Contracts that Reward Innovation: Delegated Experimentation with an Informed Principal

Yiman Sun*

August 1, 2018

Abstract

We examine the nature of contracts that optimally reward innovations in a risky environment, when the innovator is privately informed about the quality of her innovation and must engage an agent to develop it. We model the innovator as a principal who has private but imperfect information about the quality of her project: the project might be worth exploring or not, but even a project of high quality may fail. We characterize the best equilibrium for the high type principal, which is either a separating equilibrium or a pooling one. Due to the interaction between the signaling incentives of the principal and dynamic moral hazard of the agent, the best equilibrium induces inefficiently early termination of the high quality project. The high type principal is forced to share the surplus – with the agent in the separating equilibrium, or the low type principal in the pooling equilibrium. A mediator, who offers a menu of contracts and keeps the agent uncertain about which contract will be implemented, can increase the payoff of the high type principal to approximate her full information surplus.

Keywords: Contracts, Innovation, Experimentation, Informed Principal, Dynamic Moral Hazard, Signaling Game, Mechanism Design

JEL Codes: D82, D83, D86

*Department of Economics, University of Texas at Austin. Email: yiman@utexas.edu
I am very grateful to my advisor, V. Bhaskar, for his invaluable guidance and support. I am also deeply indebted to Caroline Thomas for her insightful comments and encouragement. I would like to thank Svetlana Boyarchenko, Yi Chen, Joyee Deb, Laura Doval, Johannes Hörner, Erik Madsen, Tymofiy Mylovanov, Larry Samuelson, Vasiliki Skreta, Maxwell Stinchcombe, Tatuo Sugaya, Robert Town, Thomas Wiseman, and other seminar participants in UT, Yale, 2017 North American Summer Meeting of the Econometric Society, the 28th Stony Brook International Conference on Game Theory, 2017 Texas Economic Theory Camp, and 2017 Midwest Economic Theory Conference for their excellent suggestions and helpful discussions. I thank Cowles Foundation and the Department of Economics at Yale for their hospitality.

1 Introduction

How should we design contracts in order to reward innovation, when poor ideas can masquerade as worthwhile ones? Consider an innovator who has a new idea, and has private but imperfect information about the quality of the idea. Specifically, the idea may be worth exploring, but even high quality ideas may result in failure, just as ideas that are of low quality ex ante may sometimes be successful. The innovator needs to engage an agent to explore the idea, and thus the agent's moral hazard must also be confronted. Since the agent becomes more pessimistic over time about the probability of success in the absence of a breakthrough, moral hazard is dynamic. Moreover, the innovator also needs to convince the agent that the idea's quality is high and it is worth exploring. Our question is: what are the contracts that provide maximal rewards for high quality innovations? If high quality innovations are not rewarded properly, then innovators would not invest in better ideas in the first place. The question is particularly important in the knowledge economy, where economic growth is driven by innovations. Furthermore, all the ingredients listed are important considerations in knowledge industries: innovators have private information, but even the best ones are not omniscient, and sometimes come up with unworkable ideas; the innovators themselves may not be the best ones to undertake project development, and may have to delegate the task to specialized bodies; only these specialized bodies know how intensively they are working on the innovator's project.

We analyze a model where a privately informed innovator engages an agent to work on a project. A senior professor may need to hire a junior research assistant to conduct some lab experiments for a research project; a tech startup founder may need to hire a professional team to validate the needs of the new product. In drug development, scientists in biotech companies may have strong insights into the fundamental mechanism of diseases. Other entities, like Contract Research Organizations (CRO) and Contract Manufacturing Organizations (CMO) may be specialized in other research support services.[1] When a biotech company hires a CRO or a CMO to undertake some experiments for a drug development project, the biotech company needs to not only incentivize them to exert efforts towards

[1]According to Wikipedia, the services provided by a CRO may include biopharmaceutical development, biologic assay development, commercialization, preclinical research, clinical research, clinical trials management, and pharmacovigilance. The services provided by a CMO may include pre-formulation, formulation development, stability studies, method development, pre-clinical and Phase I clinical trial materials, late-stage clinical trial materials, formal stability, scale-up, registration batches and commercial production. See `https://en.wikipedia.org/wiki/Contract_research_organization` and `https://en.wikipedia.org/wiki/Contract_manufacturing_organization`.

exploring the viability of the project, but also convince them of the quality of the project. The problem is particularly severe when high quality projects are relatively scarce, and cannot be distinguished from low quality projects. This is typically the case in research industries.[2]

This paper has two parts. In the first part, we combine a signaling game with an exponential bandit model to study the dynamic agency problem. An innovator or a principal (she) is endowed with a project. Neither she nor the hired agent (he) knows whether the project is viable, i.e. will it generate any profits. However, the principal is privately aware of the project's quality, i.e. how long the project, if it is viable, will take in expectation to generate profits.[3] Only the high quality project is worth exploring. The principal can commit to a long-term contract that specifies how payments depend on outcomes to incentivize the agent to exert private effort. The contract also serves as a signal to the quality of her project, or her type. When the agent works on the project, he also gradually learns about both the viability and the quality of the project. Hence, the principal has two tasks at the same time: signal her type and provide incentives to the agent to overcome the dynamic moral hazard problem. To isolate the conflict between signaling and providing incentives, we assume both players are risk neutral and have unlimited liability.

Since our focus is on rewarding innovation, we characterize the best equilibrium contract for the high type principal. The high type principal can never obtain her full information surplus, i.e. the payoff she would get if she was known to be a high type. Depending on the proportion of high quality projects in the population, the best equilibrium for her is either a separating equilibrium or a pooling equilibrium.

When high quality projects are scarce, a separating equilibrium gives the best outcome to the high type principal. Without limited liability, the high type principal could sell her project to the agent, leaving him to be the residual claimant that solves the dynamic moral hazard problem. Indeed, selling the project would be the optimal contract if the quality of project were publicly known. That is, the agent makes an upfront payment that equals to the expected surplus of the project, in exchange for the entire profit when the project succeeds. When only the principal knows her type, though, such a contract does not do a good job

[2]Stevens and Burley (1997) estimate that there is only one commercial success in 3,000 raw ideas of innovation across most industries. For drug development, Klees and Joines (1996) report that only one compound is approved for marketing among every 5,000 to 10,000 compounds that enter preclinical testing.

[3]To be more specific, the quality of a project is modeled as the success probability in one period when the project is viable and the agent exerts effort.

separating the two types. Both types value the upfront payment the same, and the low type principal has incentives to exaggerate the quality of her project to sell it for a good price. To separate from the low type principal and convince the agent of the project's quality, the high type principal would like to pay some positive base wage to the agent that does not depend on the outcome, as well as to share only a part of profits as bonus payments with him when the project succeeds. The base wage serves as a signal for the high type principal. The bonus payments are increasing over time to incentivize an increasingly pessimistic agent to continue working until the principal's desirable termination date. Thus, the high type principal shares a portion of the total surplus with the agent. Moreover, since the high type principal cannot ask for compensation for the dynamic moral hazard costs, she would like to terminate the project inefficiently early, and the inefficiency is relatively large. Thus, both the inefficient termination and sharing the surplus with the agent reduce the payoff of the high type principal.

On the other hand, when high quality projects are not rare, it is better for the high type principal to pool with the low type to avoid signaling costs. In that case, the best equilibrium for the high type principal is a pooling equilibrium, where both types charge the agent a positive sign-up fee, then share the profits with him once the project succeeds. Thus, instead of leaving rents to the agent, the high type principal "shares" rents with the low type principal. Moreover, the equilibrium contract still features inefficiently early termination, compared with a project whose quality is ex ante unknown.

The above results show that the high quality project is operated inefficiently, and the high type principal cannot obtain her full information surplus. This provides a role for a mediator to facilitate the agency process. Online platforms, such as Science Exchange[4], Scientist[5] and Upwork[6], are such mediators. Science Exchange and Scientist help scientists to outsource their research to other scientific institutions around the world. Upwork is a global freelancing platform that connects independent professionals.

In the second part of the paper, we analyze the mechanism design problem faced by a mediator who seeks to maximize the reward for superior innovations, i.e. who seeks to maximize the payoff of the principal with a high quality project. The designer offers a menu consisting of two contracts, one for each type of principal. This menu must satisfy incentive

[4] `https://www.scienceexchange.com`.
[5] `https://www.scientist.com`.
[6] `https://www.upwork.com`.

compatibility for the principal, i.e. it must induce truthful revelation by each type (it must also satisfy individual rationality). The agent observes the menu, and infers that each of the contracts will be implemented with a probability corresponding to his prior, and decides whether to accept or reject the menu. However, the agent is not told which element is being implemented until it is essential for him to know. In other words, the agent is confronted with an opaque contract, and remains uncertain about his exact rewards for some time. This relaxes the agent's individual rationality and incentive compatibility constraints; they only need to hold on average. In addition, the opaqueness of the contingent transfers allows the two types of principal to bet on a success, which provides an additional device for the high type principal to separate herself. In this way, the mediator designs a mechanism that improves the payoff of the high type principal. Moreover, the inefficiency costs are minimal, since the mediator could recommend the low type project running for only one period. In the optimal mechanism for the high type principal with pure recommendations, the high type principal obtains approximately all the surplus from her innovation when the period length is small. Thus the contract with the mediator allows the innovator of a superior project to appropriate almost her entire contribution to social surplus. Furthermore, an innovator who comes up with an inferior project is left with no surplus, thus the menu simultaneously minimizes the rewards for wasteful innovations.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 introduces the model. Section 4 discusses the efficient solution and other benchmarks. Section 5 characterizes the best equilibrium for the high type principal in the signaling game, and discusses its properties. Section 6 considers the third-party mechanism design problem. Section 7 concludes. The Appendix provides proofs, and studies the implementation problem of the full information benchmark.

2 Related Literature

This paper examines the incentives for experimentation when the principal is informed. It mainly relates to three strands of literature.

First, it relates to an increasing literature on incentives for experimentation. These papers, as well as the current one, build on a two-armed bandit model of learning, as in Keller, Rady, and Cripps (2005), and focus on how to provide incentives for agents to experiment through contingent contracts. Most papers, such as Bergemann and Hege (1998, 2005), and Hörner and Samuelson (2013), consider how to incentivize one party to experiment, who is subject to the moral hazard problem. They typically consider a repeated interaction between the principal and the agent, assume limited liability, and demonstrate an inefficiency result due to the agency costs. Guo (2016) studies a dynamic relationship in which a principal delegates experimentation to a biased agent who has private information about the prior belief that the state is good. Thus, the incentive problem comes from the hidden information of an informed but biased agent. Halac, Kartik, and Liu (2016) is a closely related paper. They examine an agency problem subjected to both moral hazard and adverse selection in the context of experimentation. They assume that the principal can commit to a long-term contract, and no limited liability. The major difference is that, instead of considering the private information on the side of the agent, our paper examines the case when the private information is on the side of the principal. In the screening problem of Halac, Kartik, and Liu (2016), the high type agent has an incentive to pretend to be the low type. In our signaling problem, the low type principal has an incentive to pretend to be the high type. Therefore, Halac, Kartik, and Liu (2016) show that the optimal contract has no distortion for the high type agent, but requires the low type agent to terminate the project inefficiently early.[7] By contrast, we show that it is the high type project that is terminated inefficiently early in the best equilibrium for the high type principal. Moreover, in that equilibrium, either there is no distortion for the low type project, or it is distorted towards over-experimentation, depending on the prior belief about the low type principal. Thus the economic forces underlying the analyses are very different in the two papers. Furthermore, when we allow for a mediator, we show that we can achieve approximate efficiency, a result that has no counter-part in Halac, Kartik, and Liu (2016).

Second, our paper relates to a relatively small literature on the informed principal problem with moral hazard. Myerson (1983) first considers the informed principal problem from an axiomatic point of view. Maskin and Tirole (1990, 1992) develop a noncooperative game framework to analyze the informed principal problem with no moral hazard. According to their categorization, our model is the "common value" informed principal problem, since the agent cares directly about the type of principal. Beaudry (1994) considers the informed principal problem with moral hazard. He characterizes a separating equilibrium where the principal leaves rents to the agent. More recent papers, such as Silvers (2012), Wagner, Mylovanov, and Tröger (2015), Bedard (2016), and Karle, Schumacher, and Staat (2016) also

[7]It is a typical result in screening problems that distortions occur at the bottom, but not at the top.

consider the informed principal problem in the presence of moral hazard. In all these papers, the moral hazard problem is static. Few papers examine this problem in a dynamic setting.[8] One exception is Kaya (2010). She studies a similar informed principal problem with moral hazard when the principal and the agent interact repeatedly. She considers a situation when the principal and the agent start with symmetric uncertainty about the productivity information. In each period, the principal has a choice to acquire that information without costs, but would rather delay to acquiring it in order to save the costs for signaling. However, in our paper, the agent learns both the type of the principal and the viability of the project during experimentation, which makes the incentive problem very different.[9]

Last, our paper is also related to the information design problem with moral hazard. Jehiel (2015) examines whether a principal with private signals prefers to commit to a non-transparent information disclosure policy to overcome the agent's moral hazard. He finds that full transparency is generically suboptimal under some mild conditions. Our result in the mechanism design part echoes his result - keeping the agent in the dark improves the payoff of the high type principal. However, we focus on the signaling problem of the principal after her private information is realized, while Jehiel (2015) assumes the principal can commit to an information disclosure policy before the realization of the private information. Ely and Szydlowski (2016) study a model where a principal is privately informed about the duration of required effort for completing a project. The principal's objective is to induce the agent to work as much as possible, thus they find that the optimal information disclosure policy is "moving the goalposts": at the outset, the principal tries to make the agent optimistic that the task is easy in order to induce him to start working, but persuades him that the task is hard when the difficult goal is within reach. In their setting, there is no signaling consideration, and the principal's problem is to keep the agent from quitting.

3 The Model

Time is discrete and the horizon is infinite, with a small but strictly positive period length $\Delta > 0$. The per period common discount factor is $\delta = e^{-\rho\Delta}$, where $\rho \geq 0$. To get rid of

[8] Fryer and Holden (2012) consider a two period informed principal problem with moral hazard. However, they make a behavioral assumption that the agent only learns from a noisy signal, but not from the contract proposed by the principal.

[9] Kaya (2010) uses "money burning", i.e. a pure production-irrelevant cost, as a signaling device for the principal after acquiring information. This money burning signaling device, together with the assumption that the agent has limited liability, are essential to her results. See Kaya (2010) Footnote 15.

integer problems, we will look at the case as $\Delta \to 0$ for some results.[10]

There are two risk neutral players, a principal P, and an agent A. The principal (she) hires an agent (he) to complete a project with uncertain viability. The agent, after accepting the principal's offer, could choose whether to exert efforts in every period before the game ends. The agent has flow costs $c\Delta$, where $c > 0$, when he exerts efforts in any period. The efforts are not observable or verifiable by the principal.

The principal is privately informed about the *quality* of the project, $\theta \in \Theta = \{H, L\}$. H represents a high quality project, and L represents a low quality project. The quality is persistent, and is chosen by nature at the beginning. We can also regard the quality as the *type* of the project/principal. The agent doesn't know the principal's type, but has a prior $\beta_0 \in (0,1)$ on the H type, which is common knowledge.

A project of either type may be in one of the two *states*: a good state G, or a bad state B. The state is also persistent and chosen by nature at the beginning. However, neither party knows the state of the project. They have a common prior $q_0 \in (0,1)$ on the G state, and the distribution of states is independent of the distribution of types.

The following figures summarize the information structure when nature moves. P knows the types, but not the states. A knows neither of them.

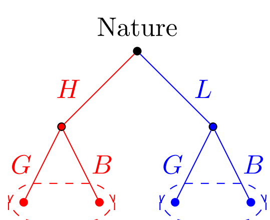

Figure 1: Principal's Information

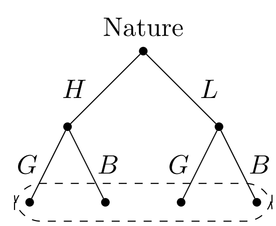

Figure 2: Agent's Information

In the B state, a project, independent of its type and the agent's efforts, never generates any profits. In the G state, when the agent exerts efforts in one period, a $\theta \in \Theta$ type project can generate a lump-sum profit $h > 0$ to the principal in that period with probability $\lambda^\theta \Delta$, and $\lambda^H > \lambda^L > 0$. It is a *success* of a project in some period when the lump-sum profit

[10]Those results approximate the circumstance when Δ is small. We do not use a continuous time model to avoid technical difficulties, since stochastic integration with respect to general transfer scheme may not be well defined. We also use the following notational conventions. N and n are used for time periods when we talk about the model with a fixed $\Delta > 0$; T and t are used for time when we examine the model in the limit $\Delta \to 0$. For a variable x_n that has time subscript n, let x_t be the limit of x_n as $\Delta \to 0$ and $n\Delta \to t$ (when it exists).

is generated, otherwise we call it a *failure*. A success of the project can be publicly observed and verified[11], and will end the game. Hence, the difference between the H and the L type projects is the arrival rate of success conditional on the project is in the G state and the agent exerts efforts. In addition, we will maintain the following assumption in this paper:

Assumption. $q_0\lambda^H h > c \geq q_0\lambda^L h$.

Thus, given the cost and the benefit of the project, and the prior belief about the state, only the H type project is worth experimenting on. However, if it was commonly known that the project's type was L, then it would exit the market.

The timing of the game is the following.

At the beginning of the game, the principal learns her type, H or L, then proposes a take-it-or-leave-it long-term contract to the agent. The principal can fully commit to her contract. A contract will specify the following:

- A termination date $N \in \mathbb{N}_0$ of the project conditional on no success.[12] This means the principal would shut down the project after N failures.
- A lump-sum payment, $W \in \mathbb{R}$ from the principal to agent at time 0. It can be positive or negative.
- A contingent payment plan during experimentation, two vectors $\boldsymbol{b} \in \mathbb{R}^N$ and $\boldsymbol{p} \in \mathbb{R}^N$, from the principal to the agent. $\boldsymbol{b}$ is for bonuses, it determines the payment from the principal to the agent when the project succeeds in any period; $\boldsymbol{p}$ is for penalties, it determines the payment from the principal to the agent when the project fails in any period. In other words, if the project succeeds in the n-th period, where $1 \leq n \leq N$, then the payments to the agent are p_k in the k-th period for $1 \leq k < n$, and b_n in the n-th period. Bonuses and penalties can be either positive or negative.

Hence, a *contract* is a quadruple $C = \{N, W, \boldsymbol{b}, \boldsymbol{p}\}$. Note that all above terms depend on

[11] The observability and verifiability of a success is a harmless assumption in the model. If a success is only observed by the agent, then the agent could choose to misreport a failure to the principal when a success actually arrives. However, the agent has no incentive to hide a success from the principal in any equilibrium contract in this paper. When a contract gives exact incentives for the agent to work if a success is contractible, it also gives enough incentives for the agent to truthfully report the success if it is not contractible. See Section 5.2.

[12] The principal, whatever type she is, never wants to experiment forever. Therefore, without loss, we do not need to consider the option to run the project forever.

