longer finds it optimal to extend the deadline upon receiving the $n^* + 1$'th bad signal, hence the only way to provide incentives to the agent is through the bonus payments upon the realization of signals. The division of this payment between good and bad signals does not matter as there is no disclosure constraints, but in expectation they have to satisfy the agent's working constraint.

Now let us compare the cases of public and privately observed signals. It is easy to see that the increase in the continuation value is higher in the public signals case. This is because the principal does not have to pay the informational rent due to the private observation of signals and can set the payments to zero for $k < n^*$. Hence, the incentives are completely back loaded to extra time in the case of public signals. In this case, the principal makes payments less often: in states $k < n^*$ he never makes any payment. Instead, the expected time in contract is longer (the jump in continuation value is higher), and the agent can only get a payment if a signal is realized in state n^*. The agent gets the payments less often, but in case he does get paid, this is usually a higher payment.

4.2 Signals are lost if not revealed right away

It is easy to see that this does not modify the results. Indeed, the possible deviations after receiving a signal are much simplified. If the agent hides a good signal, he does not find it optimal to work again. To see why this is the case: the possible actions after hiding a good signal are either to stay in the project and shirk, or to work in order to get another bad signal. However, given that now the agent knows the state is good, the probability of getting a bad signal is low enough that he does not find it profitable to work. The most profitable deviation which involves hiding the good signal is to shirk until the deadline, which provides the same minimum reward, $w_{t,k}(G) = V_{t,k}$, as in the original contract. The condition for the revelation of G is:

$$w_{t,k}(G) \geq V_{t,k}$$

which implies that the payment upon a good signal can be chosen to be the same as in the original problem. On the other hand, the agent cannot do better by hiding a bad signal either, as revelation of a bad signal increases his continuation value and does not end the relationship. Then, it is possible to set $w_{t,k}(B) = 0$ as long as $V_{t,k+1} \geq V_{t,k}$. So, the original contract remains optimal under this assumption.

5 Conclusion

This paper studied the optimal provision of incentives to an agent to acquire information about the profitability of a project. Bad news is valuable to the principal as it prevents spending resources on an unpromising project and to the agent as it reveals that he has been working. The random arrival of signals while effort is exerted implies that monitoring the agent's effort is not a trivial problem. We find that the optimal contract punishes the agent through early stopping whenever enough time passes without signals. Plus, it specifies extra time after the disclosure of bad signals as long as information about the project still has positive value. Even though a bad signal leads to a more pessimistic belief, the principal prefers providing more time as long as the value of information is still positive. However, after sufficiently many bad signals, when it also optimal to stop in the first best, the principal makes a payment to the agent and ends the contract.

Appendix

Derivation of the principal's value function

First, at $t = 0$, if the principal invests and the agent works as induced by the contract, the expected continuation value of the principal is:

$$F_{0,0} = (\lambda(\theta\rho_0(1 - w_{0,0}(G)) + (1 - \theta\rho_0)(-w_{0,0}(B) + F_{dt,1})) - c)dt + (1 - \lambda\theta dt)F_{dt,0}$$

when $dt => 0$ and replacing t instead of $t = 0$:

$$-\dot{F} + (r + \lambda)F = \lambda(\theta\rho_0(1 - w_{t,0}(G)) + (1 - \theta\rho_0)(-w_{t,0}(B) + F_{t,1})) - c$$

As T^0 is the terminating time, $F_{T^0,0} = 0$. Solving the differential equation yields:

$$F_{0,0} = \int_0^{T^0} e^{-t(\lambda+r)}[\lambda(\theta\rho_0(1 - w_{t,0}(G)) + (1 - \theta\rho_0)(-w_{t,0}(B) + F_{t,1})) - c]dt$$

At any moment t in state 0 when a bad signal arrives and is revealed, the continuation value of the principal at that moment becomes $F_{t,1}$. This gives the principal's value function.

Proof of Proposition 1

An optimal contract of this type specifies a fixed allocated time T to the agent and $w_t(S)$ which is the payment when termination happens at time t upon the disclosure of a signal S and V_t is the continuation value at time t for the agent. When a good or a bad signal are both terminal, the agent's continuation value is at any time 0 is:

$$V_{t,k} = a_t\lambda dt[E(W_{\hat{t}})] + (1 - a_t)cdt + (1 - \lambda a_t dt)(1 - rdt)V_{t+dt,k}$$

Letting dt go to 0 leads to the following equation for the agent:

$$0 = V'_{t,k} + \max_{a_t}\{-(\lambda a_t + r)V_{t,k} + a_t\lambda[E(W_t)] + (1 - a_t)c\} \tag{14}$$

By using lemma 1, given a contract C the agent will choose $a_t = 1$ if and only if:

$$E(W_{\hat{t}}) \geq \frac{c}{\lambda} + V_{t,k} \tag{15}$$

where $E(W_{\hat{t}}) = \theta\rho_k w_t(G) + (1-\theta\rho_k)w_t(B)$ is the expected payment upon a signal arrival, assuming he is willing to disclose it and that after the disclosure, the contract ends. It is optimal not to make any payment to the agent if he reaches the deadline without disclosing any signals. As we have that the incentive condition to work binds in the optimal contract, the continuation value of the agent at any t is $V_t = \frac{c}{r}(1-e^{-r(T-t)})$ and we want to find the payments that should be promised to the agent. As the agent doesn't affect the type of the signal but its arrival, the principal could also set payments upon good or bad signal identical in the optimal contract, as long as the working condition is satisfied. Hence, the working constraint gives the expected payments while the disclosure constraints provide the minimum values of each type of payment. The payoff of the principal in this type of contract is:

$$F_0 = \int_0^T e^{(-(r+\lambda)t}[\lambda(\theta\rho_0 v - \frac{c}{\lambda} - V_t) - c]dt \tag{16}$$

where $E(w_t) = \frac{c}{\lambda} + V_t$ is the expected payment to the agent upon a signal arrival when the contract ends.

