Rewarding success and failure: moral hazard and adverse selection in strategic experimentation[1]

Fahad Khalil, Jacques Lawarrée
Department of Economics, University of Washington
and
Alexander Rodivilov
School of Business, Stevens Institute of Technology

January 30, 2024

A principal hires an agent to learn about the cost of a project via a series of experiments and then to execute it (production). The agent is privately informed about the probability that the cost is low, with the high-type agent earning rent because he is more optimistic than the low type. In each experiment, the agent privately chooses his effort. The joint presence of adverse selection and moral hazard makes it possible for the low type to "double deviate" by misrepresenting his type and shirking during experimentation. This leads to both truth-telling constraints being binding. The principal uses the length and outcomes of experimentation as well as the timing of payments to screen the agent by rewarding the low type after failure in experimentation. We provide sufficient conditions for both success and failure to be rewarded. We derive the optimality of over experimentation as it makes it less likely that the agent produces without uncovering the true cost and reduces the asymmetric information after a series of failed experiments. We also consider whether it may be optimal to separate experimentation and production between two different agents. Having the same agent working on both tasks enables the principal to use the adverse selection rent to address moral hazard. Therefore, integrating experimentation and production is optimal when adverse selection is severe.

Keywords: Strategic Experimentation, Moral hazard, Adverse selection, Outsourcing, Integration and Separation.

JEL classification: D83, D86.

[1] *E-mail*: khalil@uw.edu (F. Khalil), lawarree@uw.edu (J. Lawarree), arodivil@stevens.edu (A. Rodivilov).

"I haven't failed. I've just found 10,000 ways that don't work."

Thomas Edison

1. Introduction

Many important tasks involve two stages: a preliminary stage of experimentation or learning before a production or implementation stage. Consider, for instance, a surgeon who is contemplating surgery for a patient. To decide on an appropriate surgical procedure, a surgeon relies on their assessment of the prospect for success along with their diagnosis based on medical history and a series of diagnostic tests. While the prospect for success depends largely on a surgeon's prior experience and ability, the diagnosis is a dynamic learning process. In this paper, we study such a two-stage problem with an interaction between the experimentation and production stages in a mixed model with both moral hazard and adverse selection.

We introduce a principal-agent model where a production stage is preceded by a multi-period learning stage of strategic experimentation.[2] Each period of experimentation is subject to moral hazard by an agent trying to learn the cost of the project. Success in experimentation can only occur if the agent works and takes the form of uncovering "good news", i.e., the working agent finds out whether production cost is low. If the agent works, failure to uncover good news increases the expected cost of production. Unobserved shirking makes the principal more pessimistic than the agent about the true cost leading to a moral hazard rent during experimentation.[3] With an added production stage, the possibility of shirking implies a second moral hazard rent at the production stage as the principal, unaware of the shirking, overestimates the cost of production. Furthermore, the relationship is subject to adverse selection. The agent has private information about the ex ante probability that the cost is low, which is another source of asymmetric beliefs. The high type is more optimistic than the low type that the cost is low.

The first contribution of this paper is to show that both success and failure can be rewarded in a dynamic experimentation model with production. We find that moral hazard forces the principal to reward success, but the simultaneous presence of adverse selection may

[2] The exponential bandit model has been widely used as a canonical model of learning: see Bolton and Harris (1999) or Bergemann and Välimäki (2008).