Δ, however, we omit that to save notation. A contract is a *bonus contract* if $\boldsymbol{p} = \mathbf{0}$, and a *penalty contract* if $\boldsymbol{b} = \mathbf{0}$. If $N = 0$, then the payment in the contract only consists of W. A contract is a *null contract* if $N = 0$ and $W = 0$.

After the principal proposes a contract, the agent chooses whether to reject or accept it. If the agent rejects it, then both parties obtain their reservation payoffs 0 and the game ends. If the agent accepts it, then the contract is implemented and the agent will decide whether to exert efforts in each period, until the project succeeds or the termination date is reached.

Since the informed player (the principal) moves first, this is a signaling game. We will consider Perfect Bayesian Equilibria (PBE).

A contract allows any transfers between the principal and the agent. However, the set of the contracts is unnecessarily large. We can restrict attention to a smaller set of contracts, i.e. the set of bonus contracts or the set of penalty contracts, without loss of generality.

Let β be the probability of the H type, C be some arbitrary contract with terminating date N, and $\boldsymbol{a} \in A^N = \{0,1\}^N$ be agent's action plan, where 0 means shirking and 1 means working. Given the contract, the belief and the action plan, then

- $\Pi^\theta(C, \boldsymbol{a})$ denote the expected discounted payoff for type θ principal;
- $W_A^\beta(C, \boldsymbol{a})$ denote the expected discounted payoff for the agent.

We define *payoff-equivalent* contracts as follows:

Definition. *Two contracts C and $\hat{C}$, with the same termination date N, are payoff-equivalent, if for any $\boldsymbol{a} \in A^N$, any $\theta \in \Theta$ and any $\beta \in [0,1]$, $\Pi^\theta(C, \boldsymbol{a}) = \Pi^\theta(\hat{C}, \boldsymbol{a})$ and $W_A^\beta(C, \boldsymbol{a}) = W_A^\beta(\hat{C}, \boldsymbol{a})$.*

The transfer structure is rich enough to allow many different contracts that are payoff-equivalent. If two contracts induce the same discounted transfers on any realized event, then they are payoff-equivalent, since both parties are risk neutral and share the same time preference. Furthermore, for any realized event, i.e. whether and when the project succeeds, bonuses (or penalties) together with the lump sum payment W could deliver any discounted transfers that a general contract could. Formally,

Lemma 1. *For any contract, there exist both an equivalent bonus contract and an equivalent penalty contract.*

The idea and the proof are the same as in the Proposition 1 of Halac, Kartik, and Liu (2016). Our setup differs from theirs, since we have a signaling game rather than a screening problem. The agent may form different beliefs on payoff-equivalent contracts. However, the principal and the agent would feel indifferent to payoff-equivalent contracts as long as the agent forms the same beliefs, since they would induce the same action plan of the agent and deliver the same expected payoffs. Hence, we shall have no reason to assume that the agent would regard different payoff-equivalent contracts as different signals. Formally, we assume

Assumption. *For any payoff-equivalent contracts, the agent forms the same belief about the principal's type.*

Hence, without loss of generality, we can focus only on bonus contracts or penalty contracts. In this paper, we will restrict attention to bonus contracts, or *contracts* for simplicity. Thus, a contract is a triple $C = \{N, W, \boldsymbol{b}\}$.

4 Benchmarks

4.1 Efficient Solution

We first consider the efficient solution without the agency problem. Since we have a quasi-linear environment, consider a social planner who seeks to maximize total surplus, given that he knows both the private type of the principal and the hidden actions of the agent, but not the state of the project.

Clearly, the social planner would solve an optimal stopping problem. The optimal strategy is to stop the project as the posterior belief about the state of the project being G dropping to some cutoff belief. This strategy is also equivalent to specifying how long to experiment.[13]

Let $V^{\theta}(N)$ be the expected discounted value of the $\theta \in \Theta$ project when the social planner

[13] More precisely, a strategy specifies how long to experiment without a success. Since a success ends the game, we omit repeating "without a success".

experiments N times. Then it is given by:

$$V^\theta(N) = \sum_{n=1}^{N} \delta^n f^\theta_{n-1}(q_0)(q^\theta_n \lambda^\theta h - c)\Delta, \tag{1}$$

where $f^\theta_m(q) = q(1-\lambda^\theta\Delta)^m + 1 - q$ is the probability that a θ project, with prior q being in the G state, fails m times. $(q^\theta_n\lambda^\theta h - c)\Delta$ is the expected payoff for the n-th experiment conditional on a success not having arrived yet. Here, q^θ_n is the posterior belief about the state being G for a θ project before the start of the n-th experiment,[14]

$$q^\theta_n = \frac{q_0(1-\lambda^\theta\Delta)^{n-1}}{q_0(1-\lambda^\theta\Delta)^{n-1} + 1 - q_0}.$$

Thus, the optimal policy for the social planner, with a θ project, is to conduct the project as long as

$$q^\theta_n \lambda^\theta h \geq c.$$

Given the assumption $q_0\lambda^L h \leq c$, it is easy to see that the social planner would like to abort the L type project immediately; the L type project is a lemon. Thus, the *efficient stopping time for the L type project* is 0. On the other hand, given the assumption $q_0\lambda^H h > c$, the *efficient stopping time for the H type project*, denoted by N^H_*, is strictly positive. To get rid of the integer problem of N^H_*, it will be helpful to examine the limit of the experimenting time,

$$T^H_* := \lim_{\Delta\to 0} N^H_*\Delta = \frac{\ln l_0 - \ln l^H}{\lambda^H},$$

where $l_0 = \frac{q_0}{1-q_0}$ is the likelihood ratio of the prior belief that the state of the project is G, and $l^H = \frac{c}{\lambda^H h - c}$ is the likelihood ratio of the efficient cutoff posterior belief that the state of the H project is G. We also denote $V^\theta_0(T) := \lim_{\Delta\to 0, N\Delta\to T} V^\theta(N)$.

The efficient solution is obtained assuming away both the private information of the principal and the hidden action of the agent. In the next section, we consider the benchmarks where the incentive problem is one-sided.

[14] After a project fails $n-1$ previous experiments.

4.2 One-sided Incentive Problem

Our model has both private information on the side of the principal, and hidden actions on the side of the agent. Both are crucial, because the incentive problem would be trivial without either of them.

No Private Information - Let us first consider the case when the private type of the project is public information, while the agent's actions are still private. Although the principal still needs to incentivize the agent to experiment on the commonly unknown state of the project, the incentive problem is trivial. The principal can just sell the project to the agent, since the agent has no financial constraint. After the agent becomes the residual claimant of the project, he would like to implement the efficient solution. Thus, the H type can extract all the surplus by selling her project with a price at its expected value, and the L type would exit the market.

Observable Efforts - Another simple case is when the agent's efforts are publicly observable and verifiable. Thus, the principal can contract directly on the efforts. Even though the principal has private information on the type of the project, the incentive problem is also trivial. There exists a separating equilibrium that both projects are conducted efficiently and the H type principal extracts all the surplus of her project. The H type principal proposes an "honest" contract that pays the agent each period until her efficient stopping time for his costly efforts if and only if he exerts efforts, while the L type exits the market. The agent would always like to accept such "honest" contract and exert efforts, and believe contracts other than the "honest" contract are offered by the L type.

Therefore, in the above two cases, both projects are implemented efficiently, and all the surplus goes to the principal. We call the result as the *full information benchmark*, or *FIB*. As it will be shown later, in the presence of both private information on the side of the principal and hidden actions on the side of the agent, the principal needs to signal her type and incentivize the agent to work at the same time. Then, the FIB can never be achieved.

5 Equilibrium Characterization

In this section, we first show that there is no equilibrium that achieves the FIB. Multiple equilibria exist, since a PBE does not restrict the off-equilibrium beliefs. Because a L type project is a lemon, we examine how much a H type project can be rewarded in the equilibrium. Hence, we study the best equilibrium for the H type. We characterize the best equilibrium for the H type principal, which is either a separating equilibrium or a pooling equilibrium, depending on the prior belief about the H type principal.

5.1 The Impossibility of Achieving the FIB

If an equilibrium implements the FIB, then it is a separating equilibrium, in which the L type principal aborts the project, and the H type principal incentivizes the agent to work until the efficient stopping time and extracts all the surplus. More specifically, in such an equilibrium, the H type contract should leave no rents to the agent, and incentivize him to work until period N_*^H, when he believes the principal is the H type. In addition, the L type principal cannot get strictly positive payoff from mimicking the H type.

We first characterize all contracts that satisfy the agent's binding individual rationality (IR) constraint, and the agent's incentive compatible (IC) constraints until period N_*^H, when the agent believes the principal's type is H. Then we find the worst one among that set of contracts for the L type, and show that the L type can still obtain strictly positive payoff from the worst contract. Thus, any contract that implements the FIB must violate the IC constraint for the L type. Hence, the FIB cannot be achieved in any equilibrium.

When the agent believes that the principal's type is $\theta \in \Theta$, a contract $C = \{N, W, \boldsymbol{b}\}$ satisfies the IC constraints for the agent to work from period 1 to period N, when for all $1 \leq n \leq N$,

$$\sum_{s=n}^{N} \delta^{s-n} f_{s-n}^{\theta}(q_n^{\theta})(q_s^{\theta}\lambda^{\theta} b_s - c)\Delta \geq \sum_{s=n+1}^{N} \delta^{s-n} f_{s-n-1}^{\theta}(q_n^{\theta})(q_{s-1}^{\theta}\lambda^{\theta} b_s - c)\Delta. \qquad IC_A^{\theta}(N)$$

$IC_A^{\theta}(N)$ contains N inequality constraints. The left hand side (LHS) is the agent's expected discounted payoff when he experiments from period n until the termination date N, after failing $n-1$ previous experiments. It has the same structure as the social planner's expected discounted value in expression (1), except the agent's value of a success in period s is b_s, not h. The right hand side (RHS) is the agent's expected discounted payoff when he shirks in period n but experiments from then on until the termination date N, after failing $n-1$ previous experiments. Thus, $IC_A^{\theta}(N)$ prevents a one time profitable deviation (shirking) of the agent in all histories when he never shirks before. If the agent deviates and shirks in some periods before, then it is still optimal for him to work thereafter, since he is more optimistic than he would be had he worked.

The moral hazard problem is dynamic. In the n-th period, the expected payoff of working within the current period is $(q_n^{\theta}\lambda^{\theta}b_n - c)\Delta$, and shirking gives 0 current payoff. The continuation values of working and shirking are also different, due to two effects. First, the *learning effect*: the (unexpected) shirking makes the agent more optimistic about the project than the principal. After (unexpected) shirking, the principal's belief about the state being G is $q_{n+1}^{\theta} = \frac{q_n^{\theta}(1-\lambda^{\theta}\Delta)}{1-q_n^{\theta}\lambda^{\theta}\Delta} < q_n^{\theta}$, while the agent's belief is still q_n^{θ}. Second, the *end-of-game effect*: the game has only $f_1^{\theta}(q_n^{\theta}) = 1 - q_n^{\theta}\lambda^{\theta}\Delta$ probability to continue if the agent works, while it continues for sure if the agent shirks. Those two effects would increase the continuation value of shirking compared to working. Therefore, the bonus must compensate both the current period costs, and the loss of the continuation value for the agent.

When the agent believes that the principal's type is $\theta \in \Theta$, a contract $C = \{N, W, \boldsymbol{b}\}$ satisfies the IR constraint for the agent, when

$$W_A^{\theta} = W + \sum_{n=1}^{N} \delta^n f_{n-1}^{\theta}(q_0)(q_n^{\theta}\lambda^{\theta}b_n - c)\Delta \geq 0. \qquad IR_A^{\theta}(N)$$

Here, the agent receives both an upfront payment W before the experimentation, and a bonus for success during the experimentation.

Let $S(N)$ be the set of contracts that satisfy both $IC_A^H(N)$ and a binding $IR_A^H(N)$. Hence, the set of contracts, which implement the FIB for the H type when the agent believes her type is H, is $S(N_*^H)$.

The worst contract $C^{wt} = \{N_*^H, W^{wt}, \boldsymbol{b^{wt}}\}$ for the L type in the above set of contracts solves the following Program I:

$$\min_{\boldsymbol{b},W} \Pi^L = -W + \sum_{n=1}^{N_*^H} \delta^n f_{n-1}^L(q_0) q_n^L \lambda^L (h - b_n)\Delta$$
$$\text{s.t. } IC_A^H(N_*^H) \text{ and a binding } IR_A^H(N_*^H).$$

Note that the expected payoff of the L type is determined by the payment transfered to the agent before experimentation $-W$, and the kept share of profits once the project succeeds. Here, $q_n^L \lambda^L (h - b_n)\Delta$ is expected payoff of the L type principal for the n-th experiment conditional on a success not having arrived yet.

Lemma 2. *The worst contract for the L type in the set $S(N_*^H)$ of contracts is the contract in which the agent's IC constraints bind for every period $n \in \{1, 2, \ldots, N_*^H\}$.*

We now provide intuition for why all the IC constraints must bind. The set of contracts that implement the FIB for the H type is large, because the principal has the discretion to give more high powered incentives than required, so that the agent's IC constraints are slack, and then extract the surplus so conferred via a larger sign-up fee, i.e. a lower value of W. However, such a contract is more attractive to the L type. We now explain why.

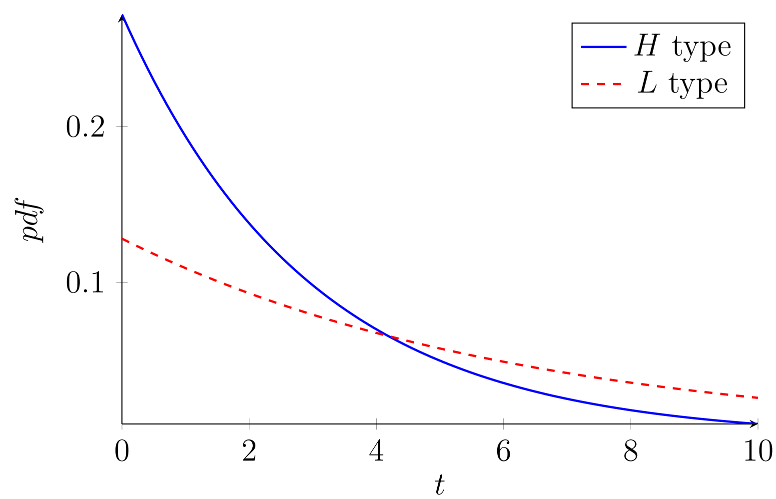

Figure 3: The probability of success: the H type and the L type

Figure 3 shows the probability of success over time for the two types. Let $\eta_n := f_{n-1}^H(q_0) q_n^H \lambda^H - f_{n-1}^L(q_0) q_n^L \lambda^L$, then $\eta_n \Delta$ denote the difference of probabilities of success for the n-th experiment between the H and L types. Though the H type is more likely to succeed than the L type conditional on the state being G, the posterior beliefs about the state being G conditional on no success decrease faster for the H type than for the L type, since the H type also learns faster. Therefore, the two curves cross only once. In the earlier periods, the H type is more likely to succeed ($\eta_n > 0$), while in the later periods, the L type is more likely to succeed ($\eta_n < 0$).

Suppose that the agent believes that the principal's type is H, but her type is actually L. Thus, the agent overvalues the bonus payments when the probability of success is greater for the H type than that for the L type, and is willing to pay a higher sign-up fee. Consequently, to minimize the incentive of the L type to mimic the H type, bonus payments should be minimized in any period where the probability of success is greater for the H type than that for the L type.

Consider now any period where the probability of success is greater for the L type than that for the H type. By the preceding argument, it would seem that the bonus payment in such a period should be increased, since they reduce the incentives of the L type to mimic the H type. However, this is not true. Consider the period right after the crossing point. Raising the bonus in that period makes the agent's continuation value of shirking, in all previous periods, higher. To incentivize the agent to work, the principal needs to raise bonus payments in all previous periods proportionally. It actually makes the L type better off, because the distribution of success for the L type first-order stochastic dominates that for the H type, i.e. $\phi_n := \sum_{s=1}^{n} \eta_s \Delta > 0$ for any $n \geq 1$. To minimize the L type's incentive to mimic the H type, the bonus payment should also be minimized in the period right after the crossing point. By induction, we can show the bonus payments in any period should be minimized in the worst contract for the L type. Thus, all IC constraints should bind in the worst contract for the L type. This result is important, and it will hold whenever the H type has a signaling concern.[15]

Even in the worst contract for the L type, the bonus scheme is an increasing sharing plan that leaves the L type with positive expected payoff during experimentation. Moreover, the agent is also willing to pay some positive sign-up fee, if he believes the principal's type is H. Clearly, the L type can obtain strictly positive payoff from that worst contract, and has incentive to mimic the H type. Hence, we conclude:

Proposition 1. *There is no equilibrium that implements the FIB.*

5.2 Separating Equilibrium

The FIB cannot be implemented in any equilibrium, since any contract that achieves the FIB for the H type must violate the IC constraint for the L type. Thus, how can the H

[15]See Lemma 6 in the Appendix.

type principal separate herself from the L type?

The simplest way is to increase the upfront payment W. Since the L type has a lower probability of success than the H type, an upfront payment W that makes the H type profitable could be unprofitable to the L type. Consider the contract C^{wt} we found. Let $\Pi^L(C^{wt})$ be the payoff of the L type from C^{wt}, when the agent believes her type is H. Now, we introduce a new contract, $C^1 = \{N_*^H, W^1, \boldsymbol{b^{wt}}\}$, with the same termination date and bonus payments, but a higher upfront payment, where $W^1 = W^{wt} + \Pi^L(C^{wt})$. Hence, the L type would obtain 0 payoff from C^1, even when the agent believes her type is H. However, the H type can obtain strictly positive payoff from C^1, when the agent believes her type is H. This can be seen from the comparison between contracts C^1 and $C^0 = \{N_*^H, 0, \boldsymbol{h}\}$, where $\boldsymbol{h} \in \mathbb{R}^{N_*^H}$ is the vector whose entries are all h. Clearly, C^0 gives the project to the agent for free, and leaves the principal 0 payoff. In addition, the argument from Lemma 2 asserts that the H type prefers C^1 to C^0, since the former has tight IC constraints for the agent, while the latter has slack IC constraints.

Therefore, C^1 can separate the H type from the L type. Furthermore, the H and L type proposing C^1 and the null contract respectively, constitute part of an equilibrium. In other words, for any other contract, we can find some belief that prevents a profitable deviation by either type. In fact, we have the following result:

Lemma 3. *For any contract, if the agent believes the principal's type is L, then the principal, whatever type she actually is, cannot obtain a strictly positive payoff.*

Hence, a pessimistic belief that assigns probability 1 to type L for any off-equilibrium path contract can support the above equilibrium.

The contract C^1 simply uses the upfront payment to separate the two types. The H type could design a more sophisticated contract to separate herself, and improve her payoff. Now, we consider the *best separating equilibrium for the H type*, or *BSEH*, i.e. the equilibrium that gives the H type the highest payoff among all separating equilibria.