Now, assume instead that the contract didn't end upon the first bad signal. As we are looking for a contradiction, let's assume instead that the contract ended upon the disclosure of the second bad signal, keeping all else constant. Let's take $w_{t,\rho_0}(G) = \frac{c}{\lambda\theta\rho_0} + V_t$ and $w_{t,\rho_0}(B) = 0$ after which the agent's continuation utility doesn't change: $V_{t,\rho_1} = V_{t,\rho_0}$. Then, check that the agent's incentive constraint to work binds:

$$\lambda(\theta\rho_0(\frac{c}{\lambda\theta\rho_0} + V_t) + (1-\theta\rho_0)V_t) = \frac{c}{\lambda} + V_{t,k} \tag{17}$$

Hence, upon the disclosure of the first bad signal, the continuation value of the agent remains constant. The principal's expected profit at time 0 is then:

$$F_0 = \int_0^T e^{-(r+\lambda)t}[\lambda(\theta\rho_0(v - \frac{c}{\lambda\theta\rho_0} - V_t) + (1-\rho_0\theta)F_{t,\rho_1}) - c]dt \tag{18}$$

Where F_{t,ρ_1} is the continuation value of the principal at the moment his belief is ρ_1. The expected payment in this state should satisfy $E(w_t) = \frac{c}{\lambda} + V_T$ as any signal in this state leads the contract to end. This means:

$$F_{t,\rho_1} = \int_t^T e^{-(r+\lambda)(s-t)}[\lambda(\theta\rho_1 v - \frac{c}{\lambda} - V_s) - c]ds \tag{19}$$

Now, let us compare the principal's payoff function in the two cases: $\Delta = (39) - (40) = \int_0^T e^{-(r+\lambda)t}\lambda[(1-\theta\rho_0)(V_s + F_{s,\rho_1})]ds$. Let us simplify F_{s,ρ_1}. All terms except V_s add up to:

$$[\frac{1-e^{-(T-s)(\lambda+r)}}{\lambda+r}](\lambda\theta\rho_1 v - 2c) \tag{20}$$

Next, we calculate $\int_t^T -e^{-(r+\lambda)(s-t)}\lambda V_s ds$ where we know V_s:

$$\int_t^T e^{-(r+\lambda)(s-t)}\frac{c}{r}(1-e^{-r(T-s)}) = \int_t^T e^{-(s-t)\lambda-(T-t)r}\frac{c}{r} = -\frac{1-e^{-(T-t)(\lambda+r)}}{\lambda+r}\frac{c\lambda}{r} - e^{-r(T-t)}[\frac{1-e^{-(T-t)\lambda}}{\lambda}]\frac{c\lambda}{r}$$

After some calculations, we get:

$$F_{t,\rho_1} = [\frac{1-e^{-(T-t)(\lambda+r)}}{\lambda+r}](\lambda\theta\rho_1 v - c) - \frac{c}{r}(1-e^{-r(T-t)}) \tag{21}$$

where the second term is equivalent to V_t, which shows that the expected payment to the agent didn't change. Then, Δ simplifies to:

$$[-\frac{1-e^{-(T-t)(\lambda+r)}}{\lambda+r}](\lambda\theta\rho_1 v - c) \tag{22}$$

which means, (39) $>$ (40) whenever $(\lambda\theta\rho_1 v - c) \leq 0$. Hence, whenever in the first best it is optimal to stop upon first bad signal, it is always optimal to do so in the second best as well, as continuing in state ρ_1 has a negative value for the principal.

Proof of Proposition 2

Consider the contract in which after one bad signal, no payment is made to the agent and the deadline doesn't change which gives the principal payoff as (40). We will show that by decreasing the payment upon good signal and keeping the incentive compatibility constraint binding, we can extend the initial deadline. The minimum payment that induces the disclosure of a good signal is:

$$w_{t,\rho_0}(G) = V_{t,\rho_0} \tag{23}$$

While the incentive constraint to work is:

$$\theta\rho_0 w_{t,\rho_0}(G) + (1-\theta\rho_0)(w_{t,\rho_0}(B) + V_{t,\rho_1}) \geq \frac{c}{\lambda} + V_{t,\rho_0} \tag{24}$$

As there is no reason to provide the agent with more continuation value than that already promised at time 0, this constraint should bind. In addition, as $w_{t,\rho_0}(G) = V_t$, we can set $w_{t,\rho_0}(B) = 0$ after which making the incentive constraint to work bind implies a jump in the continuation value after bad signal disclosure:

$$V_{t,\rho_1} = V_{t,\rho_0} + \frac{c}{\lambda(1-\theta\rho_0)} \tag{25}$$

Rather than as increased payments upon termination in state ρ_1, this jump in continuation value can be provided as an extended deadline $\hat{T} > T$, which is given by:

$$V_{t,\rho_1} = V_{t,\rho_0} + \frac{c}{\lambda(1-\theta\rho_0)} = \frac{c}{r}(1 - e^{-r(\hat{T}-t)}) \tag{26}$$

This equality ensures that the agency rent does not change as a result of this extension, as the IC condition to work still binds. Hence, upon the disclosure of the first bad signal, the continuation value of the agent increases.

The principal's expected profit from time 0 in this extended contract is then:

$$\tilde{F}_{0,\rho_0} = \int_0^T e^{-(r+\lambda)t}[\lambda(\theta\rho_0(v - V_t) + (1-\rho_0\theta)F_{t,\rho_1}) - c]dt \tag{27}$$

where F_{t,ρ_1}:

$$\tilde{F}_{t,\rho_1} = \int_t^{\hat{T}} e^{-(r+\lambda)(s-t)}[\lambda(\theta\rho_1 v - \frac{c}{\lambda} - V_s) - c]ds \tag{28}$$

which leads to:

$$[\frac{1 - e^{-(\hat{T}-t)(\lambda+r)}}{\lambda + r}](\lambda\theta\rho_1 v - c) - \frac{c}{r}(1 - e^{-r(\hat{T}-t)}) \tag{29}$$

Comparing to the contract with fixed deadline in which $w_{t,\rho_0}(G) = V_t + \frac{c}{\lambda\theta\rho_0}$ leads to the following profit:

$$F_0 = \int_0^T e^{(-r-\lambda)t}[\lambda(\theta\rho_0(v - \frac{c}{\lambda\theta\rho_0} - V_t) + (1-\rho_0\theta)F_{t,\rho_1}) - c]dt \tag{30}$$