[3] See, e.g., Bergemann and Hege (1998), Bergemann and Välimäki (2008), and Horner and Samuelson (2013).

make rewarding failure in the experimentation stage optimal. This is because failure is not only an indicator of shirking but also a more likely event for the low type. We show how the principal uses the length and outcomes of experimentation as well as the timing of payments to screen the agent. This is in contrast to recent papers with mixed models, which have emphasized pooling in one-shot models without a production stage.[4]

When experimentation is accurate (the probability of success is high), failure in experimentation is a strong indicator of the agent's type, while also reducing the cost of inducing effort. Then, we find that rewarding the low type for failure is optimal. The production stage plays a key role since the asymmetry of beliefs between the principal and the agent creates a scope for information rent based on expected production cost. While it is standard that the high type benefits from misreporting, we find that the low might have incentives to misreport as well. This is because, in a mixed model, the low type can "double deviate" by misrepresenting his type and then shirking. In our model, shirking off the equilibrium path allows the low type to remain more optimistic than the principal about the expected cost at the production stage. When both truth-telling constraints are binding, the principal finds it optimal to screen by rewarding the low type after failure, as that is a less likely event for the high type. The high type's rent is only given after success.

An alternative to rewarding the low type after failure is to shorten the high type's experimentation stage significantly. This reduces the scope for low type to create asymmetric beliefs by shirking off the path and, as a result, his incentive to misreport. Indeed, we find that the principal asks the high type to under-experiment. However, if experimentation is very accurate, this option is costly. Then, significantly shortening the high type's experimentation stage is suboptimal, and the principal prefers rewarding the low type after failure.

It may seem surprising that rewarding failure is optimal in the presence of moral hazard. The principal pays rent to the low type at the end of the experimentation stage after he fails in every prior period. Then, shirking becomes a profitable option for the agent, and rewards after success must be increased in each period of experimentation to deter the agent from delaying

[4] See, e.g., Ollier and Thomas (2013), Castro-Pires and Moreira (2021), and Gottlieb and Moreira (2022).

success. The agent is rewarded both after success and failure but gets a higher reward after success.

We also show that the principal asks the low type to *over*-experiment to mitigate the high type's rent. To understand why, note first that the lying high type does not shirk off the path, unlike the low type. A contract that induces the low type to work also induces the lying high type to work: being less pessimistic, he is more likely to collect promised rewards after success. Then, the source of the high type's rent is the asymmetry of information about expected costs after failure. By asking the low type to over-experiment, the principal makes it less likely that the lying high type produces without success and reduces the asymmetric information after a series of failed experiments.

If experimentation is costly and not so accurate (the probability of success is low), payments to induce effort are sufficiently high to address the truth-telling incentives for both types. The optimal contract is then driven by binding moral hazard constraints. We obtain the standard results in a model of experimentation without adverse selection and a production stage, where an agent is rewarded only for success and under experimentation is optimal.

Another contribution of our model is to analyze whether it is optimal to outsource experimentation to a second agent (separation) or to retain the same agent for both stages (integration). The standard model of experimentation based only on moral hazard would imply outsourcing the experimentation stage, which is typical in the literature.[5] We find that the presence of adverse selection can make integration optimal. Thus, our results suggest that isolating the experimentation stage is not without loss of generality if adverse selection is severe enough. A key benefit to the principal of having one agent for both stages is being able to use the adverse selection rent to address the moral hazard problem. If adverse selection is severe, yielding a large rent, the principal can even satisfy the moral hazard constraints 'for free' by spreading the adverse selection rent across time.

There are practical implications of our insight regarding the optimality of outsourcing experimentation depending on the relative strength of the two incentive problems. Consider first the case of drug approval trials, where a pharmaceutical company (principal) typically

[5] More precisely, by separating experimentation, the principal saves the additional moral hazard rent at the production stage mentioned above.

outsources to a clinical research organization (a separate agent) the clinical trials to demonstrate the effectiveness of a new drug. Moral hazard is the more serious issue.[6] Adverse selection is less relevant as much information about the prospective efficacy of the drug is in the public domain. Separation is optimal. In contrast, for the case of a surgeon described earlier, integrating the two stages (a series of diagnostic tests and surgery) would be optimal as seems to be the observed practice. Adverse selection is likely to be a major issue depending on each surgeon's expertise and experience. Moral hazard is less of a concern given protocols and regulations for healthcare activities required by the health insurance company or HMO.[7] As a result, integration is optimal.