Given the structure of a contract, the H type principal could use three potential devices to separate herself. First, she could increase the upfront payment W as we have seen. Second, she could reduce the experimentation by terminating it earlier. Since the L type learns slower than the H type, the value of an additional experiment decreases more slowly for the former than for the latter. Hence, an experiment could be more valuable to the L type than to the H type after some time. Thus, shutting down the project earlier makes mimicry less attractive. Third, she could delay the project, by having the agent not work in some periods before terminating the project permanently. Those devices are all costly to the H type; the problem is how to use them in the least costly way.

Though it is not clear at first whether delaying the project temporarily permits separation, we can show that this is not the case. Delaying the project essentially increases the discount factor between two consecutive experiments.[16] Each discount factor affects discounted payoffs linearly, if all terms in the payoff functions are independent with the discount factor. To make the agent follow the induced action plan, bonus payments have to be correlated with future discount factors. However, allowing for the possibility to delay the project does not change the fact that the least costly separation and the least costly pooling tend to have the agent's IC constraints bind. It, in turn, makes the current bonus payment linearly correlate with any future discount factor. Thus, discounted payoffs are linear in each discount factor between two consecutive experiments. Therefore, it is optimal to either never delay the project, or delay the project forever (i.e. terminating the project), from the H type's point of view. Hence, the H type project will never be temporarily delayed in the best equilibrium for the H type. From now on, we will not consider delaying the project temporarily in the main body of the paper. More details can be found in the Appendix.

In the BSEH, the L type must obtain a payoff of 0; let her propose the null contract in the equilibrium. We also assume that the agent assigns probability 1 to the belief about the principal being the L type for any off-equilibrium path contract. Thus, Lemma 3 ensures that both types of principal would not deviate to those contracts. We shall not worry about the H type's IR constraint and her incentives to choose the L type's equilibrium contract in the BSEH. The last constraint for the principal we need to consider, is the L type's incentives to deviate to the H type's equilibrium contract. Given the H type's contract $C = \{N, W, \boldsymbol{b}\}$, the L type's IC constraint is

$$\Pi^L = -W + \sum_{n=1}^{N} \delta^n f_{n-1}^L(q_0) q_n^L \lambda^L (h - b_n)\Delta \leq 0. \qquad IC_L^H(N)$$

The agent's IR and IC constraints are the same as before. Hence, the H's equilibrium

[16]Given the time preference δ, the discount factor between the $(n-1)$-th and n-th experiments is $\delta_n = \delta^k$, where $k \in \mathbb{N} \cup \{\infty\}$ is the number of delaying periods. $k = 1$ means no delay, and $k = \infty$ means terminating the project.

contract, $C^{sep} = \{N^{sep}, W^{sep}, \boldsymbol{b^{sep}}\}$, in the BSEH will solve the following Program II:

$$\max_{N,W,\boldsymbol{b}} \Pi^H = -W + \sum_{n=1}^{N} \delta^n f_{n-1}^H(q_0) q_n^H \lambda^H (h - b_n)\Delta$$
$$\text{s.t. } IC_A^H(N) \text{ , } IR_A^H(N) \text{ , and } IC_L^H(N)$$

To solve this program, we drop the constraint $IR_A^H(N)$, and verify it later. In the relaxed program, $IC_L^H(N)$ must bind, which determines the time zero transfer W, otherwise we can decrease W to obtain a higher payoff for the H type. For any given N, the agent's IC constraints in $IC_A^H(N)$ also bind for every period $1 \leq n \leq N$ due to the same reason as in Lemma 2, which determines the bonus scheme $\boldsymbol{b}$. The last thing is to determine the termination date N. We take the limit $\Delta \to 0$ to avoid the integer problem, and obtain the following result:

Proposition 2. *In the limit $\Delta \to 0$, there is a unique equilibrium contract for the H type $C^{sep} = \{T^{sep}, W^{sep}, \boldsymbol{b^{sep}}\}$ in the BSEH, and it has the following features:*

- *Under experimentation: $0 < T^{sep} < \frac{1}{2}T_*^H$, where T^{sep} satisfies*

$$\lambda^H e^{-\lambda^H T^{sep}} - \lambda^L e^{-\lambda^L T^{sep}} = (\lambda^H - \lambda^L) e^{-\lambda^H T_*^H} e^{(\lambda^H - \lambda^L)T^{sep}};$$

- *A positive base wage: $W^{sep} > 0$;*
- *A (weakly) increasing bonus plan: b_t^{sep} is weakly increasing in t, and $0 < b_t^{sep} < h$ for $0 \leq t \leq T^{sep}$.*

Rent Sharing - In such an equilibrium, the H type principal shares some surplus with the agent. She obtains less than $1 - \lambda^L/\lambda^H$ of the total surplus $V_0^H(T^{sep})$, and the agent obtains more than λ^L/λ^H of it. The surplus is shared through both a positive base wage W^{sep} before experimentation, and an increasing bonus scheme in the period when the success arrives. Even though the principal has all the bargaining power and the agent has no financial constraint, she leaves some rents to the agent, to signal that her type is H. This phenomenon is first shown in Beaudry (1994). The L type is left with no rents, and exits the market in this equilibrium.

Role of Learning - An increasing bonus scheme means the principal would give a larger reward to the agent if the success comes later. This is due to the decline of the posterior beliefs about the viability of the project. A failed experiment drives down the agent's belief about the viability of the project, and greater incentives are needed in later periods.

Inefficiently early termination is the second type of signaling cost. Since the initial payment W must equal the L type's expected share of the profits if the L type proposes the H type's contract, the H type principal's payoff can be seen as the difference between her expected share of the profits and that of the L type. Extending the experiment has two marginal effects.

Consider only the last period payoffs, the marginal net benefit for extending the experiment is the difference between the H type's expected share of her profits and that of the L type in that period:

$$q_0 e^{-\rho T}(\lambda^H e^{-\lambda^H T} - \lambda^L e^{-\lambda^L T})(h - b_T), \tag{2}$$

where $\lambda^\theta e^{-\lambda^\theta T}$ is the probability of success in the last period T for the $\theta \in \Theta$ type project in the G state, and $h - b_T$ is the share retained by the principal.

Consider all other periods payoffs, due to the dynamic moral hazard problem, the marginal net cost for extending the experiment is the difference between the increased accumulated bonuses that are given up by the H type and that of the L type:

$$q_0 e^{-\rho T}(e^{-\lambda^L T} - e^{-\lambda^H T})(\lambda^H b_T - c), \tag{3}$$

where $e^{-\lambda^L T} - e^{-\lambda^H T}$ is the probability difference of failures until the last period T between the H and L type in the G state, and $\lambda^H b_T - c$ represents the increased bonuses.[17]

Equalizing the above two effects gives the equilibrium terminating date T^{sep}. The inefficiency comes from the combination of the private information and actions on both sides. Moreover, it is aggravated because of the dynamic moral hazard problem due to *learning*. If the state of the project is known, then extending the experiment will not increase the agent's share of the profits in previous periods; the principal does not need to give more incentives to the agent than the current experimenting cost. Hence, the effect represented by expression (3) is zero, when there is no learning towards the state of the project. Then, the inefficiency is determined only by expression (2), and it would diminish, i.e. the terminating date goes to be infinity, as λ^L goes to 0. However, in our model, the state of the project is unknown; thus, prolonging experimentation increases the bonus payments in all periods. Even when

[17]The marginal increase of a bonus at time t by extending the experiment is $\lambda^H b_T - c$ with proper discounting, i.e. $\frac{\mathrm{d}b_t}{\mathrm{d}T} = e^{-\rho(T-t)}(\lambda^H b_T - c)$.

λ^L goes to 0, expression (3) is strictly positive. In that case when $\lambda^L = 0$, the equilibrium termination date, T^{sep}, is determined by

$$q_0 e^{-\rho T} e^{-\lambda^H T} \lambda^H (h - b_T) = q_0 e^{-\rho T} (1 - e^{-\lambda^H T})(\lambda^H b_T - c),$$

where $\lambda^H (h - b_T)$ and $(\lambda^H b_T - c)$ are the expected payoffs for the principal and the agent in the G state, respectively. Therefore, for $T = T^{sep}$

$$e^{-\lambda^H T} = \frac{\lambda^H b_T - c}{\lambda^H h - c} = \frac{l^H}{l_T} = e^{-\lambda^H (T_*^H - T)},$$

where $l^H = \frac{c}{\lambda^H h - c}$ and $l_T = \frac{c}{\lambda^H b_T - c}$ are the likelihood ratio of the posterior beliefs that the state of the H project is G at the efficient stopping time T_*^H and the stopping time T, respectively. Hence, we have $T^{sep} = \frac{1}{2} T_*^H$ when $\lambda^L = 0$.

We can show that the equilibrium termination date T^{sep} is decreasing in λ^L, therefore the H type project is only operated less than half of the efficient time.

Limited Liability - Consider a different setting where the principal is publicly known to be the H type, but the agent is protected by limited liability, i.e. W and $\boldsymbol{b}$ must be positive. Then, it is easy to show that the optimal contract in that case is the H type's equilibrium contract in the BSEH of our original model when $\lambda^L = 0$, in which $W^{sep} = 0$. Hence, our model provides an alternative explanation for the limited liability. Economists observe that the agent is sometimes protected by limited liability. The principal cannot charge anything from the agent, or sell the project to him. A standard argument is that the agent may have some financial constraints. However, in our model, the principal is willing not to charge anything from the agent, not due to any limited liability protection, but to signal her type. Without financial constraints, we still may observe the "limited liability" phenomenon in equilibrium.

Unobservable Successes - We assume a success is publicly observable and verifiable so far. However, it is sometimes very costly for the principal to verify a success. If those costs are prohibitively high, then a contract can only depend on the agent's report on successes, but not successes per se. Hence, the principal may face an additional incentive problem: the agent may delay reporting a success to obtain a higher discounted payoff if the increase of bonuses over time offsets the costs of discounting.

However, this is not the case for the equilibrium contract. The agent has no incentive to hide a success from the principal, as long as he is given the exact incentives to experiment when a success can be costlessly verified, as in our original model. In other words, given the binding IC constraints for the agent to experiment, the agent would like to report a success truthfully when it arrives.

Formally, Lemma 7 in the Appendix shows that binding IC constraints for the agent implies, for any n before the termination date,

$$b_n \geq \delta b_{n+1}.$$

When $\delta < 1$, the above inequality is strict. Since all the remaining results in this paper feature no temporary suspension and binging IC constraints for the agent, the agent has no incentive to delay reporting a success even when it cannot be observed by the principal.

Equilibrium Refinement - We have shown two different separating equilibria, there are many others. Let us consider a commonly used equilibrium refinement, i.e. the intuitive criterion from Cho and Kreps (1987). We have the following result:

Proposition 3. *The BSEH survives the intuitive criterion, and all equilibria which give a lower payoff to the H type fail the intuitive criterion.*

Note that all other separating equilibria gives less payoff to the H type, thus fail the intuitive criterion. In addition, according to Lemma 5 in Section 5.4, all other equilibria that give 0 payoff to the L type also give less equilibrium payoff to the H type than the BSEH does, thus fail the intuitive criterion. However, there may be some pooling equilibrium gives more payoff for H type. We examine them in the next section.

5.3 Pooling Equilibrium

A pooling equilibrium is an equilibrium in which both types of principal propose the same equilibrium contract, and the agent holds the prior belief β_0 about the H type after the equilibrium contract is proposed.

Let us first consider the efficient solution when the type of the project is unknown at the ex ante stage. We call it the *mixed* project. When the prior belief is β_0, the value of the mixed project with the termination date N is

$$V(N) := \beta_0 V^H(N) + (1-\beta_0)V^L(N).$$

Let N_* be the *efficient stopping time for the mixed project*, which maximizes $V(N)$, and let T_* be the limit of $N_*\Delta$ when $\Delta \to 0$. When the prior β_0 is too pessimistic, i.e. $q_0\lambda_0 h \leq c$, where $\lambda_0 := \beta_0\lambda^H + (1-\beta_0)\lambda^L$, the efficient stopping time is 0. Thus, any pooling equilibrium is trivial, and features no experiments or transfers. Otherwise, when the prior β_0 is not too small, i.e. $q_0\lambda_0 h > c$, the efficient stopping time for the mixed project is in between 0 and the efficient stopping time for the H type. We will only consider the latter case.

One simple pooling equilibrium is that both types of principal sell the mixed project to the agent, and extract all expected surplus of the mixed project. Lemma 3 ensures no one would like to deviate to any off-equilibrium path contract, where the agent believes the deviator is the L type. However, the H type could achieve a higher payoff in some other pooling equilibrium. Again, we consider the *best pooling equilibrium for the H type*, or *BPEH*.

When both types of principal pool together, the agent would update his beliefs on both the type and the state of the project during the experimentation. Furthermore, even the type and the state of the project are independent at the outset, the posterior beliefs after some failures will be correlated. Let q_n be the posterior belief about the state of the project being G, after $n-1$ failures, and β_n be the posterior belief about the type of the project being H conditional on the state being G, after $n-1$ failures. The posterior belief about the type of the project being H conditional on the state being B is the prior β_0, no matter how many failures the project has. Hence,

$$q_n = \frac{q_0[\beta_0(1-\lambda^H\Delta)^{n-1} + (1-\beta_0)(1-\lambda^L\Delta)^{n-1}]}{q_0[\beta_0(1-\lambda^H\Delta)^{n-1} + (1-\beta_0)(1-\lambda^L\Delta)^{n-1}] + 1 - q_0},$$

and

$$\beta_n = \frac{\beta_0(1-\lambda^H\Delta)^{n-1}}{\beta_0(1-\lambda^H\Delta)^{n-1} + (1-\beta_0)(1-\lambda^L\Delta)^{n-1}}.$$

Let $\lambda_n := \beta_n\lambda^H + (1-\beta_n)\lambda^L$ be the expected arrival rate of success conditional on the state being G, after $n-1$ failures. Both q_n are β_n are strictly decreasing over time towards 0, and λ_n is strictly decreasing over time towards λ^L. That is, both the posterior belief about the G state and the posterior belief about the H type conditional on the state being G go down after a failure.

Now the agent's IC constraints are more complex due to the learning process. However, those constrains can be easily transformed from the constraints in the separating equilibrium by using correct beliefs. Given $C = \{N, W, \boldsymbol{b}\}$, the agent's IC constraints become, for $1 \leq n \leq N$

$$\sum_{s=n}^{N} \delta^{s-n} f_{s-n}(q_n, \beta_n)(q_s \lambda_s b_s - c)\Delta \geq \sum_{s=n+1}^{N} \delta^{s-n} f_{s-n-1}(q_n, \beta_n)(q_{s-1}\lambda_{s-1} b_s - c)\Delta, \quad IC_A(N)$$

where $f_m(q, \beta) = q[\beta(1-\lambda^H\Delta)^m + (1-\beta)(1-\lambda^L\Delta)^m] + 1 - q$ is the probability of failing m times for a project starting with probability q being in state G, and probability β being type H conditional on being in state G.

The agent's IR constraint becomes

$$W_A = W + \sum_{n=1}^{N} \delta^n f_{n-1}(q_0, \beta_0)(q_n \lambda_n b_n - c) \geq 0. \qquad IR_A(N)$$

In the BPEH, the principal's incentive is trivially satisfied on the equilibrium path. By Lemma 3, no type has an incentive to deviate to any other contract once we assume that the agent assigns probability 1 to the belief about the principal being the L type for those contracts. We shall not worry about the IR constraint for the H type. The IR constraint for the L type is

$$\Pi^L = -W + \sum_{n=1}^{N} \delta^n f_{n-1}^L(q_0) q_n^L \lambda^L (h - b_n)\Delta \geq 0. \qquad IR_L(N)$$

Then, the equilibrium contract, $C^{pl} = \{N^{pl}, W^{pl}, \boldsymbol{b^{pl}}\}$, in the BPEH will solve the following Program III:

$$\max_{N,W,\boldsymbol{b}} \Pi^H = -W + \sum_{n=1}^{N} \delta^n f_{n-1}^H(q_0) q_n^H \lambda^H (h - b_n)\Delta$$
$$\text{s.t. } IC_A(N),\ IR_A(N), \text{ and } IR_L(N)$$

$IR_L(N)$ is slack in the above program, so we drop it and verify it after solving the relaxed program. $IR_A(N)$ must bind in the relaxed program, otherwise we can decrease W to obtain more payoff for both types. Thus, it determines the time zero transfer W. Given any N, we can show that the agent's IC constraints in $IC_A(N)$ must bind for all $1 \leq n \leq N$, which determines the bonus scheme $\boldsymbol{b}$.[18] Then, we can solve the program and find the equilibrium

[18]The reason for binding IC constraints of the agent is similar to Lemma 2 and Proposition 2. Notice that the agent obtaining no rents implies that the payoff of the H type principal equals to the sum of the

stopping time.

Proposition 4. *Assume that* $q_0\lambda_0 h > c$*. In the limit* $\Delta \to 0$*, there is a unique equilibrium contract for the* H *type* $C^{pl} = \{T^{pl}, W^{pl}, \boldsymbol{b^{pl}}\}$ *in the BPEH, and it has the following features:*

- *Under experimentation:* $0 < T^{pl} < T_*$*;*
- *Sign-up fee:* $W^{pl} < 0$*;*
- *A (weakly) increasing bonus plan:* b_t^{pl} *is weakly increasing in* t*, and* $0 < b_t^{pl} < h$ *for* $0 \leq t \leq T^{pl}$*.*

Moreover, the equilibrium payoff for the H *type in the BPEH is strictly increasing in the prior belief about the* H *type* β_0*, and converges to* $V_0^H(T_*^H)$ *as* $\beta_0 \to 1$*.*

The above contract shares some common features with the H type's equilibrium contract in the BSEH. It gives the agent exact enough incentives to work, i.e. the IC constraints for the agent always bind. When there is no discounting, i.e. $\rho = 0$, the bonus payments are constant over time. When $\rho > 0$, they are strictly increasing. Furthermore, as in the BSEH, the increase of bonus payments cannot offset the cost for discounting. Hence, even when a success is not publicly observed, the agent has no incentive to delay reporting a success.

Pooling Costs - Unlike the case in the BSEH, where the H type gives a positive base wage to the agent, in the BPEH, both types charge a sign-up fee $-W^{pl} > 0$ from the agent before experimentation, leaving no rents to the agent. Since the agent has no financial constraint, he is willing to pay the sign-up fee to exchange for some share of profits, i.e. bonuses, once the project succeeds. The amount that the H type principal extracts from the agent cannot recover her FIB surplus due to the pooling costs. When the L type pools with the H type, the latter has to compensate more costs to the agent, but only partially extracts the promised bonuses back from the agent. Both the additional compensation costs and the non-recoverable bonuses are proportional to the population of the L type, i.e. $1 - \beta_0$. Thus, those pooling costs are diminishing as the prior belief about the H type goes to 1. On the other hand, the L type is left with strictly positive rents by pooling with the H type.