When we call $\Delta = (30) - (27)$:

$$\int_0^T e^{-(r+\lambda)t}(\lambda(\theta\rho_0(-\frac{c}{\lambda\theta\rho_0}) + (1-\theta\rho_0)(F_{t,\rho_1} - \tilde{F}_{t,\rho_1})))dt \tag{31}$$

where $F_{t,\rho_1} - \tilde{F}_{t,\rho_1} = \left[\frac{1-e^{-(T-t)(\lambda+r)}}{\lambda+r} - \frac{1-e^{-(\hat{T}-t)(\lambda+r)}}{\lambda+r}\right](\lambda\theta\rho_1 v - c) + \frac{c}{\lambda(1-\rho_0\theta)}$. Finally, we get Δ:

$$\int_0^T e^{-(r+\lambda)t}(1-\theta\rho_0)\left[\frac{1-e^{-(T-t)(\lambda+r)}}{\lambda+r} - \frac{1-e^{-(\hat{T}-t)(\lambda+r)}}{\lambda+r}\right](\lambda\theta\rho_1 v - c)dt < 0 \tag{32}$$

given that $\hat{T} > T$. Hence, the total cost of the contract hasn't changed, and this extended contract improves on the previous one with fixed deadline, as $\frac{1-e^{-(\hat{T}-t)(\lambda+r)}}{\lambda+r}](\lambda\theta\rho_1 v - c) > [\frac{1-e^{-(T-t)(\lambda+r)}}{\lambda+r}](\lambda\theta\rho_1 v - c)$.

The intuition is simple: first, the extended deadline transfers the expected payment c from $w_{t,\rho_0}(G)$ into the continuation value V_{t,ρ_1} and transforms it into extra time. Hence, the expected cost doesn't increase. Plus, the principal obtains $(\lambda\theta\rho_1 v - c)$ which is strictly positive for a longer time.

Proof of Lemma 1

If the agent shirks for an infinitesimal time period of dt, no signal arrives and the index k does not change. He enjoys cdt and obtains the continuation value $V_{t+dt,k}$. The agent prefers to work rather than shirk if and only if:

(33)

$$\lambda[\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)(w_{t,k}(B) + V_{t,k+1})]dt + (1-\lambda dt)(1-rdt)V_{t+dt,k} \geq cdt + (1-rdt)V_{t+dt,k}$$

which, after letting dt go to 0, leads to:

$$\lambda[\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)(w_{t,k}(B) + V_{t,k+1})] \geq c + \lambda V_{t,k} \tag{34}$$

dividing everything by λ provides the result.

Proof of lemma 2

The condition the contract should satisfy in order to make sure that the agent prefers to work rather than shirk for an infinitesimal time period dt is:

$$V_{t,k} \geq cdt + (1-rdt)V_{t+dt,k}$$

Letting dt go to 0:

$$-V'_{t,k} + rV_{t,k} \geq c \tag{35}$$

In state k, as the contract will end as soon as the deadline T^k is reached (hence the $k+1$'th bad signal has not been disclosed), $V_{T^k,k} = 0$. Then, solving the differential equation $-V'_{t,k} + rV_{t,k} \geq c$ by making use of the boundary condition $V_{T^k,k} = 0$ we get:

$$V_{t,k} \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) \tag{36}$$

Hence, equation (36) gives the minimum continuation value that should be provided to the agent in order to make him work. As the principal has no incentive to provide more continuation value to the agent than needed, this equation will bind.

Proof of Lemma 3

First, I derive equation which makes sure that the agent is willing to reveal a good signal upon receiving it instead of delaying revelation to $t + dt$ and shirking in the meantime, which is:

$$w_{t,k}(G) \geq cdt + (1 - rdt)w_{t+dt,k}(G)$$

which as dt goes to 0, leads to:

$$-w'_{t,k}(G) + rw_{t,k}(G) \geq c \tag{37}$$

Next, I will derive the equation for bad signal. The constraint which makes sure that it is not profitable to wait for an infinitesimal time dt before revealing a bad signal is:

$$w_{t,k}(B) + V_{t,k+1} \geq cdt + (1 - rdt)(w_{t+dt,k}(B) + V_{t+dt,k+1})$$

Letting dt go to 0:

$$-w'_{t,k}(B) - V'_{t,k+1} + r(w_{t,k}(B) + V_{t,k+1}) \geq c \tag{38}$$

Finally, *disclosure constraints* which make sure that the agent is not willing to delay disclosing the signals instantaneously are:

$$w'_{t,k}(G) \leq rw_{t,k}(G) - c \tag{39}$$

$$w'_{t,k}(B) \leq r(w_{t,k}(B) + V_{t,k+1}) - V'_{t,k+1} - c \tag{40}$$

where w' and V' denote derivatives with respect to t. Equation (39) makes sure that the agent does not want to delay disclosing a good signal, and equation (40) makes sure that the agent does not want to delay disclosing a bad signal. If the local constraints are sufficient conditions for global incentive compatibility, then it is sufficient to check that these two conditions are satisfied at each t. Now, assuming this is the case, I solve the differential equation (39) using the boundary condition at T^k:

$$w_{t,k}(G) \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)} w_{T^k,k}(G) \tag{41}$$

which says that the agent should be better off disclosing the signal now rather than shirking until the deadline and revealing it then. Doing the same for the bad signal revelation by integrating equation (40) and using the condition at the deadline T^k:

$$w_{t,k}(B) + V_{t,k+1} \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)}(w_{T^k,k}(B) + V_{T^k,k+1}) \tag{42}$$

Proof of Lemma 4

I will prove that $V_{T^k,k+1} = \frac{c}{\lambda(1-\theta\rho_k)}$ and hence $w_{T^k,k}(B) = w_{T^k,k}(G) = 0$ in the optimal contract. In order to show this, I will show that $V_{t,k+1} = V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)}$ for any t. This will be done in 2 steps.

Step 1: First, let us show that $V_{t,k+1} \leq V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)}$. Assume the contrary, $V_{t,k+1} = V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)} + \Delta$. The incentive constraint at t is:

$$\theta\rho_k w_{t,k}(G) + (1 - \theta\rho_k)(w_{t,k}(B) + V_{t,k+1}) = \frac{c}{\lambda} + V_{t,k} \tag{43}$$

Then we have: $\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)w_{t,k}(B) < \theta\rho_k V_{t,k}$. This means, even when $w_{t,k}(B) = 0$, $w_{t,k}(G) < V_{t,k}$ hence the disclosure constraint is violated. Then, it cannot be possible that $V_{t,k+1} > V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)}$.