Our analysis suggests that, when the tasks are separated, incentive schemes are simple, and the agent is paid after success. When the tasks are integrated, incentives schemes are more complex, and the agent can be paid after failure. Therefore, our model provides support for why CEOs are sometimes paid hefty compensations despite failure to perform. Under pure moral hazard, CEOs should be given the lowest possible wage (typically zero) upon failure. However, some recent papers surveyed by Edmans et al. (2017) rationalize payment after failure, for example, to induce CEOs to reveal negative information or to explore risky new technologies. Our model explains that payment after failure can serve as a screening instrument.

Related Literature. Our paper is related to the literature on contracting for experimentation following Bergemann and Hege (1998). Most of that literature considers either moral hazard or adverse selection models in isolation.[8] Among the few exceptions that introduce both moral hazard and adverse selection are Gerardi and Maestri (2012), Guo (2016), and Halac et al. (2016). Unlike all those papers, we also consider a production stage and show how the rent in one stage echoes into the other stage.[9] The presence of the production stage leads to the possibility of both types benefiting from private information. While the standard result in the

[6] There are multiple examples of clinical research organizations shirking, for example, by creating fake patient profiles (see Lindblad et al. (2014), Anand et al. (2012), Pogue et al. (2013) and references therein).
[7] In addition, healthcare practitioners are required by law to record patient medical histories and retain detailed case histories. There is also little room for skipping tests or altering results since this behavior might be simply illegal and a surgeon might be subject to prosecution. Surgeons are of course also bound by the Hippocratic Oath.
[8] See, e.g., Horner and Samuelson (2013), Sadler (2021), Escobar and Zhang (2021), Rodivilov (2022), and Moroni (2022) for models of pure moral hazard, and Gomes et al. (2016) and Khalil et al. (2020) for models of adverse selection only. Bhaskar and Mailath (2019) study a dynamic moral hazard with discrete actions model where the principal can use only short-term contracts, and Bhaskar and Roketskiy (2023) allow for continuous effort choice.
[9] Except for Khalil et al. (2020) who introduce a production stage but with adverse selection only.

literature is to reward success in the experimentation stage to address moral hazard, we find that the presence of adverse selection may make rewarding failure in the experimentation stage optimal.

Gerardi and Maestri (2012) also have a result akin to rewarding failure in a model with a fixed length of experimentation. The agent is rewarded if he fails to obtain a signal that the state is good, but only if his report matches the true state observed ex-post. The reason is that information is "soft" in Gerardi and Maestri (2012) (the agent's report is not verifiable), whereas information is "hard" in our model. Manso (2011) introduces a two-period model where failure is rewarded to incentivize the agent to explore riskier projects. In our model, there is only one available project, and rewarding failure is used to screen the low type.

More broadly, the literature on mixed models has pointed out that the optimal contract can be pooling.[10] In our model, the principal uses the timing of payments along with the length of experimentation and outcomes to induce effort and screen the agent. With multiple screening instruments, mixed models do not necessarily imply pooling. See Foarta and Sugaya (2021) for an example.[11] Martimort et al. (2023) study a mixed model with limited liability in a static setting and find that separation can be optimal. They use a dichotomous setup where the effort only determines a separate additive stochastic benefit but does not affect the cost of production or the output. They find pooling may occur but only for inefficient agents. We consider a dynamic model that is not dichotomous since the agent's effort impacts the expected cost of production through learning.

This paper is also related to the literature on endogenous information gathering before production. The standard model is static, based on early papers by Crémer and Khalil (1992), Lewis and Sappington (1997), and Crémer, Khalil, and Rochet (1998), where an agent exerts effort that increases the precision of the signal of the state (relevant for a production decision).[12] By modeling effort as experimentation, we contribute to this literature by introducing the

[10] See, e.g., Gottlieb and Moreira (2022).