Learning Viability and Quality - When both types pool together, the agent needs to

project's value and $(1 - \beta_0)$ portion of the payoff difference between the H and L type. Hence, given a fixed amount of experimentation, the pooling contract in the BPEH must maximize the payoff difference.

learn both the viability and the quality of the project. A failure means the agent would be more pessimistic to both the viability and the (conditional) quality[19], thus learning creates dynamic moral hazard costs. However, learning the viability and the quality has different consequences. If what needs to learn is the viability ($0 < q_0 < 1$), it is really costly to extend the experiment, since the posterior beliefs of generating success eventually goes to 0. However, if what needs to learn is solely the quality ($q_0 = 1$, and $0 < \beta_0 < 1$), the principal can always incentivize the agent to work with bounded bonuses, i.e. $b_t \leq \frac{c}{\lambda^L}$ for any $t > 0$. Formally, let T be the termination date, then the limit of the marginal effect of extending experiment for the last period bonus is

$$\lim_{T\to\infty} \frac{\mathrm{d}b_T}{\mathrm{d}T} = \begin{cases} \infty & \text{if } q \in (0,1), \\ 0 & \text{if } q = 1. \end{cases}$$

Under Experimentation[20] - It is straightforward that the equilibrium stopping time is less than the efficient stopping time for the H type. In addition, it is also less than the efficient stopping time for the mixed project. We now provide some intuition.

Since the agent has no rents in the equilibrium, the value generated by the mixed project is shared by two types of principals. Thus, the equilibrium payoff of the H type is equal to the sum of (1) the generated value, and (2) $(1 - \beta_0)$ portion of the equilibrium payoff difference between the H and L type.

Consider the marginal effect of extending the project at the efficient stopping time for the mixed project, T_*. We will show that the marginal effect is negative at T_*, thus it is better to terminate the project earlier.

Consider the first part of the H type's payoff, i.e the value generated by the mixed project. By the definition of the efficient stopping time for the mixed project, extending the project has zero marginal effect on the generated value. For the second part of the H type's payoff, extending the project does not generate any benefit in the last period, since the required bonus (for incentivizing the agent to work) is already equal to the lump-sum profit of success at T_*. However, due to the dynamic moral hazard problem and the fact that the H type

[19] The conditional quality means the project's quality conditional on it being viable. Note that the conditional quality is what matters for the agent's incentives, and the unconditional quality could be increasing with one more failure after some point.

[20] The contract in the BPEH features under experimentation only for the H type principal, while the same contract requires over experimentation for the L type principal.

is more likely to succeed before the termination date than the L type, the costs for sharing more bonuses to the agent in all periods incurred by the H type is larger than that incurred by the L type. Thus, extending the project does generate some additional costs. Therefore, the marginal effect of extending the project at T_* is negative; the H type principal would like to terminate the project earlier than the efficient stopping time for the mixed project.

Equilibrium Refinement - Though the BPEH may give a higher payoff to the H type than the BSEH does, we can show that it fails the intuitive criterion:

Proposition 5. *The BPEH fails the intuitive criterion.*

This is because there exists a contract that enlarges the payoff difference between the two types of principal without scarifying efficiency, when the agent believes the deviator is the H type. Enlarging the payoff difference between the two types leaves room to increase the H type's payoff and decrease the L type's payoff at the same time. Keeping the efficiency makes the agent still participate voluntarily when the H type asks for more share of surplus. Hence, such a contract could make the BPEH fail the intuitive criterion.

5.4 Best Equilibrium for the High Type

We have examined both the BSEH and the BPEH. In the BSEH, the equilibrium payoff for the H type is independent of the prior belief β_0, and is less than her FIB surplus. However, in the BPEH, the equilibrium payoff for the H type is strictly increasing in β_0 when her payoff is strictly positive, i.e. when $q_0\lambda_0 h > c$, and converges to her FIB surplus as β_0 goes to 1. Obviously, we can find a cutoff prior belief about the H type, $\beta_c \in (0, 1)$, such that when $\beta_0 < \beta_c$, the BSEH gives the H type a higher payoff than the BPEH does, and when $\beta_0 > \beta_c$, the BPEH gives the H type a higher payoff than the BSEH does.

For the purpose to find the best equilibrium for the H type, we still need to examine all other equilibria, like partial separating equilibrium which allows the principal to randomize over different contracts.

To use the results we have already established, we categorize all equilibria into two classes: (1) the set of equilibria that gives the L type strictly positive payoff, and (2) the set of equilibria that gives the L type 0 payoff.

We will show that (1) among all equilibria that gives the L type strictly positive payoff, the BPEH gives the H type the highest payoff, and (2) among all equilibria that gives the L type 0 payoff, the BSEH gives the H type the highest payoff. Thus, we can conclude that the best equilibrium for the H type is either BSEH or BPEH, depending on β_0.

First, consider the set of equilibria that gives the L type strictly positive payoff. This means that no equilibrium path contract can fully reveal the L type. In other words, the L never chooses any contract by herself alone in the equilibrium, she always pools with the H type. Moreover, there must be at least one contract proposed by the L type in the equilibrium, such that the agent's belief about the principal being the H type does not exceed the prior β_0. Together with Proposition 4, we can show that:

Lemma 4. *For any equilibrium such that the L type obtains strictly positive payoff, the H type's payoff cannot exceed her payoff in the BPEH.*

Now, consider the set of equilibria that gives the L type 0 payoff. Fix any such equilibrium E, and any equilibrium contract C proposed by the H type in E. Suppose the agent forms belief $\beta \in [0, 1]$ about H type after C is proposed. Then C must satisfy the agent's IC and IR constraints, and the L type cannot get strictly positive payoff by proposing the same contract. We solve the best contract for the H type subjected to the above constraints, and show that the best outcome the H type can obtain is increasing in β. In addition, we come back to our BSEH result when $\beta = 1$. Hence, we show the following result:

Lemma 5. *For any equilibrium such that the L type obtains 0 payoff, the H type's payoff cannot exceed her payoff in the BSEH.*

With the above two lemmas, we can conclude:

Proposition 6. *There exits a cutoff belief $\beta_c \in (0, 1)$,*

- *when $\beta_0 < \beta_c$, the BSEH gives the H type the highest payoff among all equilibria;*
- *when $\beta_0 > \beta_c$, the BPEH gives the H type the highest payoff among all equilibria.*

Thus, in the best equilibrium for the H type, the stopping time of the H type project is inefficiently early, and she has to share rents with either the agent (in the BSEH), or the L type (in the BPEH). Since the equilibrium contract for the H type in the BSEH also maximizes the payoff difference between the H and L type when the agent assigns probability 1 to the belief about the principal being the H type, together with Lemma 8 in the Appendix, we can conclude that the BSEH is the equilibrium that maximizes the payoff difference between the two types of principal. Hence, if we consider a pre-game where the principal can invest in the quality of a project before the signaling game, namely exerting efforts to increase the probability β_0 of having a H type project, her investment decision, i.e. the prior β_0, is endogenously determined by the payoff difference between the two types in the equilibrium of the signaling game. Thus, the BSEH permits the highest prior, as long as the marginal investment cost is increasing in β_0.

6 Introducing a Mediator

The above results show the conflicts between signaling and providing incentives to the agent. We now consider the implication of allowing for a mediator who designs a mechanism (i.e. a menu of contracts).

As in the rest of the paper, we assume that the goal of the mediator is to maximize the payoff of the H type principal. This is a natural benchmark that facilitates the comparison across the signaling game and the mechanism design problem. The mediator needs to attract the H type project, which generates the social surplus. Thus, it must provide her with a higher payoff than she can obtain in the signaling game without the mediator. Furthermore, if there are multiple profit-oriented mediators compete in the Bertrand fashion, they would maximize the H type's profits in the equilibrium. We also consider the FIB as the objective of the mediator in the Appendix.

The basic role of the mediator is to communicate with both the principal and the agent, and disclose information at a proper time. Thus, we can achieve immediate separation of the two types of the principal, but this information can be concealed from the agent, at least temporarily. In this way, the agent's IR and IC constraints only need to hold in expectation, though different principals offer different contracts.

Science Exchange, an online platform that connects scientific researchers with experimental service providers, is an example of a mediator. The researcher corresponds to the principal in our model, while the service provider is the agent. The platform has its own system to verify the qualification of service providers, so one shall not worry much about their unobservable abilities. The major problem for service providers and for the platform is that the quality of ideas brought in by researchers is, by their nature, hard to evaluate. Furthermore, moral hazard on the part of service providers is likely to be important, given the uncertainties associated with the research process. The role played by the platform is not only to reduce search and transaction costs, but also to design contracts, protect intellectual properties and confidentiality, and facilitate communications.[21] Hence, researchers' confidential information can be concealed from service providers, and disclosed to them only when necessary.

Let us now formally examine the mediator's mechanism design problem. A (direct) mechanism, $\mathcal{M} = \{\mathcal{C}^H, \mathcal{C}^L\}$, is an extensive form game containing a menu of contracts. A contract, $\mathcal{C}^\theta = \{N^\theta, W^\theta, \boldsymbol{b}^\theta\}$, is a triple as before, where $\theta \in \Theta$, $N^\theta \in \mathbb{N}_0$, $W^\theta \in \mathbb{R}$, and $\boldsymbol{b}^\theta \in \mathbb{R}^{N^\theta}$.

Given a mechanism, $\mathcal{M} = \{\mathcal{C}^H, \mathcal{C}^L\}$, the agent and the principal simultaneously make their participation decision, and the principal reports her type to the mediator. If both parties accept the mechanism and the principal reports $\theta \in \Theta$, then the mediator implements the contract $\mathcal{C}^\theta$ in the following way:

- In any period $n \leq N^\theta$ (before success), the mediator will recommend the agent to work. If the project succeeds in that period, the principal will transfer b_n^θ to the agent, and the game ends; otherwise, there is no payment.
- In the period $n = N^\theta + 1$ (before success), the mediator will recommend the agent not to work, and the game ends.
- The principal will transfer W^θ, which is measured in time-zero discounted value, to the agent when the game ends.

The agent's action set is the same as before. He can choose whether to follow the mediator's recommendations (i.e. whether to exert effort), which is unobservable to both the principal and the mediator.

Before we go any further, we should explain the space of mechanisms that we study. They are direct mechanisms with two restrictions. First, recommendations for delaying a project temporarily are not considered, but this is without loss for the purpose to find the *optimal mechanism for the H type principal.* Second, random recommendations are not allowed. It has some loss. However, when the period length shrinks, we show that the H type principal

[21] See `https://www.scienceexchange.com/trust` and other sites under the same domain for details.

can obtain a payoff that converges to her FIB surplus via pure recommendations. In the Appendix, we will introduce random recommendations, through which the H type's payoff could be further improved for any fixed $\Delta > 0$.

Now, consider the IR and IC constraints in this new environment. We say a mechanism $\mathcal{M} = \{\mathcal{C}^H, \mathcal{C}^L\}$ is *feasible*, if both types of the principal and the agent are willing to participate, both types of the principal are willing to report truthfully, and the agent is willing to follow recommendations. In other words, a feasible mechanism satisfies the following IR and IC constraints.

The type $\theta \in \Theta$ principal's IR constraint is

$$\Pi^\theta = -W^\theta + \sum_{n=1}^{N^\theta} \delta^n f_{n-1}^\theta(q_0) q_n^\theta \lambda^\theta (h - b_n^\theta)\Delta \geq 0. \qquad IR_\theta(\mathcal{M})$$

The agent's IR constraint is

$$W_A = \sum_{\theta \in \Theta} \beta_0^\theta [-W^\theta + \sum_{n=1}^{N^\theta} \delta^n f_{n-1}^\theta(q_0)(q_n^\theta \lambda^\theta b_n^\theta - c)\Delta] \geq 0, \qquad IR_A(\mathcal{M})$$

where $\beta_0^H = \beta_0$ and $\beta_0^L = 1 - \beta_0$.

The type $\theta \in \Theta$ principal's IC constraint is

$$-W^\theta + \sum_{n=1}^{N^\theta} \delta^n f_{n-1}^\theta(q_0) q_n^\theta \lambda^\theta (h - b_n^\theta)\Delta \geq -W^{\theta'} + \sum_{n=1}^{N^{\theta'}} \delta^n f_{n-1}^\theta(q_0) q_n^\theta \lambda^\theta (h - b_n^{\theta'})\Delta, \quad IC_\theta^{\theta'}(\mathcal{M})$$

where $\theta' \in \Theta \backslash \{\theta\}$.

The above constraints are straightforward extensions from the signaling game. The only difference is that the transfers now depend on the reported type.

The agent's IC constraints are more involved, since signaling takes place during the experimentation. Suppose the H type experiments longer than the L type, i.e. $N^H \geq N^L$.[22] Since recommendations for both types are the same until N^L, the agent's beliefs about the type of the principal and the state of the project evolve in the same way as in the pooling equilibrium. However, if $N^H > N^L$ and the agent is recommend to work in the period $N^L + 1$, his posterior belief about the type of principal being H will jump to 1, but his posterior belief

[22]When $N^H < N^L$, we can define the IC constraints for the agent in the same way, though it is irrelevant for the optimal mechanism for the H type principal.

about the state of the project being G will jump down, since failures are worse news about the state being G for the H type than for the L type. Specifically, the agent's IC constraints are

$$\begin{cases} \nu_n \chi_n^H + (1-\nu_n)\chi_n^L \geq 0, & \text{for } 1 \leq n \leq N^L \ , \\ \chi_n^H \geq 0, & \text{for } N^L + 1 \leq n \leq N^H \text{ if } N^H > N^L, \end{cases} \qquad IC_A(\mathcal{M})$$

where $\nu_n = q_n\beta_n + (1-q_n)\beta_0$ is the posterior belief about the principal's type being H after $n-1$ failures, and χ_n^θ is the difference of the expected payoffs for the agent between always following recommendations and shirking in period n, but following all remaining recommendations thereafter, conditional on the principal's type is θ. Hence, $\chi_n^\theta = \sum_{s=n}^{N^\theta} \delta^{s-n} f_{s-n}^\theta(q_n^\theta)(q_s^\theta \lambda^\theta b_s^\theta - c)\Delta - \sum_{s=n+1}^{N^\theta} \delta^{s-n} f_{s-n-1}^\theta(q_n^\theta)(q_{s-1}^\theta \lambda^\theta b_s^\theta - c)\Delta$.

Remark. *Clearly, the equilibrium contracts in the signaling game remain feasible. Moreover, Proposition 1 implies there is no feasible mechanism that implements the FIB. Thus, allowing for a mediator cannot help the H type achieve her FIB surplus.*

We now show how a mediator can achieve a strictly higher payoff for the H type than both the BPEH and the BSEH by relaxing only the IR constraint of the agent.

Consider the feasible mechanisms in which the contract $\mathcal{C}^L$ recommends not experimenting at all. Hence, the agent would know the type of the principal before he conducts the project, but after he accepts the offer. The reason that a mediator can achieve a strictly higher payoff for the H type than the BPEH is straightforward. First, by recommending not experimenting for the L type, the total social surplus increases. Second, the bonus payments that the H type needs to incentivize the agent are less without the L type pooling during the experimentation. For the BSEH, let us consider the equilibrium contracts in which the H type leaves rents to the agent. Now, the mediator can slightly lower the lump-sum transfers in both contracts by the same amount, while keeping the agent's IR constraint satisfied in expectation before he knows the type of the principal. This will not change the incentive of the L type, but increase the payoffs of both types. Thus, separation is less costly when it occurs right after the offer acceptance. Hence, the mediator can achieve a strictly higher payoff for the H type than the BSEH.

The mediator can further improve the H type's payoff by also relaxing the IC constraints of the agent. This requires that the mediator also recommends experimenting on the L type project for some time. Thus, separation takes place in the period when recommendations differ for different types of principal.

In the quasi-linear environment, the total surplus generated in any feasible mechanism is shared by both types of principal and the agent. Let Π^θ be the payoff obtained by the $\theta \in \Theta$ type of principal, and W_A be the payoff obtained by the agent. Then, we have

$$\beta_0 V^H(N^H) + (1-\beta_0)V^L(N^L) = \beta_0 \Pi^H + (1-\beta_0)\Pi^L + W_A.$$

On the one hand, feasibility requires that both types of principal and the agent must obtain positive payoffs. On the other hand, the value generated by the L type, $V^L(N^L)$, is strictly decreasing in N^L for $N^L \geq 1$, and the value generated by the H type, $V^H(N^H)$, is maximized at $N^H = N^H_*$. Thus, for any $N^L \geq 1$, we have

$$\begin{aligned}\Pi^H =& V^H(N^H) + \frac{1-\beta_0}{\beta_0}(V^L(N^L) - \Pi^L) - \frac{1}{\beta_0}W_A \\ \leq & V^H(N^H_*) + \frac{1-\beta_0}{\beta_0}V^L(1).\end{aligned}$$

That means, in any feasible mechanism with $N^L \geq 1$, the payoff of H type principal has an upper bound, which is strictly positive for a small Δ.

We now show that there exists a feasible mechanism with $N^L \geq 1$ that achieves that upper bound for the H type's payoff. Clearly, such a mechanism must (1) satisfy feasibility, (2) give both the agent and the L type principal 0 payoffs, and (3) implement the H type project efficiently ($N^H = N^H_*$) and the L type project almost efficiently ($N^L = 1$). Thus, the last condition determines the efficiency of both projects. Together with the second condition, they determine the lump-sum transfers,

$$W^L = \delta q_0 \lambda^L (h - b^L_1)\Delta, \tag{4}$$

$$W^H = -[V^H(N^H_*) + \frac{1-\beta_0}{\beta_0}V^L(1)] + \sum_{n=1}^{N^H_*} \delta^n f^H_{n-1}(q_0) q^H_n \lambda^H (h - b^H_n)\Delta. \tag{5}$$

Now, we show how to use bonus payments to incentivize the agent (to work) and the L type principal (to report truthfully) at the same time.

From period 2, the agent will learn that the principal's type is H, if the mediator still recommends experimenting on the project. Thus, his IC constraints from period 2 to N^H_* are the same as in the BSEH. However, in the first period, the agent's belief about the principal being type H is β_0, and his IC constraint can be simplified to

$$q_0(\beta_0\lambda^H b_1^H + (1-\beta_0)\lambda^L b_1^L) - c \geq \beta_0 q_0 \lambda^H \sum_{s=2}^{N_*^H} \delta^{s-1}(1-\lambda^H\Delta)^{s-2}(\lambda^H b_s^H - c)\Delta. \tag{6}$$

Note that the LHS of the inequality (6) is linear and strictly increasing in b_1^L. Thus, given $\{b_n^H\}_{1\leq n\leq N_*^H}$, the inequality determines a lower bound for b_1^L.

The L type's IC constraint is $-W^H + \sum_{n=1}^{N_*^H} \delta^n f_{n-1}^L(q_0)q_n^L\lambda^L(h-b_n^H)\Delta \leq 0$, i.e.

$$V^H(N_*^H) + \frac{1-\beta_0}{\beta_0}V^L(1) - \sum_{n=1}^{N_*^H}\delta^n\eta_n(h-b_n^H)\Delta \leq 0. \tag{7}$$

The LHS of the inequality (7) is also linear and strictly increasing in b_1^H. Hence, given $\{b_n^H\}_{2\leq n\leq N_*^H}$, this inequality determines an upper bound for b_1^H.