Step 2:

Now I will show that $V_{t,k+1} \geq V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)}$. Assume the contrary, that $V_{t,k+1} = V_{t,k} + \frac{c}{\lambda(1-\theta\rho_k)} - \Delta$. Look at the incentive constraint:

$$\theta\rho_k w_{t,k}(G) + (1 - \theta\rho_k)(w_{t,k}(B) + V_{t,k+1}) = \frac{c}{\lambda} + V_{t,k}$$

For this to hold, we have $\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)w_{t,k}(B) = \theta\rho_k V_{t,k} + (1-\theta\rho_k)\Delta$. This means

$w_{t,k}(G) > V_{t,k}$ or $w_{t,k}(B) > 0$. Hence, at least one disclosure constraint is now slack. Then, consider increasing $V_{t,k+1}$ by Δ while keeping the constraint binding, hence decreasing the expected payment. The right hand side of the incentive constraint decreases by $\Delta(1-\theta\rho_k)$, implying the expected payments at t decreases by $\lambda\Delta(1-\theta\rho_k)$. As the payments are decreased by the same amount for any t, the revelation constraints are still satisfied. I will show that the decrease in payments at time t is exactly equal to the expected agency rent during the extended horizon, and hence the net effect of this extended time period is positive given that the value of experimentation in state $k+1$ is positive. Then, I will conclude that this modification at any t increases profits.

Before the extension of time at T, the profit during a period $(t, T^{k+1}(t))$ when the state is $k+1$ is:

$$\int_t^{T^{k+1}(t)} e^{-(s-t)(\lambda+r)}[\lambda(\theta\rho_{k+1}(1-\frac{c}{r}(1-e^{-r(T^{k+1}-t)}))+(1-\theta\rho_{k+1})F_{s,k+2})-c]ds$$

which is equal to:

$$(44)\quad (1-e^{-(T^{k+1}-t)(\lambda+r)})\frac{[\lambda\theta\rho_{k+1}]}{\lambda+r}+\int_t^{T^{k+1}(t)} e^{-(s-t)(\lambda+r)}(1-\theta\rho_{k+1})F_{s,k+2}ds-\frac{c}{r}(1-e^{-r(T^{k+1}-t)})$$

where $\frac{c}{r}(1-e^{-r(T^{k+1}-t)}) = V_{t,k+1}$. After increasing $V_{t,k+1}$ by Δ, $\hat{T}^{k+1}$ is such that $\frac{c}{r}(1-e^{-r(\hat{T}^{k+1}-t)}) = V_{t,k+1}+\Delta$.

Then, the principal's profit during the extended horizon is:

$$(45)\quad (1-e^{-(\hat{T}^{k+1}-t)(\lambda+r)})\frac{\lambda\theta\rho_{k+1}}{\lambda+r}+\int_t^{\hat{T}^{k+1}(t)} e^{-(s-t)(\lambda+r)}(1-\theta\rho_{k+1})\hat{F}_{s,k+2}ds-(V_{t,k+1}+\Delta)$$

the increase in cost at t, $\Delta(1-\theta\rho_k)$, is equal to the expected payment to the agent during the extended horizon. Finally, we show that the expected profit has increased. $(1-e^{-(\hat{T}^{k+1}-t)(\lambda+r)})\frac{[\lambda\theta\rho_{k+1}+(1-\theta\rho_{k+1})]}{\lambda+r}$ increases in T^{k+1}, hence the first term has increased as $\hat{T}^{k+1} > T^{k+1}$. Then, $\int_t^{\hat{T}^{k+1}(t)} e^{-(s-t)(\lambda+r)}(1-\theta\rho_{k+1})\hat{F}_{s,k+2}ds > \int_t^{T^{k+1}(t)} e^{-(s-t)(\lambda+r)}(1-\theta\rho_{k+1})F_{s,k+2}ds$, because $\hat{T}^{k+1} > T^{k+1}$, and $\hat{F}_{s,k+2} > F_{s,k+2}$.

Proof of Lemma 5

Proof. I make the revelation constraint for good signal in equation (41) bind in order to find the minimum $w_{t,k}(G)$:

$$w_{t,k}(G) = \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)} w_{T^k,k}(G) \tag{46}$$

where $w_{T^k,k}(G) = 0$, hence $w_{t,k}(G) = \frac{c}{r}(1 - e^{-r(T^k - t)})$.

Given that the incentive constraint for working binds for any (t, k), the agent's continuation value is given by:

$$V_{t,k} = \frac{c}{r}(1 - e^{-r(T^k - t)})$$

which implies: $w_{t,k}(G) = V_{t,k}$. Then, the working constraint given in lemma 1 becomes:

$$(1 - \theta\rho_k)(w_{t,k}(B) + V_{t,k+1}) \geq \frac{c}{\lambda} + (1 - \theta\rho_k)V_{t,k}$$

It is optimal that this constraint binds. Now I claim that it is optimal to set $w_{t,k}(B) = 0$ and increase the continuation value as experimentation is still profitable when the belief is ρ_{k+1}, which implies:

$$V_{t,k+1} = V_{t,k} + \frac{c}{\lambda(1 - \theta\rho_k)}$$

To see that $w_{t,k}(B) = 0$ satisfies the disclosure constraint, replace $w_{T^k,k}(B) = 0$ in the revelation constraint (42):

$$w_{t,k}(B) + V_{t,k+1} \geq V_{t,k} + e^{-r(T^k - t)} V_{T^k,k+1}$$

after replacing $V_{T^k,k+1} = \frac{c}{\lambda(1-\theta\rho_k)}$:

$$w_{t,k}(B) \geq V_{t,k} - V_{t,k+1} + (1 - e^{-r(T^k - t)}) \frac{c}{\lambda(1 - \theta\rho_k)} \tag{47}$$

As $V_{t,k+1} - V_{t,k} = \frac{c}{\lambda(1-\theta\rho_k)}$, the right hand side is negative, which means this constraint is slack. It then concludes that it is optimal to set $w_{t,k}(B) = 0$ for any t and $k \leq n^*$.