[11] Castro-Pires et al. (2024) study a mixed model when the agent is risk averse and provide sufficient conditions under which the moral hazard problem can be decoupled from the adverse selection problem. Our setting does not satisfy those sufficient conditions since our problem is multi-dimensional as the optimal contract sets the wage, the length of experimentation and the output.

[12] For recent papers, see citations in Krähmer and Strausz (2011), Rodivilov (2021), Downs (2021), and Häfner and Taylor (2022).

dynamics of learning, and especially the possibility of learning with asymmetric speeds. In our model, the principal endogenously determines the degree of asymmetric information in the production stage by choosing the length of experimentation. Unlike the rest of the literature, we show that the principal may find it optimal to reward failure and to over-experiment to screen the types.

We also contribute to the literature on the power of incentives for innovation. Manso (2011) shows that an optimal incentive scheme may exhibit a reward for early failure for a risk averse agent. Benabou and Tirole (2003) show that using high-powered incentives may be detrimental to intrinsic motivation. In a laboratory experiment, Ederer and Manso (2013) find that a combination of rewards for both failure and success can be effective in incentivizing innovation. Sadler (2021) illustrates that high-powered incentives may discourage creativity. We contribute to this literature by showing theoretically that the coexistence of low- and high-powered incentive schemes can be optimal to mitigate the effect of adverse selection when failures to innovate are informative for the subsequent production decision.

Finally, our paper is also related to the extensive literature on integration and separation of tasks between agents. See, for instance, Schmitz (2005), Khalil et al. (2006), Iossa and Martimort (2012), Hoppe and Schmitz (2013 and 2021), and Li et al. (2015). Our dynamic model of learning allows us to pinpoint the relative importance of moral hazard and adverse selection in determining the optimal organization structure.

2. The Model and the First Best

A principal hires an agent to implement a project. The (marginal) cost of the project, c, is initially unknown to both the principal and the agent, but it is common knowledge that the cost can be low, $c = \underline{c} > 0$, with probability $\beta_0 \in (0,1)$, or high, $c = \overline{c}$, with probability $1 - \beta_0$, where $\overline{c} - \underline{c} = \Delta c > 0$. Both the principal and the agent are risk neutral, and, for simplicity, we assume that their discount factor is one. Before the actual production occurs, the agent gathers information regarding the production cost, which we model as a standard *experimentation stage*.[13] In the *production stage*, the agent produces based on what is learned about cost during the experimentation stage.

[13] See, e.g., Halac et al. (2016).

2.1. The Experimentation Stage

The length of the experimentation $T \in N$ is chosen by the principal. In each period $t \in \{1, 2, \dots, T\}$, the principal must address a moral hazard problem. At each period t, the agent privately chooses whether to perform an experiment, i.e., "work," $e_t = 1$, or not to experiment, i.e., "shirk," $e_t = 0$. Experimentation at t costs γe_t to the agent, where $\gamma > 0$.

The principal must also address an adverse selection problem. The agent is privately informed about the probability that the cost is low,

$$\beta_0^\theta = Pr\left(c = \underline{c} \middle| \theta\right),$$

where, the agent's type is denoted by the superscript $\theta \in \{H, L\}$, with $0 < \beta_0^L < \beta_0^H < 1$. In other words, a high type is more optimistic that the cost is low before experimentation starts, i.e., the *high type has a lower expected cost than the low type*. The principal believes the agent is a high type ($\theta = H$) with probability $\nu \in (0,1)$ and a low type ($\theta = L$) with probability $(1 - \nu)$.

We assume that information gathering takes the form of looking for good news. We say that the experimentation was successful if it reveals that the cost is low (*good news*). If the cost is actually low and the agent works at period t, success occurs with probability $0 < \lambda < 1$. Success is publicly observable. Once success occurs in a period t, the experimentation stage stops, and production takes place based on $c = \underline{c}$. Success cannot occur in a period t if the cost is high, or if the agent shirks. Thus, it is optimal for the principal to induce $e_t = 1$ in every period of the experimentation stage.