Thus, distinct bonus payments for the first experiment can be used to incentivize both the L type principal and the agent. We can make b_1^H small to prevent the L type from mimicking the H type, but keep b_1^L large such that the expected bonus payment is high enough for the agent to work. This virtually resolves the conflicts between signaling and providing incentives to the agent. When the period length is small, the inefficiency is negligible, and the H type principal would obtain approximately her FIB payoff. Thus, we have the following result:

Proposition 7. *There exists a $\bar{\Delta} > 0$, such that for any $\Delta \in (0, \bar{\Delta})$, there exists an optimal mechanism for the H type, in which the H type obtains a payoff of $V^H(N_*^H) + \frac{1-\beta_0}{\beta_0}V^L(1)$. As $\Delta \to 0$, the H type's payoff converges to her FIB surplus.*

In this mechanism, the H type principal can obtain a payoff close to her FIB surplus for a small Δ; while the L type principal and the agent are left with no rents. The L type only needs to conduct one experiment, but provides high enough incentives, to compensate the H type's low bonus, for the agent to exert efforts in the first period. The reason that the H type needs to provide a very low bonus for the first experiment, is to make sure that the L type has no interest to pretend the H type. The chance to succeed in the first experiment for the H type is larger than that for the L type. Thus, the two contracts are like two bets on the success of the first experiment. The H type's contract is a bet with a higher reward, as well as a higher price, on the success of the first experiment; while the L type's contract is a bet with a lower reward, as well as a lower price. Hence, the L type would like to report truthfully, and provide a high bonus for the first experiment. Thus, the agent is also given enough expected bonus to work. He would gain some rents from working for the H type, and lose some rents from working for the L type, but break even on average.

We also introduce random recommendations and study the implementation problem of the FIB in the Appendix. We show that, for any fixed $\Delta > 0$, the FIB can be virtually implemented by recommending the agent that experiments once with some probability $\mu \in (0, 1)$ for the L type. Thus, introducing a mediator with a menu of contracts helps the H type keep virtually all the surplus of her innovation, while the L type who has an inferior innovation is left with nothing. Hence, this mechanism provides proper incentives for innovators to invest in better ideas ex ante.[23]

The reader may ask, can we achieve approximate efficiency without a mediator? Suppose we consider a game where the principal can propose an arbitrary menu of contracts, with the provision that the exact contract that is to be implemented will depend upon a private message sent by the principal. Furthermore, the menu provides a disclosure date, when it will be revealed to the agent which contract has been chosen. Is there a PBE where both types of principal propose the above mechanism? This is an appropriate generalization, to our context, of the question that has been examined in the three-stage mechanism proposing game of Maskin and Tirole (1992). The key problem here is ensuring that deviations, by either type of principal to a different menu of contracts, are unprofitabe. In our context, the set of possible deviating contracts is extremely large and complex, and we have been unable to show that every such deviation is unprofitable. More specifically, we need to show that for any other menu, there exists a belief of the agent about the type of the principal that prevents a profitable deviation by either type. Our Lemma 3 shows that a pessimistic belief that assigns probability 1 about the principal being the L type can prevent a profitable deviation in the original signaling game, but such a simple belief does not work in the menu proposing game.[24] An explicit construction of beliefs that prevents profitable deviations appears to be intractable. In this context, the paper by Wagner, Mylovanov, and Tröger (2015) is relevant. They examine moral hazard in a static setting and show that when the FIB is feasible, it is an equilibrium outcome. However, they do not study the case where

[23] This mechanism is also approximately ex ante optimal for the principal.

[24] For example, consider a menu of contracts in which one of the contracts is a "bet" (between the two types for the first experiment) that would not be taken by the L type. If the agent believes it is the L type who proposes such a mechanism, he would believe that the principal would not commit to that "bet" contract. However, such a contract might be very attractive to the H type.

the FIB is not feasible. The key feature of the present model is that the FIB is not feasible, even though we can approximate it.[25] Consequently, their results and methods cannot be applied in our context.

7 Conclusion

We analyze a model where an informed principal engages an agent to explore the viability of her project. The principal has private information about the quality of her project. Thus, in addition to providing incentives to the agent to experiment on the project, the high type principal has to convince the agent of her project's quality. We examine the best outcome for the high type principal in equilibrium.

When there is a large prior probability that the project is of low quality, the best equilibrium for the high type principal is a separating equilibrium. The high type principal separates himself from the low type by leaving rents to the agent, and also by terminating the project inefficiently early. When the prior probability of a low quality project is small, the best equilibrium for the high type is a pooling equilibrium. We find that the best equilibrium, from the point of view of rewarding innovations, is either a separating equilibrium or a pooling equilibrium.

In neither equilibrium of the signaling game is the innovator with a superior project able to capture her contribution to social surplus. This leads us to study the role of a mediator. The mediator offers a menu of two contracts, and does not disclose the type of the principal to the agent unless it is necessary to do so. We find that an opaque contract is able to approximately implement the optimal outcome, in terms of both inducing efficient experimentation, and ensuring high rewards for the high type principal. It is necessary to induce experimentation by the low type only for a single period, and as the period length becomes small, this is approximately efficient.

[25] Recall that the full information benchmark requires immediate termination of the L types project and a zero payoff for this type. This is not feasible.

Appendices

A Temporary Delay of the Project

In this section, we formalize the idea of delaying the project temporarily. Lemma 6 will show that it is never optimal for the H type to delay her project temporarily.

Given a contract $C = \{N, W, \boldsymbol{b}\}$, and the agent's belief β_0 about the principal being H, after accepting the contract, the agent would solve a dynamic decision problem, i.e. when to exert efforts. Suppose the agent's optimal action plan is to exert efforts only for periods in $R \subseteq \{1, 2, \cdots, N\}$, then the contract must satisfy the agent's IC constraints[26] for those periods, i.e. for all $1 \leq n \leq \#R$ ($\#R$ is the size of R),

$$\sum_{s=n}^{\#R} \Pi_{j=n+1}^{s} \delta_j f_{s-n}(q_n, \beta_n)(q_s \lambda_s b_{r_s} - c)\Delta \geq \sum_{s=n+1}^{\#R} \Pi_{j=n+1}^{s} \delta_j f_{s-n-1}(q_n, \beta_n)(q_{s-1} \lambda_{s-1} b_{r_s} - c)\Delta$$

where $\delta_j = \delta^{r_j - r_{j-1}}$, r_j is the j-th minimal item in R, and $r_0 = 0$.

In such a contract, $\#R \leq N$ is the actual number of experiments. Everything remains the same in the IC constraints, except for the discount factor for each experiment. Therefore, we can emphasize the contract in a different way: the principal has to incentivize the agent to work for every experiment, but she can choose how long to suspend the project between two consecutive experiments, or equivalently, choose the discount factors between every two consecutive experiments.

From the perspective of the induced actions, we define the following direct contracts. A *direct contract* is a quadruple, $\mathscr{C} = \{N, W, \boldsymbol{b}, \boldsymbol{\delta}\}$, where $N \in \mathbb{N}_0$ is the number of experiments, W is still the lump-sum time zero transfer from the principal to the agent, $\boldsymbol{b} \in \mathbb{R}^N$ is the vector of bonus scheme for success that *must* incentivize the agent to work for *all* N experiments, and $\boldsymbol{\delta} \in [0, \delta]^N$ is the vector of discount factors between two consecutive experiments, and for $1 \leq n \leq N$, $\delta_n = \delta^k$ for some $k \in \mathbb{N} \cup \{\infty\}$, therefore $0 \leq \delta_n \leq \delta$. The principal

[26]The notations are borrowed from the section 5.3. The constraints are the agent's IC constraints when the agent believes that the H and L type pool with the prior belief β_0. The degenerate case $\beta_0 = 0$ (or $\beta_0 = 1$) means the agent believes the principal is the L (or H) type. The bonus scheme also needs to satisfy the IC constraints that the agent dose not work in the periods $n \in \{1, 2, \cdots, N\} \backslash R$. We omit the requirement of b_n for $n \in \{1, 2, \cdots, N\} \backslash R$, since they are payoff irrelevant in the equilibrium as long as they are low enough to discourage the agent to experiment.

can freely choose any discount factors and any number of experiments. However, given the chosen discount factors and the number of experiments, the bonus payments must satisfy the IC constraints for the agent for all N experiments.

With a direct contract $\mathscr{C} = \{N, W, \boldsymbol{b}, \boldsymbol{\delta}\}$, nothing changes except for discount factors. We can easily rewrite the payoffs of all parties, and their IR and IC constraints by properly adjusting their discount factors. For example, given the agent's belief about the principal's type being H is β_0, his IC constraints become, for $1 \leq n \leq N$

$$\sum_{s=n}^{N} \Pi_{j=n+1}^{s} \delta_j f_{s-n}(q_n, \beta_n)(q_s \lambda_s b_s - c)\Delta \geq \sum_{s=n+1}^{N} \Pi_{j=n+1}^{s} \delta_j f_{s-n-1}(q_n, \beta_n)(q_{s-1} \lambda_{s-1} b_s - c)\Delta.$$

Therefore, we can use direct contracts to analyze the signaling game and the third-party mechanism design problem. However, Lemma 6 implies that delaying the H type project temporarily never happens in the best equilibrium for the H type. As we will see later, this is also true for the optimal mechanism for the H type.

B Three Useful Lemmas

We provide three lemmas that we need to prove other results in the Appendix. Lemma 6 shows that the IC constraints for the agent always bind whenever there are signaling concerns, and temporary delay of the project is not a good way to separate the H type. Lemma 7 characterizes the recursive relations between two consecutive bonuses when the IC constraints for the agent bind. Lemma 8 explores the limit of the payoff difference between the two types of principal from the same contract when the period length shrinks, and shows that it is monotonic in the belief about the type of principal when the stopping time is less than the efficient stopping time for the H type.

B.1 Lemma 6

Lemma 6. *Fix any $N > 1$, let $\boldsymbol{b}^*, \boldsymbol{\delta}^*$ be a solution to the following program:*

$$\max_{\boldsymbol{b},\{0\le\delta_n\le\delta\}_{n=1}^N} \sum_{n=1}^{N} \Pi_{j=1}^n \delta_j [f_{n-1}^H(q_0) q_n^H \lambda^H \Delta - f_{n-1}^L(q_0) q_n^L \lambda^L \Delta](h - b_n)$$

$$s.t. \sum_{s=n}^{N} \Pi_{j=n+1}^s \delta_j f_{s-n}(q_n, \beta_n)(q_s \lambda_s b_s - c)\Delta \ge \sum_{s=n+1}^{N} \Pi_{j=n+1}^s \delta_j f_{s-n-1}(q_n, \beta_n)(q_{s-1}\lambda_{s-1} b_s - c)\Delta$$

Let $m^* + 1 = \min\{n : \delta_n^* = 0\}$, or $m^* = N$ if $\min\{n : \delta_n^* = 0\} = \varnothing$, then for any $1 \le n \le m^*$, the above constraints bind and $\delta_n^* = \delta$.

The constraints are the agent's IC constraints. The objective function is the payoff difference between the H and L type when they propose the same contract. To find the best equilibrium for the H type, the lump-sum payment W is pinned down by either a binding IC constraint for the L type as in the BSEH or a binding IR constraint for the agent as in the BPEH. In other words, the H type has no reason to leave rents to the agent, unless she wants to prevent the L type mimicking her. If the L type obtains 0 payoff, then the objective function is equal to the H type's payoff. If the agent obtains 0 payoff, then the objective function represents the part of the H type's payoff that involves the bonus payments.

We show that given any termination date N, the IC constraints for the agent must bind for all periods even when arbitrary delay of the project is allowed. Moreover, we also show that temporary delay of the project is never optimal for the H type. Furthermore, the result holds with respect to the following transformation of the objective function: $H \longmapsto \xi H + G(\delta_1, \dots, \delta_n)$, where $\xi \in \mathbb{R}_{++}$, and $G(\delta_1, \dots, \delta_n)$ is linear in every δ_n and independent with every b_n for $1 \le n \le N$.

Proof. If $m^* = 1$, the statement is trivial. Consider $m^* > 1$. From $m^* + 1$, all terms in the objective function is 0, we can rewrite the program in which the last term is in the m^* period.

Note that $\eta_n := f_{n-1}^H(q_0) q_n^H \lambda^H - f_{n-1}^L(q_0) q_n^L \lambda^L$, and $\eta_n \Delta$ is the difference of the probabilities to succeed for the n-th experiment between the H and L type. It is easy to see that when $n \le n^* = \lfloor \frac{\log \lambda^H - \log \lambda^L}{\log(1-\lambda^L\Delta) - \log(1-\lambda^H\Delta)} \rfloor + 1$, $\eta_n \ge 0$. When $n > n^*$, $\eta_n < 0$.

Therefore, for $1 \le n \le n^*$, we would like to set b_n as small as possible. Since b_n is bounded from below, the IC constraints for $1 \le n \le n^*$ must bind.

Given the binding constraints for $1 \leq n \leq n^*$, consider $n = n^* + 1$. Note that a change of b_{n^*+1}, call it y_{n^*+1}, will change all the previous bonuses according to the binding constraints from period 1 to period n^*.

We can show that the change of b_n for $1 \leq n \leq n^*$, call it y_n, from y_{n^*+1} is

$$y_n = \Pi_{j=n+1}^{n^*+1} \delta_j [\lambda^H - (\lambda^H - \lambda^L) \frac{\lambda^L (1 - \beta_{n^*})}{\lambda_{n^*}}] \Delta y_{n^*+1}.$$

Clearly, $\zeta_{n^*} := \lambda^H - (\lambda^H - \lambda^L) \frac{\lambda^L(1-\beta_{n^*})}{\lambda_{n^*}}$ is a positive constant. When $\beta_0 = 1$, it is λ^H, and when $\beta_0 = 0$, it is λ^L. When $\beta_0 \in (0, 1)$, it is in between λ^L and λ^H, and converges to λ^L when n^* goes to infinity.

Thus, the change of the objective function from y_{n^*+1} is

$$\begin{aligned} -\sum_{n=1}^{n^*+1} \Pi_{j=1}^{n} \delta_j \eta_n \Delta y_n &= -\sum_{n=1}^{n^*} \Pi_{j=1}^{n^*+1} \delta_j \eta_n \Delta \zeta_{n^*} \Delta y_{n^*+1} - \Pi_{j=1}^{n^*+1} \delta_j \eta_{n^*+1} \Delta y_{n^*+1} \\ &= -\Pi_{j=1}^{n^*+1} \delta_j y_{n^*+1} (\zeta_{n^*} \Delta \sum_{n=1}^{n^*} \eta_n \Delta + \eta_{n^*+1} \Delta) \end{aligned} \tag{8}$$

Note that the distribution on success for L type first-order stochastic dominates that for H type; the H project is more likely to succeed than the L project in the first $s \geq 1$ experiments. Formally, for any $s \geq 1$

$$\phi_s := \sum_{n=1}^{s} \eta_n \Delta = -[f_s^H(q_0) - f_s^L(q_0)] = q_0[(1 - \lambda^L \Delta)^s - (1 - \lambda^H \Delta)^s] > 0.$$

Therefore, the terms in the parenthesis of equation (8) becomes

$$\zeta_{n^*} \Delta q_0 [(1 - \lambda^L \Delta)^{n^*} - (1 - \lambda^H \Delta)^{n^*}] + q_0 [(1 - \lambda^H \Delta)^{n^*} \lambda^H \Delta - (1 - \lambda^L \Delta)^{n^*} \lambda^L \Delta]$$

which is large than $q_0(\lambda^H - \lambda^L)\Delta(1 - \lambda^H \Delta)^{n^*} > 0$ since $\zeta_{n^*} \geq \lambda^L$.

Therefore, to maximize the objective function, we would like to make b_{n^*+1} as small as possible. Thus, the IC constraints in period $n^* + 1$ also binds.

By induction, we can get a necessary condition for the solution of the program: the IC constraints to incentivize agent to work must bind in each period.

In addition, since every discount factor enter objective function linearly, we must have $\delta_n^* = \delta$

for $1 \leq n \leq m^*$. The reason is the following.

Consider the discount factor in $n+1$ period, δ_{n+1}. It will not affect any b_s for $s \geq n+1$, then δ_{n+1} enters into the terms linearly beyond period $n+1$ as a discount factor. It will linearly affect every b_s for $s \leq n$, but they are not discounted by δ_{n+1}. Therefore, the discount factor δ_{n+1} enters objective function linearly.

Thus, the optimal discount factor must be a corner solution, i.e. $\delta_n^* \in \{\delta, 0\}$ for any $1 \leq n \leq m^*$. Since $m^*+1 = \min\{n : \delta_n^* = 0\}$, we have $\delta_n^* = \delta$ for all $1 \leq n \leq m^*$. □

B.2 Lemma 7

Lemma 7. *Given the binding IC constraints for the agent when he has prior β_0 about the H type, i.e. for $1 \leq n \leq N$,*

$$\sum_{s=n}^{N} \delta^{s-n} f_{s-n}(q_n, \beta_n)(q_s \lambda_s b_s - c)\Delta = \sum_{s=n+1}^{N} \delta^{s-n} f_{s-n-1}(q_n, \beta_n)(q_{s-1}\lambda_{s-1} b_s - c)\Delta.$$

Then

$$q_n \lambda_n b_n - c = \delta(q_n \lambda_n b_{n+1} - c)$$

for $1 \leq n \leq N-1$, and

$$q_N \lambda_N b_N - c = 0.$$

Proof. Immediately, the binding IC constraint for the last period N implies

$$q_N \lambda_N b_N - c = 0.$$

For the IC constraint in period $N-1$, it gives us

$$q_{N-1}\lambda_{N-1} b_{N-1} - c + \delta f_1(q_{N-1}, \beta_{N-1})(q_N \lambda_N b_N - c) = \delta(q_{N-1}\lambda_{N-1} b_N - c).$$

Since the second term on the LHS is 0, it implies

$$q_{N-1}\lambda_{N-1} b_{N-1} - c = \delta(q_{N-1}\lambda_{N-1} b_N - c).$$

Suppose for all $m+1 \leq n \leq N-1$, we have $q_n \lambda_n b_n - c = \delta(q_n \lambda_n b_{n+1} - c)$. Consider the IC constraint in the period m,

$$\sum_{s=m}^{N-1} \delta^{s-m} f_{s-m}(q_m, \beta_m)(q_s \lambda_s b_s - c)\Delta = \sum_{s=m+1}^{N} \delta^{s-m} f_{s-m-1}(q_m, \beta_m)(q_{s-1}\lambda_{s-1} b_s - c)\Delta.$$

Note that the LHS has no terms in the period N since $q_N \lambda_N b_N - c = 0$. Relabeling the terms on the LHS gives us

$$\sum_{s=m+1}^{N} \delta^{s-m-1} f_{s-m-1}(q_m, \beta_m)(q_{s-1}\lambda_{s-1} b_{s-1} - c)\Delta = \sum_{s=m+1}^{N} \delta^{s-m} f_{s-m-1}(q_m, \beta_m)(q_{s-1}\lambda_{s-1} b_s - c)\Delta.$$

Since $q_n \lambda_n b_n - c = \delta(q_n \lambda_n b_{n+1} - c)$ for $m+1 \leq n \leq N-1$, it is clear that the recursive relation remains valid for $n = m$, which concludes this lemma. □

The following observations will be useful for other proofs:

1. By iteration, for $1 \leq n \leq N$

$$b_n = \delta^{N-n} \frac{c}{q_N \lambda_N} + (1-\delta) \sum_{s=n}^{N-1} \delta^{s-n} \frac{c}{q_s \lambda_s};$$

2. For $n < N$, $q_n \lambda_n b_n > c$;

3. For $n < N$, $\delta b_{n+1} \leq b_n \leq b_{n+1}$, and both inequalities are strict when $\delta < 1$.

B.3 Lemma 8

With a slight abuse of notation, we denote q_t, λ_t, η_t, ϕ_t, $f_t^\theta(q)$ and $f_t(q, \beta)$, where $\theta \in \Theta$, as the limit of q_n, λ_n, η_n, ϕ_n, $f_n^\theta(q)$ and $f_n(q, \beta)$ when $\Delta \to 0$ and $n\Delta \to t$.