□

Proof of Lemma 6

Consider the incentive constraint right before the deadline T^{n^*} which is the last moment of experimentation. By replacing $V_{T^{n^*},n^*} = V_{T^{n^*},n^*+1} = 0$, the no shirking constraint in (15) simplifies to:

$$\theta\rho_n^* w_{T^{n^*},n^*}(G) + (1-\theta\rho_{n^*})w_{T^{n^*},n^*}(B) \geq \frac{c}{\lambda} \tag{48}$$

As the revelation constraints are irrelevant at the deadline T^{n^*}, the principal will make the incentive constraint (48) bind. In case this constraint were slack, the payments could be decreased without modifying the agent's incentives and the principal's profits would have increased. Then, for $t < T^{n^*}$:

$$\theta\rho_{n^*} w_{t,n^*}(G) + (1-\theta\rho_n^*)w_{t,n^*}(B) \geq \frac{c}{\lambda} + V_{t,n^*} \tag{49}$$

after replacing $V_{t,n^*+1} = 0$ in the incentive constraint given by lemma 1. The first term, $\frac{c}{\lambda}$, represents the compensation for the instantaneous flow cost of working for the agent, and V_{t,n^*} is the future payoff foregone after revealing a signal that leads to project termination. Before concluding that equation (49) binds, it is necessary to check the *disclosure constraints*. Multiplying the constraint for the revelation of G given in equation (41) and the constraint for B given in equation (42) respectively by their probabilities $\theta\rho_n^*$ and $1-\theta\rho_n^*$ gives:

$$\begin{aligned} \theta\rho_n^* w_{t,n^*}(G) + (1-\theta\rho_n^*)w_{t,n^*}(B) \geq \frac{c}{r}(1-e^{-r(T^{n^*}-t)}) + \\ e^{-r(T^{n^*}-t)}(w_{T,n^*}(S)\theta\rho_n^* w_{T^{n^*},n^*}(G) + (1-\theta\rho_n^*)w_{T^{n^*},n^*}(B)) \end{aligned} \tag{50}$$

It is easy to conclude that equation (50) is slack when the constraint in equation (49) binds, given that $V_{t,n^*} = \frac{c}{r}(1-e^{-r(T^{n^*}-t)})$.

Finally, the payments can be set to any value which make the no shirking constraint (49) bind and satisfy the *disclosure constraints*. Then, the payments in state n^* in the optimal contract are as given by lemma (6).

Proof of Proposition 4

Initially I assume that there is a belief ρ_n^* at which experimentation ends upon the revelation of an additional signal B, then will show that this condition is indeed independent of the calendar time t and the history. The principal's profit at t when the belief is ρ_n^* is:

$$F_{t,n^*} = \int_0^{T^{n^*}} e^{-t(\lambda+r)}(\lambda(\theta\rho_n^* - w_{t,n^*}) - c)dt \tag{51}$$

in case experimentation ends at the n^*+1'th good signal where $w_{t,n^*} = \theta\rho_n^* w_{t,n^*}(G) + (1 - \theta\rho_n^*) w_{t,n^*}(B) = \frac{c}{\lambda} + \frac{c}{r}(1 - e^{-r(T^{n^*}-t)})$ and n^* is the last state. If the contract does not end at n^*+1'th bad signal, then it will end at the bad signal n^*+2. In that case, F_{t,n^*} becomes:

$$F_{t,n^*} = \int_0^{T^{n^*}} e^{-t(\lambda+r)}[\lambda(\theta\rho_n^*(1 - w_{t,n^*}(G)) + (1-\theta\rho_{n^*})F_{t,n^*+1}) - c]dt \tag{52}$$

where $w_{t,n^*}(G) = \frac{c}{r}(1 - e^{-r(T^{n^*}-t)}) = V_{t,n^*}$. Now, the condition for 51 > 52 is:

$$\int_0^{T^{n^*}} -c - \lambda V_{t,n^*} dt \geq \int_0^{T^{n^*}} -\lambda\theta\rho_{n^*} V_{t,n^*} + \lambda(1-\theta\rho_{n^*})F_{t,n^*+1} dt$$

which, after replacing the payments and integrating, simplifies to:

$$-\frac{c}{\lambda(1-\theta\rho_{n^*})} - V_{t,n^*} \geq F_{t,n^*+1}$$

if this holds, then it is optimal to end experimentation at the n^*+1'st bad signal at any t. Replacing $V_{t,n^*+1} = \frac{c}{\lambda(1-\theta\rho_{n^*})} + V_{t,n^*}$:

$$-V_{t,n^*+1} \geq F_{t,n^*+1} \tag{53}$$

Let us calculate the right hand side:

$$F_{t,n^*+1} = \int_t^{T^{n^*+1}} e^{-(s-t)(\lambda+r)}[\lambda(\theta\rho_{n^*+1} - \frac{c}{\lambda} - w_{s,n^*+1}) - c]ds$$

where $\frac{c}{r}(1 - e^{-r(T^{n^*+1}-t)}) = w_{t,n^*+1}$. Integrating this expression:

$$\frac{(1 - e^{-(T^{n^*+1}-t)(\lambda+r)})}{\lambda+r}(\lambda\theta\rho_{n^*+1} - c) - \frac{c}{r}(1 - e^{-r(T^{n^*+1}-t)})$$

hence the condition (53) holds if and only if:

$$\frac{(1 - e^{-(T^{n^*+1}-t)(\lambda+r)})}{\lambda+r}(\lambda\theta\rho_{n^*+1} - c) \leq 0$$

Finally, the following is the stopping belief:

$$\rho_{n^*+1} \leq \frac{c}{\lambda\theta}$$

Proof of proposition 5

Let us write down the principal's problem:

$$F_{0,0} = \int_0^{T^0} e^{-t(\lambda+r)}[\lambda(\theta\rho_0(1 - w_{t,0}(G)) + (1 - \theta\rho_0)F_{t,1}) - c]dt$$

where $w_{t,k}(G) = \frac{c}{r}(1 - e^{-r(T^0-t)})$.

$$F_{t,1} = \int_t^{T^1(t)} e^{-(s-t)(\lambda+r)}[\lambda(\theta\rho_1(1 - w_{s,1}(G)) + (1 - \theta\rho_1)F_{s,2}) - c]ds$$

and it continues for all k until $k = n^*$. Then the derivative of $F_{0,0}$ with respect to T_0 is:

$$e^{-rT^0}[e^{-T^0\lambda}(\lambda(\theta\rho_0 - (1 - \theta\rho_0)(c + F_{T^0,1})) - \theta\rho_0 c)] + (1 - e^{-T^0(\lambda+r)})\frac{\lambda(1 - \theta\rho_0)F'_{T^0,1}}{\lambda + r}$$