If the agent fails to learn that the cost is low in a period t, we say that experimentation failed in that period. Then, experimentation resumes if $t < T$, but both the agent and the principal become more pessimistic about the likelihood of the cost being low. Production takes place after the experimentation stage ends, either if the agent succeeds or if he fails all T times.

2.2. Updating Beliefs

Given that $e_t = 1$ for all t, and that experimentation has failed in previous periods, we denote by β_t^θ the updated belief of a type θ agent that the cost is low at the beginning of period t (after $t - 1$ failures). We have $\beta_t^\theta = \frac{\beta_{t-1}^\theta(1-\lambda)}{\beta_{t-1}^\theta(1-\lambda)+(1-\beta_{t-1}^\theta)}$, which can be re-written in terms of β_0^θ as follows:

$$\beta_t^\theta = \frac{\beta_0^\theta(1-\lambda)^{t-1}}{\beta_0^\theta(1-\lambda)^{t-1}+1-\beta_0^\theta}.$$

The expected cost for a type θ agent at the beginning of period t is then

$$c_t^\theta = \beta_t^\theta \underline{c} + (1-\beta_t^\theta)\,\overline{c}.$$

After each failure, the agent becomes more pessimistic about the true cost being low (β_t^θ falls), and the expected cost rises. For the same number of failures during the experimentation stage, a low type always remains *more pessimistic* than a high type and has a higher expected cost ($c_t^L > c_t^H$). However, both c_t^H and c_t^L approach $\overline{c}$ in the limit.

For future use, we note that the difference in the expected cost,

$$\Delta c_t \equiv c_t^L - c_t^H = (\beta_t^H - \beta_t^L)(\overline{c} - \underline{c}) = (\beta_t^H - \beta_t^L)\Delta c > 0,$$

is either decreasing in time (if $\beta_0^H \leq 1 - \beta_0^L$) or is *non-monotonic* with one peak (if $\beta_0^H > 1 - \beta_0^L$).[14] Two examples of Δc_t are presented in Figure 1.