Lemma 8. *Let $T = N\Delta$. Define*

$$H^{\beta_0}(N) := \sum_{n=1}^{N} \delta^n \eta_n (h - b_n)\Delta,$$

where $b_n = \delta^{N-n} \frac{c}{q_N \lambda_N} + (1-\delta) \sum_{s=n}^{N-1} \delta^{s-n} \frac{c}{q_s \lambda_s}$. Then

$$\lim_{\Delta \to 0} H^{\beta_0}(N) = H_0^{\beta_0}(T) := e^{-\rho T}(h - \frac{c}{q_T \lambda_T})\phi_T + \rho \int_0^T e^{-\rho t}(h - \frac{c}{q_t \lambda_t})\phi_t \,\mathrm{d}t.$$

In addition, $H_0^{\beta_0}(T)$ is strictly increasing in $\beta_0 \in [0,1]$ for any $T \in (0, T_*^H]$.

Proof. Note that

$$\begin{aligned}
\sum_{n=1}^{N} \delta^n \eta_n h \Delta &= \sum_{n=1}^{N} \delta^n h(\phi_n - \phi_{n-1}) = \delta^N h \phi_N + (1-\delta) \sum_{n=1}^{N-1} \delta^n h \phi_n, \\
\sum_{n=1}^{N} \delta^n \eta_n b_n \Delta &= \delta^N \frac{c}{q_N \lambda_N} \phi_N + (1-\delta) \sum_{n=1}^{N} \sum_{s=n}^{N-1} \delta^s \eta_n \frac{c}{q_s \lambda_s} \Delta \\
&= \delta^N \frac{c}{q_N \lambda_N} \phi_N + (1-\delta) \sum_{s=1}^{N-1} \sum_{n=1}^{s} \delta^s \eta_n \frac{c}{q_s \lambda_s} \Delta \\
&= \delta^N \frac{c}{q_N \lambda_N} \phi_N + (1-\delta) \sum_{s=1}^{N-1} \delta^s \frac{c}{q_s \lambda_s} \phi_s.
\end{aligned}$$

Hence, we have

$$H^{\beta_0}(N) = \delta^N (h - \frac{c}{q_N \lambda_N}) \phi_N + \frac{(1-\delta)}{\Delta} \sum_{n=1}^{N-1} \delta^n (h - \frac{c}{q_n \lambda_n}) \phi_n \Delta.$$

Let $T = N\Delta$, we have

$$\lim_{\Delta \to 0} \delta^N (h - \frac{c}{q_N \lambda_N}) \phi_N = e^{-\rho T} (h - \frac{c}{q_T \lambda_T}) \phi_t \text{ and } \lim_{\Delta \to 0} \frac{1-\delta}{\Delta} = \rho.$$

In addition, we can show that for any $t \in [0, T]$,

$$|\delta^n (h - \frac{c}{q_n \lambda_n}) \phi_n - e^{-\rho t} (h - \frac{c}{q_t \lambda_t}) \phi_t| \underset{\Delta \to 0}{=} \mathcal{O}(\Delta),$$

where $t = n\Delta$. Thus,

$$\lim_{\Delta \to 0} \sum_{n=1}^{N-1} |\delta^n (h - \frac{c}{q_n \lambda_n}) \phi_n - e^{-\rho t} (h - \frac{c}{q_t \lambda_t}) \phi_t| \Delta = 0.$$

Hence, we have

$$\lim_{\Delta \to 0} \sum_{n=1}^{N-1} \delta^n (h - \frac{c}{q_n \lambda_n}) \phi_n \Delta = \lim_{\Delta \to 0} \sum_{n=1}^{N-1} e^{-\rho t} (h - \frac{c}{q_t \lambda_t}) \phi_t \Delta = \int_0^T e^{-\rho t} (h - \frac{c}{q_t \lambda_t}) \phi_t \, \mathrm{d}t,$$

where the last equality comes from Riemann integral. Therefore, $\lim_{\Delta \to 0} H^{\beta_0}(N) = H_0^{\beta_0}(T)$.

Moreover, we have

$$\frac{\mathrm{d}H_0^{\beta_0}(T)}{\mathrm{d}\beta_0} = -[e^{-\rho T}\frac{\mathrm{d}\frac{c}{q_T\lambda_T}}{\mathrm{d}\beta_0}\phi_T + \rho\int_0^T e^{-\rho t}\frac{\mathrm{d}\frac{c}{q_t\lambda_t}}{\mathrm{d}\beta_0}\phi_t\,\mathrm{d}t].$$

Note that for all $0 \leq t \leq T_*^H$,

$$\frac{\mathrm{d}q_t\lambda_t}{\mathrm{d}\beta_0} = \frac{f_t^H(q_0)f_t^L(q_0)}{f_t^2(q_0,\beta_0)}(q_t^H\lambda^H - q_t^L\lambda^L) > 0,$$

since $q_t^H\lambda^H \geq \frac{c}{h} \geq q_t^L\lambda^L$ and at least one inequality is strict. Thus, for all $0 \leq t \leq T_*^H$, $\frac{\mathrm{d}\frac{c}{q_t\lambda_t}}{\mathrm{d}\beta_0} < 0$. Therefore, $\frac{\mathrm{d}H_0^{\beta_0}(T)}{\mathrm{d}\beta_0} > 0$ for $T \in (0, T_*^H]$, since $\phi_t > 0$ for $t > 0$. □

Actually, we can show that when $\Delta \to 0$, $n\Delta \to t$ and $N\Delta \to T$, b_n converges to

$$b_t = e^{-\rho(T-t)}\frac{c}{q_T\lambda_T} + \rho\int_t^T e^{-\rho(\tau-t)}\frac{c}{q_\tau\lambda_\tau}\,\mathrm{d}\tau,$$

and $H_0^{\beta_0}(T) = \int_0^T e^{-\rho t}\eta_t(h-b_t)\,\mathrm{d}t$. The latter can be seen by differentiating the RHS,

$$\frac{\mathrm{d}\int_0^T e^{-\rho t}\eta_t(h-b_t)\,\mathrm{d}t}{\mathrm{d}T} = e^{-\rho T}\eta_T(h-b_T) - \int_0^T e^{-\rho t}\eta_t\frac{\mathrm{d}b_t}{\mathrm{d}T}\,\mathrm{d}t = e^{-\rho T}\frac{\mathrm{d}(h-\frac{c}{q_T\lambda_T})\phi_T}{\mathrm{d}T},$$

and applying the fundamental theorem of calculus and integration by parts,

$$\begin{aligned}\int_0^T e^{-\rho t}\eta_t(h-b_t)\,\mathrm{d}t &= \int_0^T e^{-\rho t}\frac{\mathrm{d}(h-\frac{c}{q_t\lambda_t})\phi_t}{\mathrm{d}t}\,\mathrm{d}t \\ &= e^{-\rho T}(h-\frac{c}{q_T\lambda_T})\phi_T + \rho\int_0^T e^{-\rho t}(h-\frac{c}{q_t\lambda_t})\phi_t\,\mathrm{d}t.\end{aligned}$$

Hence, $H_0^{\beta_0}(T) = \int_0^T e^{-\rho t}\eta_t(h-b_t)\,\mathrm{d}t$. Moreover, we have

$$\frac{\mathrm{d}H_0^{\beta_0}(T)}{\mathrm{d}T} = e^{-\rho T}[\eta_T h + \phi_T c - cq_0\frac{(\lambda^H-\lambda^L)\lambda_0 e^{-(\lambda^L+\lambda^H)T}}{q_T\lambda_T(\beta_0\lambda^H e^{-\lambda^H T} + (1-\beta_0)\lambda^L e^{-\lambda^L T})}].$$

C Proof of Lemma 2 and Proposition 1

Proof. The worst contract, $C^{wt} = (N_*^H, W^{wt}, \boldsymbol{b^{wt}})$, for the L type solves Program I. Since $IR_A^H(N_*^H)$ bind,

$$-W = \sum_{n=1}^{N_*^H} \delta^n f_{n-1}^H(q_0)(q_n^H \lambda^H b_n - c)\Delta,$$

and the objective function in the Program I becomes

$$-\sum_{n=1}^{N_*^H} \delta^n \eta_n (h - b_n)\Delta + V^H(N_*^H).$$

According to Lemma 6, the solution to Program I must have binding constraints in $IC_A^H(N_*^H)$ for any $1 \leq n \leq N_*^H$, which gives us Lemma 2.

Furthermore, from Lemma 7, for $1 \leq n < N_*^H$,

$$q_n^H \lambda^H b_n - c = \delta(q_n^H \lambda^H b_{n+1} - c),$$

and

$$q_{N_*^H}^H \lambda^H b_{N_*^H} - c = 0.$$

Clearly, $0 < b_n \leq b_{N_*^H} \leq h$, and $q_n^H \lambda^H b_n - c > 0$ for all $n < N_*^H$. It means that the L type can obtain positive payment from the agent before experimentation, i.e. $-W^{wt} \geq 0$, and positive share of profits during experimentation, i.e. $\sum_{n=1}^{N_*^H} \delta^n f_{n-1}^L(q_0) q_n^L \lambda^L (h - b_n)\Delta \geq 0$. Additionally, at least one of the two payoffs is strictly positive. If $-W^{wt} = 0$, then $N_*^H = 1$. Thus, $b_{N_*^H} = c/q_0\lambda^H < h$ and the share of profits during experimentation is strictly positive.

Therefore, by choosing the worst contract for the L type, the L type can get strictly positive payoff, which gives us Proposition 1. □

D Proof of Lemma 3

Proof. Obviously, the L type cannot generate a strictly positive payoff. When the agent believes that the principal is the L type, he will either reject the contract which gives the agent negative payoff, or accept the contract which gives the L type principal negative payoff.

In either case, the L type principal cannot get strictly positive payoff.

Now, consider the H type. We will solve the best contract for the H type if the agent believes that she is a L type but is still willing to accept the contract. The best contract will solve the following Program V:

$$\max_{N,W,\boldsymbol{b}} \Pi^H = -W + \sum_{n=1}^{N} \delta^n f_{n-1}^H(q_0) q_n^H \lambda^H (h - b_n)\Delta$$

$$\text{s.t.} \sum_{s=n}^{N} \delta^{s-n} f_{s-n}^L(q_n^L)(q_s^L \lambda^L b_s - c)\Delta \geq \sum_{s=n+1}^{N} \delta^{s-n} f_{s-n-1}^L(q_n^L)(q_{s-1}^L \lambda^L b_s - c)\Delta$$

$$W + \sum_{n=1}^{N} \delta^n f_{n-1}^L(q_0)(q_n^L \lambda^L b_n - c)\Delta \geq 0$$

The IR constraint for agent must bind, so the program becomes

$$\max_{N,\boldsymbol{b}} \Pi^H = \sum_{n=1}^{N} \delta^n \eta_n (h - b_n)\Delta + V^L(N)$$

$$\text{s.t.} \sum_{s=n}^{N} \delta^{s-n} f_{s-n}^L(q_n^L)(q_s^L \lambda^L b_s - c)\Delta \geq \sum_{s=n+1}^{N} \delta^{s-n} f_{s-n-1}^L(q_n^L)(q_{s-1}^L \lambda^L b_s - c)\Delta$$

According to Lemma 6 and 7, the IC constraints must bind, and

$$b_n = \delta^{N-n} \frac{c}{q_N^L \lambda^L} + (1-\delta) \sum_{s=n}^{N-1} \delta^{s-n} \frac{c}{q_s^L \lambda^L}$$

We will show that reducing one experiment makes the H type better off; thus, the termination date of the best contract for the H type is 0 and her maximal profit is 0. Let b_n^N be the bonus in period n, and $\Pi^H(N)$ be the profit of the H type, when the termination date is N. We can show that for $N \geq 1$

$$b_n^{N+1} - b_n^N = \delta^{N+1-n}\left(\frac{c}{q_{N+1}^L \lambda^L} - \frac{c}{q_N^L \lambda^L}\right).$$

Furthermore, for $N \geq 0$,

$$\Pi^H(N) - \Pi^H(N+1) = \delta^{N+1}\left[\left(\frac{c}{q_{N+1}^L \lambda^L} - \frac{c}{q_N^L \lambda^L}\right)\phi_N + f_N^H(q_0) q_{N+1}^H \lambda^H \left(\frac{c}{q_{N+1}^L \lambda^L} - h\right)\Delta\right] \geq 0,$$

since $\frac{c}{q_{N+1}^L \lambda^L} - \frac{c}{q_N^L \lambda^L} \geq 0$, $\phi_N \geq 0$ and $\frac{c}{q_{N+1}^L \lambda^L} \geq \frac{c}{q_0 \lambda^L} \geq h$. □

E Proof of Proposition 2

Proof. The H type's equilibrium contract, $C^{sep} = \{N^{sep}, W^{sep}, \boldsymbol{b^{sep}}\}$, in the BSEH solves Program II. We will solve the relaxed program without the agent's IR constraint, and verify it later.

First, in the relaxed program, $IC_L^H(N)$ must bind, otherwise we can decrease W without violating any other constraint, and obtain higher payoff.

The binding $IC_L^H(N)$ determines $W = \sum_{n=1}^{N} \delta^n f_{n-1}^L(q_0) q_n^L \lambda^L (h - b_n)\Delta$. Thus, the objective function becomes

$$\sum_{n=1}^{N} \delta^n \eta_n (h - b_n)\Delta$$

From Lemma 6, all IC constraints should bind. Furthermore, Lemma 7 gives an explicit characterization of the bonus scheme. In addition, Lemma 8 gives the limit form of the above objective function. Let $\Pi_{sep}^H(T)$ be the limit of the H type's payoff when her stopping time $N\Delta$ goes to T and Δ goes to 0. Thus, we can get[27]

$$\Pi_{sep}^H(T) = V_0^H(T) - V_0^L(T) - c\frac{\lambda^H - \lambda^L}{\lambda^H}[(1-q_0)\frac{1-e^{-(\rho+\lambda^L-\lambda^H)T}}{\rho+\lambda^L-\lambda^H} + q_0\frac{1-e^{-(\rho+\lambda^L)T}}{\rho+\lambda^L}].$$

Take the first derivative with respect to T,

$$\frac{\mathrm{d}\Pi_{sep}^H(T)}{\mathrm{d}T} = \frac{q_0}{\lambda^H}(\lambda^H h - c)e^{-(\rho+\lambda^L)T}[\lambda^H e^{(\lambda^L-\lambda^H)T} - (\lambda^H - \lambda^L)\frac{l^H}{l_0}e^{\lambda^H T} - \lambda^L].$$

Note that $[\lambda^H e^{(\lambda^L-\lambda^H)T} - (\lambda^H - \lambda^L)\frac{l^H}{l_0}e^{\lambda^H T} - \lambda^L]$ is strictly decreasing in T. At $T = 0$, it is $\frac{q_0}{\lambda^H}(\lambda^H h - c)(\lambda^H - \lambda^L)(1 - \frac{l^H}{l_0}) > 0$, and it goes to negative infinity when T goes to infinity. Therefore, there exists a T^{sep} such that

$$\lambda^H e^{-\lambda^H T^{sep}} - (\lambda^H - \lambda^L)\frac{l^H}{l_0}e^{(\lambda^H-\lambda^L)T^{sep}} - \lambda^L e^{-\lambda^L T^{sep}} = 0.$$

When $T < T^{sep}$, $\Pi_{sep}^H(T)$ is strictly increasing in T; when $T > T^{sep}$, $\Pi_{sep}^H(T)$ is strictly decreasing in T. T^{sep} maximizes $\Pi_{sep}^H(T)$.

[27] We assume that $\rho + \lambda^L - \lambda^H \neq 0$ so that the denominator is not 0. However, when $\rho + \lambda^L - \lambda^H = 0$, we can obtain the same results. The proofs in Section K of the Appendix also apply, where the term $\rho + \lambda^L - \lambda^H$ is in the denominator.

Moreover, since $\frac{l^H}{l_0} = e^{-\lambda^H T_*^H}$, we have

$$\lambda^H e^{-\lambda^H T_*^H} - (\lambda^H - \lambda^L)\frac{l^H}{l_0} e^{(\lambda^H - \lambda^L)T_*^H} - \lambda^L e^{-\lambda^L T_*^H} = \lambda^H (e^{-\lambda^H T_*^H} - e^{-\lambda^L T_*^H}) < 0.$$

Therefore, $T^{sep} < T_*^H$, there is always under experimentation. In addition, since $T^{sep} < T_*^H$ and b_t^{sep} is increasing, we have $b_t^{sep} \leq b_{T^{sep}}^{sep} < h$. Furthermore, $W^{sep} > 0$ since all $b_t^{sep} < h$.

In addition, we can show that T^{sep} is decreasing in λ^L. When $\lambda^L = 0$, $T^{sep} = \frac{1}{2}T_*^H$, and $W^{sep} = 0$. Therefore, $T^{sep} < \frac{1}{2}T_*^H$ in the main model since $\lambda^L > 0$.

The last thing is to check the agent's IR constraint. The agent's payoff in the limit is $V_0^H(T^{sep}) - \Pi_{sep}^H(T^{sep})$. Thus, it is equal to

$$\begin{aligned}
&V_0^L(T^{sep}) + c\frac{\lambda^H - \lambda^L}{\lambda^H}[(1-q_0)\frac{1 - e^{-(\rho+\lambda^L-\lambda^H)T^{sep}}}{\rho + \lambda^L - \lambda^H} + q_0 \frac{1 - e^{-(\rho+\lambda^L)T^{sep}}}{\rho+\lambda^L}] \\
=&\frac{\lambda^L}{\lambda^H}V_0^H(T^{sep}) + c\frac{\lambda^H - \lambda^L}{\lambda^H}(1-q_0)[\frac{1 - e^{-(\rho+\lambda^L-\lambda^H)T^{sep}}}{\rho + \lambda^L - \lambda^H} - \frac{1 - e^{-\rho T^{sep}}}{\rho}] \\
+&\frac{\lambda^L}{\lambda^H}q_0(\lambda^H h - c)[\frac{1 - e^{-(\rho+\lambda^L)T^{sep}}}{\rho+\lambda^L} - \frac{1 - e^{-(\rho+\lambda^H)T^{sep}}}{\rho+\lambda^H}] > 0.
\end{aligned}$$

All terms above are strictly positive. The first term is strictly positive since $0 < T^{sep} < T_*^H$. The other two terms are strictly positive since $\frac{1-e^{-xt}}{x}$ is strictly decreasing when $t > 0$. Therefore, the solution we find is the H type's equilibrium contract in the BSEH. In such an equilibrium, the agent gets rents more than $\frac{\lambda^L}{\lambda^H}V_0^H(T^{sep})$, i.e. $\frac{\lambda^L}{\lambda^H}$ proportion of the total surplus. The H type principal gets payoff less than $(1 - \frac{\lambda^L}{\lambda^H})V_0^H(T^{sep})$, i.e. $\frac{\lambda^H - \lambda^L}{\lambda^H}$ proportion of the total surplus. □

F Proof of Proposition 3

Proof. (a) First, we show that the BSEH survives the intuitive criterion.