Here, $F'_{T^0,1} = 0$, which is the derivative of $F_{T^0,1}$ with respect to T^0 as $F_{T^0,1}$ does not depend on T^0. Then, after rearranging we have:

$$e^{-T^0(\lambda+r)}[\lambda(\theta\rho_0 + (1 - \theta\rho_0)F_{T^0,1}) - c] - \theta\rho_0 c e^{-rT^0}(1 - e^{-T^0\lambda})$$

where the first term denotes the marginal benefit from extending experimentation for an instant at T^0, and the second term denotes the cost of increasing experimentation time due to increased payments that should be promised in all the previous periods. After rearranging:

$$e^{-rT^0}[e^{-T^0\lambda}(\lambda(\theta\rho_0 + (1 - \theta\rho_0)F_{T^0,1}) - (1 - \theta\rho_0)c) - \theta\rho_0 c]$$

second derivative:

$$-(\lambda + r)e^{-T^0(r+\lambda)}[\lambda\theta\rho_0 + \lambda(1 - \theta\rho_0)F_{T^0,1} - (1 - \theta\rho_0)c] + \theta\rho_0 e^{-rT^0} rc$$

The first derivative has a single root, and it can be verified that at this point, the second derivative is negative which means it is a local maximum. As the first order condition has no other root, this function has only one reflection point, hence T^0 is indeed a global maximum. The optimal T^0 is then found as:

$$T^0 = \frac{ln(\frac{\lambda(\theta\rho_0 + (1-\theta\rho_0)F_{T^0,1}) - (1-\theta\rho_0)c}{\theta\rho_0 c})}{\lambda} \tag{54}$$

It can be checked that the second derivative is negative at $T = 0$ and at the optimal T^0, which implies that the value function of the principal is not convex in any region until T^0. This also proves that randomizing on the stopping time cannot be optimal for the principal, justifying the initial restriction to deterministic deadlines. In addition, the value function is decreasing for $t \geq T^0$. The second derivative becomes positive as T goes to ∞. However, as there is no other point at which the first derivative is zero and the second derivative is negative for $T > T^0$, I conclude that this value can never go above T^0.

Proof of Proposition (6)

The incentive constraint which makes sure that the agent works is identical to the one with private signals:

$$\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k)w_{t,k}(B) = \frac{c}{\lambda} + V_{t,k} - (1-\theta\rho_k)V_{t,k+1}$$

The difference is that there are no more disclosure constraints, which implies that the above condition will bind in the optimal contract. In addition, it is still optimal for the principal to extend the horizon of experimentation upon the revelation of bad signals, and even more this time by setting $w_{t,k}(G) = w_{t,k}(B) = 0$ for $k \leq n^*$, which leads to:

$$V_{t,k+1} = \frac{c}{\lambda(1-\theta\rho_k)} + \frac{V_{t,k}}{(1-\theta\rho_k)}$$

and in state n^*, as $V_{t,n^*+1} = 0$:

$$\theta\rho_n^* w_{t,n^*}(G) + (1-\theta\rho_n^*)w_{t,n^*}(B) = \frac{c}{\lambda} + V_{t,n^*}$$

In the optimal contract $w_{t,n^*}(G)$ and $w_{t,n^*}(B)$ can take on any values that satisfy this equation.

Sufficiency of local disclosure constraints:

This section verifies that the local no deviation constraints for disclosure in lemma 1 are indeed sufficient to account for global incentive compatibility. Let us check that there is no profitable deviation for the agent after receiving a bad signal, such as hiding one or more signals. The agent should be compensated at least for the change in his continuation value upon disclosing a bad signal when the state moves from k to $k+1$:

$$w_{t,k}(B) \geq max[0, V^B_{t,k} - V_{t,k+1}] \tag{55}$$

where $V^B_{t,k}$ is the continuation value of the agent after hiding a bad signal. Now, let us find what would be the continuation value of the agent after hiding this signal. After the deviation, the gain in the continuation value of the agent upon the disclosure of the next bad signals will be respectively $\frac{c}{\lambda(1-\theta\rho_k)}$ and $\frac{c}{\lambda(1-\theta\rho_{k+1})}$ which are independent of when they are disclosed. Also, the reward upon the disclosure of a good signal is always higher for higher k: $w_{t,k}(G) < w_{t,k+1}(G)$, which means hiding a bad signal can only decrease the payment from a good signal disclosure. Then, hiding and delaying the disclosure of a bad signal is weakly dominated by a strategy in which the agent reveals the bad signal upon receiving it and obtains the increase in his continuation value earlier. Then, equation (55) is satisfied.

Finally, the deviation to hide a good signal and continue experimenting in order to get another bad signal cannot be profitable either. This deviation could only be attractive if $w_{t,k+1}(G)$ were sufficiently high compared to $w_{t,k}(G)$. While the agent works, he only gets an increased continuation value when he receives a bad signal. We know that the agent's IC constraint is always binding. Then, once the agent hides a good signal, his belief about a bad signal arrival is lower than what the contract indicates. Hence, his expected payoff from a bad signal is lower which means he would no longer work. After deviating to hide a good signal, the agent is better off shirking rather than working. The gain from an additional bad signal is not high enough that an agent who has already acquired a good signal would be willing to hide it and experiment in order to get an additional bad signal. Finally, we conclude that the local constraints are sufficient for global incentive compatibility.

Non optimality of contracts having k such that $T^{k+1}(t_{k+1}) \leq T^k$

In this section I will verify that it is never optimal to have $T^{k+1}(t_{k+1}) \leq T^k$. First, I will solve for the optimal payment schedule under this condition. Then, by replacing the payments in the principal's objective function, I will find that the profits always increase in the deadline T^{k+1} justifying the initial restriction to contracts with $T^{k+1}(t_{k+1}) \geq T^k$.