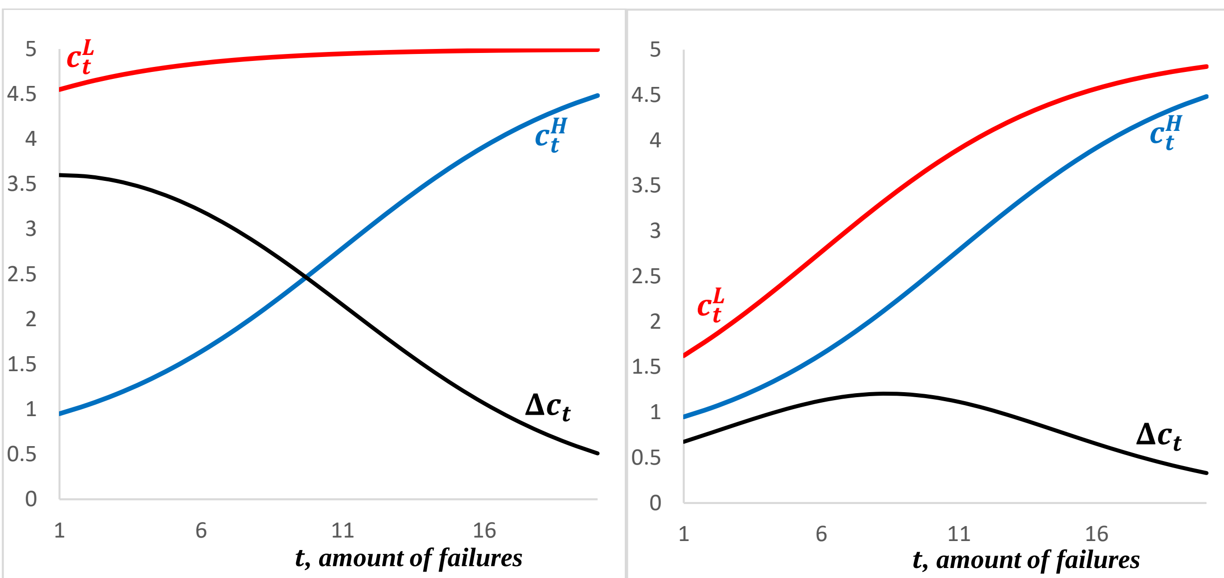

Figure 1. Expected cost with $\beta_0^H = 0.9$, $\lambda = 0.2$, $\underline{c} = 0.5$, $\overline{c} = 5$, $\beta_0^L = 0.1$ *(left)* and $\beta_0^L = 0.75$ *(right)*.

2.3. The Production Stage

Production takes place if experimentation succeeds in some t, or if it fails all T times. Since success publicly reveals low cost, the output after success is chosen under complete

[14] See Claim B in Appendix B for a formal proof.

information. The interesting case occurs when the agent has failed to learn during the entire experimentation stage since production then occurs under asymmetric information.[15]

Our assumption of a productive decision after failure is a significant departure from the standard literature on strategic experimentation. In the literature, the quantity after failure is implicitly assumed to be zero. Then, asymmetric beliefs after failure between the principal and agent do not matter for the incentives. As we will illustrate, the difference in beliefs matters whenever there is a decision at the end of the experimentation stage, and we capture it by assuming an explicit production stage even after failure. Thus, asymmetric information generated during the experimentation stage echoes into the production stage. Moreover, the anticipation of asymmetric information during production impacts the experimentation stage.

A simple way to capture the impact of asymmetric beliefs in production after failure is to assume that the output after failure is fixed at $q_F > 0$. We relax this assumption in an extension.[16] To be consistent with the extension section, we assume that the principal's value of the project is given by $V(q)$, which is strictly increasing and strictly concave. The output after success, q_S, is determined by $V'(q_S) = \underline{c}$. We assume $V(q_S) > V(q_F) > 0$. The cost of production after success is therefore $\underline{c}q_S$ and the expected cost after failure is $c_t^\theta q_F$, where $c_t^\theta = \beta_t^\theta \underline{c} + \left(1 - \beta_t^\theta\right) \overline{c}$.

2.4. The First Best Length of Experimentation

Suppose the agent's type θ is common knowledge and the agent's effort choice is publicly observable. The first-best length of experimentation T^θ for a type-θ agent determines the maximum expected surplus net of costs denoted by:

$$\Omega^\theta = \beta_0^\theta \sum_{t=1}^{T^\theta} (1-\lambda)^{t-1} \lambda \left[V(q_S) - \underline{c} q_S - \sum_{s=1}^{t} \gamma\right] \\ + \left(1 - \beta_0^\theta + \beta_0^\theta (1-\lambda)^{T^\theta}\right) \left[V(q_F) - c_{T^\theta+1}^\theta q_F - \sum_{s=1}^{T^\theta} \gamma\right].$$

Since the expected cost is rising until success is obtained, the first-best solution is characterized by a termination date T_{FB}^θ, the maximum number of periods an agent of type θ is allowed to experiment:

[15] We assume that the agent will learn the exact cost later, but it is not contractible.
[16] In the extension, the output is determined such that the marginal benefit of output equals its (expected) marginal cost. The key results are unaffected, except that variation in output after failure is an additional screening device.

$$T_{FB}^{\theta} \in arg \max_{T^{\theta}} \Omega^{\theta}.$$

Note that T_{FB}^{θ} is bounded and it is the highest t^{θ} such that

$$\beta_{t^{\theta}}^{\theta}\lambda\left[V(q_S) - \underline{c}q_s\right] + \left(1 - \beta_{t^{\theta}}^{\theta}\lambda\right)\left[V(q_F) - c_{t^{\theta}+1}^{\theta}q_F\right] \geq \gamma + \left[V(q_F) - c_{t^{\theta}}^{\theta}q_F\right].$$

The intuition is that, by extending the experimentation stage by one additional period, the type θ agent learns that $c = \underline{c}$ with probability $\beta_{t^{\theta}}^{\theta}$. If the experimentation stage is not extended, q_F is produced at the expected cost $c_{t^{\theta}}^{\theta}$, given in the RHS. The LHS describes the net benefit of extending by one period to t^{θ}. There is a chance to produce q_s at cost $\underline{c}$ if there is success, or to produce q_F at the updated expected cost at period $t^{\theta} + 1$, denoted by $c_{t^{\theta}+1}^{\theta}$ if the agent fails.