For any contract C, let z be a response of the agent, which includes both the acceptance/rejection decision and the action plan conditional on acceptance, and $BR(C, \beta)$ be the set of best responses when his belief about the principal being H type is $\beta \in [0, 1]$. Let Π_{sep}^θ be the equilibrium payoff of the $\theta \in \Theta$ type principal in the BSEH, $\Pi^\theta(C, z)$ be the payoff of the $\theta \in \Theta$ type principal when she proposes contract C and the agent's response is

z, and

$$J(C) = \{\theta \in \Theta : \Pi^{\theta}_{sep} > \max_{z \in \underset{\beta \in [0,1]}{\cup} BR(C,\beta)} \Pi^{\theta}(C, z)\}$$

be the set of types whose payoffs from deviating to contract C is strictly less than their equilibrium payoffs in the BSEH, for any beliefs that the agent could possibly form. intuitive criterion allows the agent to form beliefs about the types only on $\Theta \backslash J(C)$. When $J(C) = \Theta$, the agent could form any beliefs.

For a contract C that the agent could "reasonably" form belief about the principal being type L, i.e $J(C) = \Theta$, $J(C) = \varnothing$ or $J(C) = \{H\}$, the minimal payoff the principal could get is no more than 0, according to Lemma 3. Thus, the BSEH cannot fail the intuitive criterion for such contract C, since both types could obtain non-negative payoffs in the BSEH and they would not deviate to C.

For a contract C that the agent could only "reasonably" form belief about the principal being type H, i.e $J(C) = \{L\}$, then the L type could not get positive payoff from deviating to contract C even when the agent believes her type is H. However, by definition, the BSEH already gives the highest payoff to the H type, subjected to the L type cannot get strictly positive payoff by proposing the same contract and making the agent believe her type is H. Hence, the H type would not deviate to C either. Thus, the BSEH cannot fail the intuitive criterion for such contract C.

Therefore, the BSEH must survive the intuitive criterion.

(b) Second, we show that any equilibria that give less equilibrium payoff for the H type than the BSEH fail the intuitive criterion.

For any such equilibrium, let $\bar{\Pi}^{\theta}$ be the equilibrium payoff for $\theta \in \Theta$ type principal. Then we have $\Pi^H_{sep} > \bar{\Pi}^H$, and $\bar{\Pi}^L \geq 0$. Let $\bar{\epsilon} \in (0, \Pi^H_{sep} - \bar{\Pi}^H)$. Consider the following contract $\bar{C}^H = \{T^{sep}, W^{sep} + \bar{\epsilon}, \boldsymbol{b^{sep}}\}$. The agent will always accept this contract since $W^{sep} + \bar{\epsilon} > 0$, given any belief about the principal. In addition, for any type of the principal, to experiment in every period is the best action plan for the principal, since the bonus scheme $\boldsymbol{b^{sep}}$ is an increasing sharing rule.

However, the L type will get strictly negative payoff from $\bar{C}^H$ even in this best case. Thus, the L type principal would never deviate to that contract $\bar{C}^H$, i.e. $L \in J(\bar{C}^H)$.

On the other hand, the H type can obtain a higher payoff $\Pi^H_{sep} - \bar{\epsilon} > \bar{\Pi}^H$ if the agent accepts the contract $\bar{C}^H$ and works until the termination date, which is the best response of the agent when he believes her type is H. Thus, the H type might deviate to the contract $\bar{C}^H$, i.e. $H \notin J(\bar{C}^H)$.

Hence, the agent would believe that it is the H type who deviates to $\bar{C}^H$, and the H type can obtain a higher payoff than $\bar{\Pi}^H$, given such belief. Therefore, the equilibrium with $\bar{\Pi}^H < \Pi^H_{sep}$ fails the intuitive criterion. □

G Proof of Proposition 4

Proof. The equilibrium contract, $C^{pl} = \{N^{pl}, W^{pl}, \boldsymbol{b^{pl}}\}$, in the BPEH solves Program III. We solve the relaxed program without $IR_L(N)$, and verify it later. In the relaxed program, $IR_A(N)$ must bind, otherwise we can decrease W without violating any other constraint, and obtain higher payoff.

Given the binding $IR_A(N)$, we can determine W. Take it into the objective function, then the H type's payoff is

$$(1-\beta_0)\sum_{n=1}^{N}\delta^n\eta_n(h-b_n)\Delta + V(N)$$

From Lemma 6, all IC constraints should bind. Furthermore, Lemma 7 gives an explicit characterization of the bonus scheme. In addition, Lemma 8 gives the limit form of the first part of the objective function. Let $\Pi^H_{pl}(T)$ denote the limit of the H type's payoff when her stopping time $N\Delta$ goes to T and Δ goes to 0. Then,

$$\Pi^H_{pl}(T) = (1-\beta_0)H_0^{\beta_0}(T) + \beta_0 V_0^H(T) + (1-\beta_0)V_0^L(T).$$

Therefore, taking the first derivative with respect to T, $\frac{\mathrm{d}\Pi^H_{pl}(T)}{\mathrm{d}T}$ is equal to

$$e^{-\rho T}[-(1-\beta_0)q_0 c\frac{(\lambda^H-\lambda^L)\lambda_0 e^{-(\lambda^L+\lambda^H)T}}{q_T\lambda_T(\beta_0\lambda^H e^{-\lambda^H T}+(1-\beta_0)\lambda^L e^{-\lambda^L T})} + q_0(\lambda^H h - c)e^{-\lambda^H T} - (1-q_0)c].$$

Note that $\frac{\mathrm{d}\Pi^H_{pl}(T)}{\mathrm{d}T} > 0$ is equivalent to

$$q_0(\lambda^H h - c) > (1-\beta_0)q_0 c\frac{(\lambda^H-\lambda^L)\lambda_0}{q_T\lambda_T(\beta_0\lambda^H e^{(\lambda^L-\lambda^H)T}+(1-\beta_0)\lambda^L)} + (1-q_0)ce^{\lambda^H T}.$$

Clearly, the RHS of above inequality is strictly increasing in T, and goes to infinity when T goes to infinity. In addition, when $T = 0$, the RHS equals

$$
\begin{aligned}
(1-\beta_0)q_0(\lambda^H-\lambda^L)\frac{c}{q_0\lambda_0}+(1-q_0)c &< (1-\beta_0)q_0(\lambda^H-\lambda^L)h+(1-q_0)c \\
&= c - q_0\lambda_0 h + q_0(\lambda^H h - c) \\
&< q_0(\lambda^H h - c),
\end{aligned}
$$

where both inequities come from the fact that $q_0\lambda_0 h > c$. Hence, there exits a $T^{pl} > 0$, such that $\frac{\mathrm{d}\Pi^H_{pl}(T)}{\mathrm{d}T}|_{T=T^{pl}} = 0$. When $T < T^{pl}$, $\Pi^H_{pl}(T)$ is strictly increasing; when $T > T^{pl}$, $\Pi^H_{pl}(T)$ is strictly decreasing. T^{pl} maximizes $\Pi^H_{pl}(T)$.

To show $T^{pl} < T_*$, we show that $\frac{\mathrm{d}\Pi^H_{pl}(T)}{\mathrm{d}T}|_{T=T_*} < 0$. Since $q_{T_*}\lambda_{T_*}h = c$, we have

$$
\beta_0\frac{l_0}{l^H}e^{-\lambda^H T_*}+(1-\beta_0)\frac{l_0}{l^L}e^{-\lambda^L T_*}=1.
$$

Thus, $\frac{\mathrm{d}\Pi^H_{pl}(T)}{\mathrm{d}T}|_{T=T_*}$ is equal to

$$
e^{-\rho T}(1-\beta_0)q_0(e^{-\lambda^H T_*}-e^{-\lambda^L T_*})\frac{[\beta_0(\lambda^H-\lambda^L)(\lambda^H h-c)e^{-\lambda^H T_*}+\lambda^L\frac{c}{l_0}]}{(\beta_0\lambda^H e^{-\lambda^H T_*}+(1-\beta_0)\lambda^L e^{-\lambda^L T_*})}<0.
$$

Furthermore, $T^{pl} < T_*$ implies $b_{T^{pl}} < h$, since $q_{T^{pl}}\lambda_{T^{pl}}b_{T^{pl}} = q_{T_*}\lambda_{T_*}h = c$.

Let us explore W^{pl}. Note that the agent's expected payoff is 0, i.e.

$$
W^{pl}+\int_0^{T^{pl}} e^{-\rho t}f_t(q_0,\beta_0)(q_t\lambda_t b_t - c)\,\mathrm{d}t = 0,
$$

and $q_t\lambda_t b_t - c > 0$ for all $t < T^{pl}$. Obviously, $W^{pl} < 0$. This also means the L type must obtain strictly positive expected payoff in the equilibrium, because her payoff before experimentation is strictly positive, i.e. $-W^{pl} > 0$, and her payoff during the experimentation is also strictly positive since $b_t < h$ for all t. Therefore, the L type's IR constraint is satisfied, and the contract we solved is indeed the best pooling equilibrium for the H type.

At last, we will show that $\frac{\mathrm{d}\Pi^H_{pl}(T^{pl})}{\mathrm{d}\beta_0} > 0$. Note that

$$
\frac{\partial \Pi^H_{pl}(T)}{\partial\beta_0} = (1-\beta_0)\frac{\partial H_0^{\beta_0}(T)}{\partial\beta_0} + V_0^H(T) - V_0^L(T) - H_0^{\beta_0}(T).
$$

From Lemma 8, for $0 < T \leq T^H_*$, $(1-\beta_0)\frac{\partial H^{\beta_0}_0(T)}{\partial \beta_0} > 0$, and

$$\begin{aligned} V^H_0(T) - V^L_0(T) - H^{\beta_0}_0(T) &= \int_0^T e^{-\rho t}(\eta_t h + \phi_t c)\,\mathrm{d}t - \int_0^T e^{-\rho t}\eta_t(h - b_t)\,\mathrm{d}t \\ &= \int_0^T e^{-\rho t}(\eta_t b_t + \phi_t c)\,\mathrm{d}t \\ &= e^{-\rho T}\frac{c}{q_T\lambda_T}\phi_T + \rho\int_0^T e^{-\rho t}\frac{c}{q_t\lambda_t}\phi_t\,\mathrm{d}t + \int_0^T e^{-\rho t}\phi_t c\,\mathrm{d}t > 0. \end{aligned}$$

Hence, $\frac{\partial \Pi^H_{pl}(T)}{\partial \beta_0} > 0$ for $0 < T \leq T^H_*$. Therefore, according to the Envelope Theorem, $\frac{\mathrm{d}\Pi^H_{pl}(T^{pl})}{\mathrm{d}\beta_0} = \frac{\partial \Pi^H_{pl}(T^{pl})}{\partial \beta_0} > 0$ since $0 < T^{pl} < T_* < T^H_*$. Clearly, when $\beta_0 \to 1$, $T^{pl} \to T^H_*$, and $\Pi^H_{pl}(T^{pl}) \to V^H_0(T^H_*)$. □

H Proof of Proposition 5

Proof. Consider the case where $q_0\lambda_0 h > c$, since otherwise the H type must obtain 0 payoff in any pooling equilibria. Let Π^θ_{pl} be the equilibrium payoff of the $\theta \in \Theta$ type principal in the BPEH. Clearly, $\Pi^H_{pl} < V^H_0(T^{pl})$. Let $\check{\epsilon} = V^H_0(T^{pl}) - \Pi^H_{pl} > 0$.

According to Lemma 8, since $0 < T^{pl} < T_* < T^H_*$, we have

$$\Pi^H_{pl} - \Pi^L_{pl} = \int_0^{T^{pl}} e^{-\rho t}\eta_t(h - b^{pl}_t)\,\mathrm{d}t = H^{\beta_0}_0(T^{pl}) < H^1_0(T^{pl}) = \int_0^{T^{pl}} e^{-\rho t}\eta_t(h - \mathring{b}_t)\,\mathrm{d}t,$$

where $\mathring{\boldsymbol{b}}$ are defined by the binding IC constraints for the agent when he believes the principal is the H type, i.e. $\mathring{b}_t = e^{-\rho(T^{pl}-t)}\frac{c}{q^H_{T^{pl}}\lambda^H} + \rho\int_t^{T^{pl}} e^{-\rho(s-t)}\frac{c}{q^H_s\lambda^H}\,\mathrm{d}s$.

Let $\acute{\epsilon} = H^1_0(T^{pl}) - H^{\beta_0}_0(T^{pl}) > 0$, $\mathring{\epsilon} = 1/2\min\{\acute{\epsilon}, \check{\epsilon}\} > 0$, and $\mathring{W} = \int_0^{T^{pl}} e^{-\rho t} f^H_t(q_0) q^H_t \lambda^H (h - \mathring{b}_t)\,\mathrm{d}t - \Pi^H_{pl} - \mathring{\epsilon}$.

Consider the contract $\mathring{C} = \{T^{pl}, \mathring{W}, \mathring{\boldsymbol{b}}\}$. If the agent believes that it is the H type who proposes $\mathring{C}$, then clearly his IC constraints are satisfied. Moreover, if he accepts the offer, he will obtain a positive payoff, i.e.

$$\mathring{W} + \int_0^{T^{pl}} e^{-\rho t} f^H_t(q_0)(q^H_t\lambda^H\mathring{b}_t - c)\,\mathrm{d}t = V^H_0(T^{pl}) - \Pi^H_{pl} - \mathring{\epsilon} = \check{\epsilon} - \mathring{\epsilon} > 0.$$

Hence, his IR constraint is also satisfied.

When the H type proposes $\mathring{C}$, and the agent believes that she is the H type, her payoff would be higher than Π^H_{pl}, i.e.

$$\mathring{\Pi}^H := -\mathring{W} + \int_0^{T^{pl}} e^{-\rho t} f^H_t(q_0) q^H_t \lambda^H (h - \mathring{b}_t)\, \mathrm{d}t = \Pi^H_{pl} + \mathring{\epsilon} > \Pi^H_{pl}.$$

When the L type proposes $\mathring{C}$ and the agent accepts the offer, the best action plan for the L type is to experiment on the project until the termination date T^{pl}, since the bonus payment is an increasing sharing scheme. Nevertheless, the L type's payoff would be less than Π^L_{pl}, i.e.

$$\mathring{\Pi}^L : = -\mathring{W} + \int_0^{T^{pl}} e^{-\rho t} f^L_t(q_0) q^L_t \lambda^L (h - \mathring{b}_t)\, \mathrm{d}t = \Pi^L_{pl} + \mathring{\epsilon} - \acute{\epsilon} < \Pi^L_{pl}.$$

If the agent rejects the offer, the L type would obtain 0 payoff, which is still strictly less than Π^L_{pl}.

Thus, the only possible type who proposes $\mathring{C}$ must be the H type, and the H type could obtain a higher payoff from $\mathring{C}$ if the agent believes she is the H type. Therefore, BPEH fails the intuitive criterion. □

I Proof of Lemma 4

Proof. Fix any equilibrium E that gives the L type strictly positive payoff. Let CL be the set of equilibrium contracts chosen by the L type, and CH be the set of equilibrium contracts chosen by the H type. Then $CL \subseteq CH$, and the L type can obtain the same payoff by proposing any $C \in CL$, the H type can obtain the same payoff by proposing any $C \in CH$.

For any contract $C \in CL$, let $x^C \in (0, 1]$ be the probability that the L type chooses C in the equilibrium, and $y^C \in (0, 1]$ be the probability that the H type chooses C in the equilibrium.[28] Then $\sum_{C \in CL} y^C \leq \sum_{C \in CL} x^C = 1$.

Clearly, there are only two possible cases.

- $CL = CH$, and $x^C = y^C$ for all $C \in CL = CH$;

[28]For technical simplicity, we assume that every contract over which one principal randomizes, she chooses that contract with strictly positive probability.

- There is at least one contract C, such that $x^C > y^C$.

In the first case, for any $C \in CL = CH$, when the principal proposes the contract C, the agent's belief about the principal's type is just the prior β_0. In addition, given this belief and the contract C, we know that the H's payoff in such an equilibrium cannot exceed the payoff she can obtain in the BPEH when the prior is β_0, by definition.

In the second case, for the contract C such that $x^C > y^C$, the agent's belief about the principal' type β is strictly below the prior β_0 by Bayesian rule. Given such belief and the contract C, the H type's payoff cannot exceed the payoff she can obtain in the BPEH when the prior is β, by definition. From Proposition 4, the H type's payoff in the BPEH is increasing in the prior. Thus, we can conclude that the H type's payoff also cannot exceed the payoff she can obtain in the BPEH when the prior is β_0. □

J Proof of Lemma 5

Proof. For the equilibrium such that the L type gets zero payoff, the best contract for the H type, $\hat{C} = \{\hat{N}, \hat{W}, \hat{\boldsymbol{b}}\}$, when the agent's belief about the H type is β_0 solves the following Program VI:

$$\max_{N,W,\boldsymbol{b}} \Pi^H = -W + \sum_{n=1}^{N} \delta^n f_{n-1}^H(q_0) q_n^H \lambda^H (h - b_n)\Delta$$
$$\text{s.t. } IC_A(N) \text{ , } IR_A(N) \text{ , } IC_L^H(N)$$

As in the BPEH, when β_0 is small, i.e. $\frac{c}{q_0\lambda_0} \geq h$, the optimal stopping time is 0, and the payoff of the H type is 0. Now, consider the non-trivial case when $\frac{c}{q_0\lambda_0} < h$.

We can still drop the agent's IR constraint, and solve the relaxed program, then verify it later. However, for the purpose to prove this lemma, we don't have to. The reason is that the solution to the relaxed program cannot be smaller than that to the original program, and observe that when $\beta_0 = 1$, the relaxed program is the same program solves the BSEH. What we need to show is just the value function for this relaxed program is increasing in β_0.