The optimal payments and continuation values in a contract in which $T^{k+1}(t_{k+1}) \leq T^k$ are:

$$\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k) w_{t,k}(B) = \frac{c}{\lambda} + V_{t,k}$$
$$V_{t,k} = V_{t,k+1}$$
$$\theta\rho_k w_{T^k,k}(G) + (1-\theta\rho_k) w_{T^k,k}(B) = \frac{c}{\lambda}$$
$$T^{k+1}(t_{k+1}) = T^k$$

where $V_{t,k} = \frac{c}{r}(1 - e^{-r(T^k - t)})$. It is not optimal to shorten the horizon of experimentation, hence when restricted to $T^{k+1}(t_{k+1}) \geq T^k$, it is found that $T^{k+1} = T^k$.

Now let us show this result. When the deadline is T^k and $T^{k+1} \leq T^k$, $V_{T^k,k+1} = V_{T^k,k} = 0$ and the no shirking condition (15) simplifies to:

$$\theta\rho_k w_{T^k,k}(G) + (1-\theta\rho_k) w_{T^k,k}(B) \geq \frac{c}{\lambda} \tag{56}$$

The revelation constraints are irrelevant at the deadline T^k. It is then optimal that the equation (56) binds, which provides the payments at the deadline. Then, for $t < T^k$, using $V_{T^k,k+1} = V_{T^k,k} = 0$, the revelation constraint for B becomes:

$$w_{t,k}(B) + V_{t,k+1} \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)} w_{T^k,k}(B)$$

The revelation constraint for G is:

$$w_{t,k}(G) \geq \frac{c}{r}(1 - e^{-r(T^k - t)}) + e^{-r(T^k - t)} w_{T^k,k}(G) \tag{57}$$

multiplying $w_{t,k}(G)$ and $w_{t,k}(B)$ respectively by their weights $\theta\rho_k$ and $1-\theta\rho_k$, we get:

$$\begin{aligned}\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k) w_{t,k}(B) \geq V_{t,k} + e^{-r(T^k - t)}(\theta\rho_k w_{T^k,k}(G) \\ + (1-\theta\rho_k) w_{T^k,k}(B)) - (1-\theta\rho_k) V_{t,k+1}\end{aligned} \tag{58}$$

where the right hand side is equal to $V_{t,k} + e^{-r(T^k - t)} \frac{c}{\lambda} - (1-\theta\rho_k) V_{t,k+1}$. The incentive constraint to work is:

$$\theta\rho_k w_{t,k}(G) + (1-\theta\rho_k) w_{t,k}(B) \geq \frac{c}{\lambda} + V_{t,k} - (1-\theta\rho_k) V_{t,k+1} \tag{59}$$

Comparing the equations (58) and (59), it is easy to conclude that the incentive constraint is the binding one. Then, the payments $w_{T^k,k}(G)$ and $w_{T^k,k}(B)$ should be chosen such that the incentive constraint (59) binds and the revelation constraints are satisfied. Next I will show that $T_{k+1}(t_{k+1}) < T_k$ cannot be optimal for any k by replacing the payments into the principal's value function. First, I will look at F_{0,n^*-1} (normalizing the starting time of state n^*-1 to 0) and F_{t,n^*} to show that $T^{n^*}(t_{n^*}) \geq T^{n^*-1}$. After this, I will show that it also holds for any $k < n^*$. The optimal payment schedule in state n^* was already provided for any optimal contract, given that it is the last possible state. Replacing the values $w_{t,k}(G)$ and $w_{t,k}(B)$ into the principal's problem as well as $V_{t,k} = \frac{c}{r}(1-e^{-r(T^k-t)})$ and rearranging:

$$(60) \quad F_{0,n^*-1} = \int_0^{T^{n^*-1}} e^{-t(\lambda+r)} \Big[\lambda(\theta\rho_{n^*-1} - \frac{c}{\lambda} - \theta\rho_k\frac{c}{r}(1-e^{-r(T^{n^*}-t)}) - \frac{c}{r}(e^{-r(T^{n^*}-t)} - e^{-r(T^{n^*-1}-t)}) + F_{t,n^*}) - c\Big]dt$$

In state n^*:

$$F_{t,n^*} = \int_t^{T^{n^*}(t)} e^{-(s-t)(\lambda+r)} \left[\lambda(\theta\rho_{n^*} - \frac{c}{\lambda} - \frac{c}{r}(1-e^{-r(T^{n^*}-s)})) - c\right]$$

maximizing F_{0,n^*-1} with respect to T^{n^*}:

$$\int_0^{T^{n^*-1}} e^{-t(\lambda+r)}(1-\theta\rho_{n^*-1})\Big[ce^{-r(T^{n^*}-t)} + e^{-(T^{n^*}-t)(\lambda+r)}[\lambda\theta\rho_{n^*} - c] - ce^{-r(T^{n^*}-t)}\Big]dt > 0$$

The first and the third terms cancel out, and we are left with $e^{-(T^{n^*}-t)(\lambda+r)}[\lambda\theta\rho_{n^*} - c]$. Then, as $\lambda\theta\rho_{n^*} > c$ this expression is always positive for $k < n^*$. Hence the profit of the principal is always increasing in T^{n^*} when restricted to $T^{k+1}(t_{k+1}) \leq T^k$. This implies that the optimal contract has the feature that $T^{n^*}(t) \geq T^{n^*-1}$. Finally, I need to show that a shortening of the time horizon is not optimal when the next state, $k+1$ is such that $T^{k+2}(t_{k+2}) \geq T_{k+1}$ either, in other words given that in the next state the deadline is extended upon revelation of a bad signal. I replace the payment schedule for state $k+2$:

$$F_{0,k} = \int_0^{T^k} e^{-t(\lambda+r)}[\lambda(\theta\rho_k(1 - \frac{c}{\lambda\theta\rho_k} - \frac{c}{r}(1-e^{-r(T^k-t)})) + (1-\theta\rho_k)[-\frac{c}{r}e^{rt}(e^{-rT^{k+1}} - e^{-rT^k}) + F_{t,k+1}]) - c]dt$$

where:

$$F_{t,k+1} = \int_0^{T^{k+1}} e^{-(s-t)(\lambda+r)}[\lambda(\theta\rho_{k+1}(1-\frac{c}{r}(1-e^{-r(T^{k+1}-s)}))+$$
$$(1-\theta\rho_{k+1})[F_{s,k+2}])-c]ds$$

The derivative of this whole term with respect to T^{k+1}:

$$\int_0^{T^k} e^{-t(\lambda+r)}(1-\theta\rho_k)\big[ce^{-r(T^{k+1}-t)}+e^{-(T^{k+1}-t)(\lambda+r)}$$
$$[\lambda\theta\rho_{k+1}+\lambda(1-\theta\rho_{k+1})F_{T^{k+1},k+2}-c]-\theta\rho_{k+1}c(e^{-(T^{k+1}-t)(\lambda+r)}-e^{-r(T^{k+1})+t(\lambda+r)})\big]dt$$

the terms $ce^{-r(T^{k+1}-t)}$ and $\theta\rho_{k+1}c(e^{-(T^{k+1}-t)(\lambda+r)}-e^{-r(T^{k+1})+t(\lambda+r)})$ are positive and hence the whole expression is also positive as long as $\lambda\theta\rho_{k+1}+\lambda(1-\theta\rho_{k+1})F_{T^{k+1},k+2}-c\geq 0$ which is the case as $\lambda\theta\rho_{k+1}-c\geq 0$. The final one is the condition for experimentation to be profitable initially. I then conclude that it is always profit enhancing to increase T^{k+1} in the region when $T^{k+1}\leq T^k$. This means, $T^{k+1}(t_{k+1})=T^k$ and $V_{t,k+1}=V_{t,k}$. The optimal payment schedule follows from the constraints. Hence, there cannot be an optimal contract whose time horizon shortens after the release of a bad signal. Now I can conclude that in the optimal contract there cannot be any $T^{k+1}(t_{k+1})<T^k$, which justifies the initial restriction to contracts having $T^{k+1}(t_{k+1})\geq T^k$.

References

[Akcigit and Liu (2016)] Ufuk Akcigit and Qingmin Liu. 2016. "The Role of Information in Innovation and Competition." *Journal of the European Economic Association*, 14(4), 828-870.

[1] BBC News Technology. (2014) ''Secret Google lab 'rewards staff for failure."

[Bergemann and Hege (1998)] Dirk Bergemann and Ulrich Hege. 1998. "Venture Capital Financing, Moral Hazard, and Learning." *Journal of Banking and Finance,* 22, 700-735.

[Bergemann and Hege (2005)] Dirk Bergemann and Ulrich Hege. 2005. "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics*, 36, 719-752.

[Bolton and Harris (1999)] Patrick Bolton and Christopher Harris (1999): "Strategic Experimentation" *Econometrica.*

[Bonatti and Horner (2011)] Alessandro Bonatti and Johannes Horner (2011): "Collaborating." *American Economic Review*, 101, 632-663.

[Bonatti and Horner (2016)] Alessandro Bonatti and Johannes Horner (2016): "Career Concerns with Exponential Learning." *Theoretical Economics.*

[Chade and Kovrijnykh (2016)] Hector Chade and Natalia Kovrijnykh (2016): "Information Acquisition, Moral Hazard, and Rewarding for Bad News". *Journal of Economic Theory.*

[Demarzo and Sannikov (2016)] Peter DeMarzo and Yuliy Sannikov (2016):"Learning, Termination and Payout Policy in Dynamic Incentive Contracts.", Working paper.

[Gerardi and Maestri (2012)] Dino Gerardi and Lucas Maestri (2012): "A Principal-Agent Model of Sequential Testing." *Theoretical Economics*, 7, 425-463.

[Gomes, Gottlieb and Maestri (2016)] Renato Gomes, Daniel Gottlieb and Lucas Maestri (2016): "Experimentation and Project Selection: Screening and Learning." *Games and Economic Behavior*,96, 145-169.

[Green and Taylor (2016)] Brett Green and Curtis R.Taylor (2016): "Breakthroughs, Deadlines and Severance: Contracting for Multistage Projects". *American Economic Review.*

[Guo (2016)] Yingni Guo (2016): "Dynamic Delegation of Experimentation". *American Economic Review*, 106(8).

[2] Harvard Business Review. “Strategies for Learning from Failure”, 2011. April.

[Halac, Kartik and Liu (2016)] Marina Halac, Navin Kartik, and Qingmin Liu (2016): “Optimal Contracts for Experimentation” *Review of Economic Studies*, 83, 1040-1091.

[Horner and Samuelson (2013)] Johannes Horner and Larry Samuelson (2013): “Incentives for Experimenting Agents.” *RAND Journal of Economics*, 4,632-663.

[Keller and Rady (2010)] Godfrey Keller and Sven Rady (2010): “Strategic Experimentation with Poisson Bandits.” *Theoretical Economics*, 5: 275-311.

[Keller, Rady and Cripps (2005)] Godfrey Keller, Sven Rady and Martin Cripps (2005). “Strategic Experimentation with Exponential Bandits.” *Econometrica*, 73: 39-68.

[Klein (2016)] Nicolas Klein. (2016):“The Importance of Being Honest.” *Theoretical Economics*, 11,773-811.

[3] Suehyun Kwon (2014): “Dynamic Moral Hazard with Persistent States.” Working paper.

[4] Jean-Jacques Laffont, and Jean Tirole (1988): “The Dynamics of Incentive Contracts.” *Econometrica*, 56,1153-1175.

[5] Tracy R. Lewis. (2011): “A Theory of Delegated Search for the Best Alternative.” *RAND Journal of Economics.*

[6] Gustavo Manso (2011): “Motivating Innovation.” *Journal of Finance*, 66, 1823-1860.

[7] Robin Mason and Juuso Valimaki. (2011): “Dynamic Moral Hazard and Stopping”. Working Paper.

[8] Sofia Moroni. (2016): “Experimentation in Organizations”. Working paper.

[9] Pauli Murto and Juuso Valimaki. (2011): “Learning and Information Aggregation in an Exit Game”. *Review of Economic Studies.*

[10] Yuliy Sannikov (2008): “A Continuous-Time Version of the Principal-Agent Problem.” *The Review of Economic Studies*, 75, 957-984.

[11] Stephen E. Spear and Sanjay Srivastava. (1987): “On Repeated Moral Hazard with Discounting.” *The Review of Economic Studies*, 54, 599.

[Ulbricht (2016)] Robert Ulbricht. (2016): "Optimal Delegated Search with Adverse Selection and Moral Hazard." *Theoretical Economics*, 253-278.