Consider the relaxed program without $IR_A(N)$. Clearly, $IC_L^H(N)$ must bind. Solve for W

and take it into the objective function, which is

$$\sum_{n=1}^{N} \delta^n \eta_n (h - b_n)\Delta.$$

By Lemma 6 and 7, the agent's IC constraints in $IC_A(N)$ must bind. Let $\hat{\Pi}^H(T)$ denote the limit of the H type's payoff when the stopping time $N\Delta$ goes to T and Δ goes to 0. From Lemma 8,

$$\hat{\Pi}^H(T) = e^{-\rho T}\phi_T(h - \frac{c}{q_T\lambda_T}) + \rho\int_0^T e^{-\rho t}\phi_t(h - \frac{c}{q_t\lambda_t})\,\mathrm{d}t.$$

For any $T > T^*$, since $q_T\lambda_T h < q_{T_*}\lambda_{T_*}h = c$, we have

$$\hat{\Pi}^H(T^*) - \hat{\Pi}^H(T) = -e^{-\rho T}\phi_T(h - \frac{c}{q_T\lambda_T}) - \rho\int_{T^*}^T e^{-\rho t}\phi_t(h - \frac{c}{q_t\lambda_t})\,\mathrm{d}t > 0.$$

Hence, the principal never wants to extend the project beyond the efficient stopping time of the mixed project; we can focus solutions in the compact set $[0, T_*]$. Since the objective function is continuous, there exists a solution $\hat{T} \in [0, T_*]$ to the relaxed program. We can further show that $\hat{T} \in (0, T_*)$ by showing $\frac{\mathrm{d}\hat{\Pi}^H(T)}{\mathrm{d}T}|_{T=0} > 0$ and $\frac{\mathrm{d}\hat{\Pi}^H(T)}{\mathrm{d}T}|_{T=T_*} < 0$.

Lemma 8 also shows $\frac{\partial\hat{\Pi}^H(T)}{\partial\beta_0} > 0$ for $0 < T \leq T_*^H$. Thus, Envelope Theorem implies $\frac{\mathrm{d}\hat{\Pi}^H(\hat{T})}{\mathrm{d}\beta_0} = \frac{\partial\hat{\Pi}^H(\hat{T})}{\partial\beta_0} > 0$, since $0 < \hat{T} < T_* \leq T_*^H$.

Hence, the value function of the relaxed program is 0 for $\frac{c}{q_0\lambda_0} \leq h$, and then strictly increasing in β_0 for $\frac{c}{q_0\lambda_0} > h$. Note that when $\beta_0 = 1$, it is the equilibrium payoff for the H type in the BSEH. Therefore, all the equilibrium that gives the L type 0 payoff, cannot give the H type higher payoff than the BSEH. □

K Proof of Proposition 7

In this section, we find the optimal mechanism for the H type principal when Δ is small. First, we restrict to feasible mechanisms when $N^L = 0$. Then, we focus on feasible mechanisms when $N^L \geq 1$. Last, we compare them when Δ is small.

Proof. (a) Consider the optimal mechanism for the H type when $N^L = 0$, $\mathcal{M}^0 = \{\mathcal{C}^{H0}, \mathcal{C}^{L0}\}$.

It will solve the following Program IV:

$$\max_{\mathcal{M}} \Pi^H(\mathcal{M}) = -W^H + q_0 \sum_{n=1}^{N^H} \delta^n (1-\lambda^H\Delta)^{n-1}\lambda^H(h-b_n^H)\Delta$$
$$\text{s.t. } IR_H(\mathcal{M}) \text{ , } IR_L(\mathcal{M}) \text{ , } IR_A(\mathcal{M}) \text{ , } IC_H^L(\mathcal{M}) \text{ , } IC_L^H(\mathcal{M}) \text{ , } IC_A(\mathcal{M}) \text{ , and } N^L = 0$$

Clearly, $IR_H(\mathcal{M})$ is automatically satisfied since we maximize the H type's payoff. We drop the constraints $IR_L(\mathcal{M})$ and $IC_H^L(\mathcal{M})$, and verify them later. In the relaxed program, $IC_L^H(\mathcal{M})$ and $IR_A(\mathcal{M})$ must bind, otherwise we can adjust W^H and W^L slightly to increase the payoff of the H type.

From the binding $IC_L^H(\mathcal{M})$ and $IR_A(\mathcal{M})$, we can obtain W^H and W^L, which in turn pins down

$$\Pi^H = \beta_0 V^H(N^H) + (1-\beta_0)\sum_{n=1}^{N^H}\delta^n\eta_n(h-b_n^H)\Delta,$$
$$\Pi^L = \beta_0 V^H(N^H) - \beta_0\sum_{n=1}^{N^H}\delta^n\eta_n(h-b_n^H)\Delta.$$

For any given N^H, according to Lemma 6 and 7, the agent's IC constraints in $IC_A(\mathcal{M})$ must also bind, and we can obtain b_n^H for $1 \le n \le N^H$.

Let $\Pi_0^\theta(T)$ denote the limit of the $\theta \in \Theta$ type's payoff when Δ goes to 0, and the stopping time of the H type $N^H\Delta$ goes to T, where the subscript 0 is reminiscent of the L type's experimenting time. Thus,

$$\Pi_0^H(T) = V_0^H(T) - (1-\beta_0)V_0^L(T) - (1-\beta_0)c\frac{\lambda^H-\lambda^L}{\lambda^H}[(1-q_0)\frac{1-e^{-(\rho+\lambda^L-\lambda^H)T}}{\rho+\lambda^L-\lambda^H} + q_0\frac{1-e^{-(\rho+\lambda^L)T}}{\rho+\lambda^L}],$$

$$\Pi_0^L(T) = \beta_0 V_0^L(T) + \beta_0 c\frac{\lambda^H-\lambda^L}{\lambda^H}[(1-q_0)\frac{1-e^{-(\rho+\lambda^L-\lambda^H)T}}{\rho+\lambda^L-\lambda^H} + q_0\frac{1-e^{-(\rho+\lambda^L)T}}{\rho+\lambda^L}].$$

The last thing is to find the optimal stopping time for the H type. Take the derivative with respect to T, we have

$$\frac{\mathrm{d}\Pi_0^H(T)}{\mathrm{d}T} = (1-\beta_0)\frac{q_0}{\lambda^H}(\lambda^H h - c)e^{-(\rho+\lambda^L)T}\alpha(T),$$

where $\alpha(T) := \frac{1}{1-\beta_0}\lambda^H e^{(\lambda^L-\lambda^H)T} - \frac{\beta_0}{1-\beta_0}\lambda^H\frac{l^H}{l_0}e^{\lambda^L T} - (\lambda^H-\lambda^L)\frac{l^H}{l_0}e^{\lambda^H T} - \lambda^L$ is strictly decreasing in T. When $T=0$, it is $\alpha(0) = (\frac{1}{1-\beta_0}\lambda^H - \lambda^L)(1-\frac{l^H}{l_0}) > 0$. When $T \to \infty$, it goes to

$-\lambda^L$. Hence, there exists a unique solution T^{H0} to the relaxed program, which satisfies $\alpha(T^{H0}) = 0$. We can also show that $\alpha(T^{sep}) > 0 > \alpha(T_*^H)$, thus $T^{sep} < T^{H0} < T_*^H$.

Now, let us check $IR_L(\mathcal{M})$. Note that

$$\frac{\mathrm{d}\Pi_0^L(T)}{\mathrm{d}T} = \beta_0 q_0(\lambda^H h - c)e^{-\rho T}[\frac{\lambda^L}{\lambda^H} e^{-\lambda^L T}(1 - \frac{l^H}{l_0} e^{\lambda^H T}) + \frac{l^H}{l_0}(e^{-(\lambda^L - \lambda^H)T} - 1)].$$

When $T < T_*^H$, $e^{\lambda^H T} < e^{\lambda^H T_*^H} = \frac{l_0}{l^H}$. Hence, $1 - \frac{l^H}{l_0} e^{\lambda^H T} > 0$ for $T < T_*^H$. We also have $e^{-(\lambda^L - \lambda^H)T} - 1 > 0$ for $T > 0$. Thus, $\frac{\mathrm{d}\Pi_0^L(T)}{\mathrm{d}T} > 0$ for $0 \leq T \leq T_*^H$, i.e. $\Pi_0^L(T)$ is strictly increasing in T. Hence, $\Pi_0^L(T^{H0}) > \Pi_0^L(0) = 0$, and $IR_L(\mathcal{M})$ is slack.

We also need to check $IC_H^L(\mathcal{M})$. By choosing $\mathcal{C}^{L0}$, the H type will get $\Pi_0^L(T^{H0})$, thus deviating to $\mathcal{C}^{L0}$ will get

$$\Pi_0^L(T^{H0}) - \Pi_0^H(T^{H0}) = -\int_0^{T^{H0}} e^{-\rho t} \eta_t (h - b_t^H)\,\mathrm{d}t < 0,$$

since $\int_0^{T^{H0}} e^{-\rho t} \eta_t (h - b_t^H)\,\mathrm{d}t = H_0^1(T^{H0}) > 0$. Thus, $IC_H^L(\mathcal{M})$ is also slack.

Therefore, the solution to the relaxed program is the optimal mechanism for the H type when the L type does not experiment. Moreover, $\Pi_0^L(T^{H0}) = 0$ implies $\Pi_0^H(T^{H0}) < V_0^H(T^{H0}) < V_0^H(T_*^H)$. Hence, the payoff of the H type is strictly less than her FIB payoff in the limit $\Delta \to 0$.

(b) Consider the optimal mechanism for the H type when $N^L \geq 1$. As we have shown, for a small Δ, the payoff of the H type has an upper bound $V^H(N_*^H) + \frac{1-\beta_0}{\beta_0} V^L(1) > 0$ when $N^L \geq 1$. We now show the following mechanism $\mathcal{M}^1$ is feasible and achieves this upper bound.

Recommendations - The mediator recommends the agent working until N_*^H if the principal reports H, but working once if the principal reports L.

Bonus schemes - Let $\{b_n^H\}_{2 \leq n \leq N_*^H}$ be the bonus transfers such that satisfy the agent's binding IC constraints[29], i.e. for $2 \leq n \leq N_*^H$,

$$b_n^H = \delta^{N_*^H - n} \frac{c}{q_{N_*^H}^H \lambda^H} + (1-\delta) \sum_{s=n}^{N_*^H - 1} \delta^{s-n} \frac{c}{q_s^H \lambda^H}.$$

Fix $\{b_n^H\}_{2 \leq n \leq N_*^H}$, the inequality (7) determines an upper bound for b_1^H. We choose some b_1^H which is below the upper bound. Fix the chosen $\{b_n^H\}_{1 \leq n \leq N_*^H}$, the inequality (6) determines a lower bound for b_1^L. Now, choose some b_1^L which is above both the lower bound and h. Therefore, the above bonus scheme satisfies both the agent and the L type's IC constraints. Moreover, because $b_1^L \geq h$, if the H type reports she is L, she will obtain negative payoff:

$$-W^L + \delta q_0 \lambda^H (h - b_1^L)\Delta = \delta q_0 (\lambda^H - \lambda^L)(h - b_1^L)\Delta \leq 0.$$

Thus, the above bonus scheme also satisfies the IC constraint of the H type.

Lump-sum transfers - Give the chosen bonus schemes, let W^θ for $\theta \in \Theta$ satisfy equation (4) and (5) to make the L type and the agent's obtain 0 payoffs. At last, the H type's IR constraint is automatically satisfied.

Hence, $\mathcal{M}^1$ is feasible. Clearly, the H type could achieve her payoff upper bound $V^H(N_*^H) + \frac{1-\beta_0}{\beta_0} V^L(1)$ from $\mathcal{M}^1$, since all the surplus is retained by the H type. When $\Delta \to 0$, her payoff converges to $V_0^H(T_*^H)$, since $\lim_{\Delta \to 0} V^L(1) = 0$.

(c) We have shown that in the limit $\Delta \to 0$, the H type's payoff in $\mathcal{M}^0$ is strictly below $V_0^H(T_*^H)$, while her payoff in $\mathcal{M}^1$ converges to $V_0^H(T_*^H)$. Therefore, when Δ is small enough, $\mathcal{M}^1$ is the optimal mechanism for the H type. □

L Implementation of the FIB

In this section, we examine the implementation problem of the FIB for any fixed $\Delta > 0$. We also allows that the mediator randomizes over recommendations to the agent. Remember, the FIB is that both types of projects are conducted efficiently, and the principal obtains all the surplus of her project. Obviously, Proposition 1 implies that the FIB cannot be fully

[29] Whether the IC constraints for the agent bind or not does not affect the result here. However, binding IC constraints are robust in the sense that they also provide the proper incentives for the agent to truthfully report a success when a success cannot be observed by the principal.

implemented by any mechanism. However, we can show that it can be virtually implemented through random recommendations.

Proposition 8. *Fix any $\Delta > 0$, the FIB can be virtually implemented.*

Proof. Consider the following mechanism $\mathcal{M}^{\mu}$.

Recommendations - The agent is recommended to work until N_*^H if the principal reports H. He is recommended to work with probability μ (and not to work with probability $1-\mu$) in period 1, where μ is small but strictly positive, and not to work thereafter, if the principal reports L. We have $V^H(N_*^H) + \frac{1-\beta_0}{\beta_0}\mu V^L(1) > 0$ when μ is close to 0.

Lump-sum transfers - Now the L type's lump-sum transfer can also depend on the random event - whether the mediator recommends the agent to work in the first period. Let W_R^L and W_{NR}^L be the lump-sum transfers when the mediator does recommend and does not recommend working in the first period, respectively. Choose the lump-sum transfers as $W_{NR}^L = 0$, $W_R^L = \delta q_0 \lambda^L (h - b_1^L)\Delta$, and

$$W^H = -[V^H(N_*^H) + \frac{1-\beta_0}{\beta_0}\mu V^L(1)] + \sum_{n=1}^{N_*^H} \delta^n f_{n-1}^H(q_0) q_n^H \lambda^H (h - b_n^H)\Delta.$$

Those lump-sum transfers gives the agent and L type binding IR constraints. The H type's IR constraint is automatically satisfied.

Bonus schemes - Let $\{b_n^H\}_{2\leq n\leq N_*^H}$ be the bonus transfers such that satisfy the agent's binding IC constraints for $2 \leq n \leq N_*^H$. Fix $\{b_n^H\}_{2\leq n\leq N_*^H}$, the L type's IC constraint is

$$V^H(N_*^H) + \frac{1-\beta_0}{\beta_0}\mu V^L(1) - \sum_{n=1}^{N_*^H} \delta^n \eta_n (h - b_n^H)\Delta \leq 0.$$

The above inequality determines an upper bound for b_1^H. We choose some b_1^H which is below the upper bound.

Fix the chosen $\{b_n^H\}_{1\leq n\leq N_*^H}$, the agent's IC constraint for $n=1$ conditional on being recommended to work is

$$q_0(\beta_0\lambda^H b_1^H + (1-\beta_0)\mu\lambda^L b_1^L) - c \geq \beta_0 q_0 \lambda^H \sum_{s=2}^{N_*^H} \delta^{s-1}(1-\lambda^H\Delta)^{s-2}(\lambda^H b_s^H - c)\Delta.$$

Table 1: When the principal's type is H

	Termination date	P's expected payoff	A's expected payoff
FIB	N_*^H	$V^H(N_*^H)$	0
$\mathcal{M}^\mu$	N_*^H	$V^H(N_*^H) + \frac{1-\beta_0}{\beta_0}\mu V^L(1)$	$-\frac{1-\beta_0}{\beta_0}\mu V^L(1)$

Table 2: When the principal's type is L

	Termination date	P's expected payoff	A's expected payoff	
FIB	0	0	0	
$\mathcal{M}^\mu$	0	0	0	w/ prob. $1-\mu$
	1	0	$V^L(1)$	w/ prob. μ

The above inequality determines a lower bound for b_1^L. We choose some b_1^L which is above both the lower bound and h. Therefore, the above bonus scheme satisfies both the agent and the L type's IC constraints. If the H type reports she is L, then she will get

$$\mu[-W_R^L + \delta q_0 \lambda^H (h - b_1^L)\Delta] = \mu\delta q_0(\lambda^H - \lambda^L)(h - b_1^L)\Delta \leq 0.$$

Therefore, the above bonus scheme also satisfies the H type's IC constraint.

Hence, the above mechanism $\mathcal{M}^\mu$ is feasible. Table 1 and 2 compare $\mathcal{M}^\mu$ with the FIB, regarding to the termination date of the project, the expected payoff of the principal, and that of the agent. Obviously, the outcome of $\mathcal{M}^\mu$ and the FIB can be arbitrary close as μ goes to 0. Thus, $\mathcal{M}^\mu$ virtually implements the FIB. □

References

Beaudry, Paul (1994), "Why an informed principal may leave rents to an agent." *International Economic Review*, 821–832.

Bedard, Nicholas Charles (2016), "Contracts in informed-principal problems with moral hazard." *Economic Theory Bulletin*, 1–14.

Bergemann, Dirk and Ulrich Hege (1998), "Venture capital financing, moral hazard, and learning." *Journal of Banking & Finance*, 22, 703–735.

Bergemann, Dirk and Ulrich Hege (2005), "The financing of innovation: Learning and stopping." *RAND Journal of Economics*, 719–752.

Cho, In-Koo and David M. Kreps (1987), "Signaling games and stable equilibria." *The Quarterly Journal of Economics*, 179–221.

Ely, Jeffrey C. and Martin Szydlowski (2016), "Moving the Goalposts." Technical report, working paper, Northwestern University.

Fryer, Roland G. and Richard T. Holden (2012), "Multitasking, learning, and incentives: a cautionary tale." Technical Report w17752, National Bureau of Economic Research, Cambridge, MA.

Guo, Yingni (2016), "Dynamic delegation of experimentation." *American economic review*.

Halac, Marina, Navin Kartik, and Qingmin Liu (2016), "Optimal contracts for experimentation." *The Review of Economic Studies*.

Hörner, Johannes and Larry Samuelson (2013), "Incentives for experimenting agents." *The RAND Journal of Economics*, 44, 632–663.

Jehiel, Philippe (2015), "On Transparency in Organizations." *The Review of Economic Studies*, 82, 736–761.

Karle, Heiko, Heiner Schumacher, and Christian Staat (2016), "Signaling quality with increased incentives." *European Economic Review*, 85, 8–21.

Kaya, Ayça (2010), "When does it pay to get informed?" *International Economic Review*, 51, 533–551.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic experimentation with exponential bandits." *Econometrica*, 73, 39–68.

Klees, Julia E. and Ron Joines (1996), "Occupational health issues in the pharmaceutical research and development process." *Occupational medicine (Philadelphia, Pa.)*, 12, 5–27.

Maskin, Eric and Jean Tirole (1990), "The principal-agent relationship with an informed principal: The case of private values." *Econometrica*, 379–409.

Maskin, Eric and Jean Tirole (1992), "The principal-agent relationship with an informed principal, II: Common values." *Econometrica*, 1–42.

Myerson, Roger B. (1983), "Mechanism design by an informed principal." *Econometrica*, 1767–1797.

Silvers, Randy (2012), "The value of information in a principal–agent model with moral hazard: The ex post contracting case." *Games and Economic Behavior*, 74, 352–365.

Stevens, Greg A. and James Burley (1997), "3,000 raw ideas= 1 commercial success!" *Research-Technology Management*, 40, 16–27.

Wagner, Christoph, Tymofiy Mylovanov, and Thomas Tröger (2015), "Informed-principal problem with moral hazard, risk neutrality, and no limited liability." *Journal of Economic Theory*, 159, 280–289.