The first-best length of experimentation T_{FB}^{θ} is a *monotonic* function of the agent's type.[17] The reason is that the high type is more likely to learn $c = \underline{c}$ (conditional on the actual cost being low) since $\beta_t^H\lambda > \beta_t^L\lambda$ for any t. This implies that the principal should allow the high type to experiment longer. As is standard, we assume that it is always optimal to experiment at least once in the first-best case, where the principal observes effort and knows β_0^{θ}.[18] This restriction does not apply in the optimal contract under asymmetric information, where the principal is free to choose not to experiment.

We now consider the general model with both moral hazard and adverse selection.

2.5. The Principal's Problem: contract and payoffs

Before experimentation takes place, the principal offers the agent a menu of dynamic contracts. We restrict attention to deterministic contracts. A contract specifies, for each type of agent, the length of the experimentation stage and the transfer and output as a function of whether experimentation is successful in any period.

Without loss of generality, we use a direct truthful mechanism, where the agent is asked to announce his type, denoted by $\hat{\theta}$. A contract is defined formally by

---

[17] This is different from Halac et al. (2016) and Khalil et al. (2020), where the first-best termination date is non-monotonic in type and plays a key role. The reason for the non-monotonicity in those papers is that agent's type is given by λ, and the conditional probability of success is higher for the high type early but becomes lower as the length of experimentation increases.

[18] In particular, we assume that the principal would not choose q_S without experimenting.

$$\varpi^{\hat{\theta}} = \left(T^{\hat{\theta}}, \left\{w_t^{\hat{\theta}}(\underline{c}), w_t^{\hat{\theta}}(c_{t+1}^{\hat{\theta}})\right\}_{t=1}^{T^{\hat{\theta}}}\right),$$

where $T^{\hat{\theta}}$ is the (maximum) duration of the experimentation stage for the announced type $\hat{\theta}$, $w_t^{\hat{\theta}}(\underline{c})$ is the agent's wage if he observed $c = \underline{c}$ in period $t \leq T^{\hat{\theta}}$, and $w_t^{\hat{\theta}}(c_{t+1}^{\hat{\theta}})$ is the agent's wage if he fails $t \leq T^{\hat{\theta}}$ times.

An agent of type θ, announcing his type as $\hat{\theta}$, chooses the periods in which he works or shirks, where the number of periods the agent works is written as $\sum_{s=1}^{T^{\hat{\theta}}} e_s^{\theta}$. Denoting the effort profile by $\vec{e}^{\theta} = \{e_t^{\theta}\}_{t=1}^{t=T^{\hat{\theta}}}$, the agent receives expected utility $U^{\theta}(\varpi^{\hat{\theta}})$ at time zero from a contract $\varpi^{\hat{\theta}}$:

$$U^{\theta}(\varpi^{\hat{\theta}}, \vec{e}^{\theta}) = (1 - \beta_0^{\theta}) \left[\sum_{t=1}^{T^{\hat{\theta}}} [w_t^{\hat{\theta}}(c_{t+1}^{\hat{\theta}}) - \gamma e_t^{\theta}] - c^{\theta}_{\sum_{s=1}^{T^{\hat{\theta}}} e_s^{\theta} + 1} q_F \right]$$

$$+\beta_0^{\theta} \sum_{t=1}^{T^{\hat{\theta}}} \left(\prod_{s=1}^{t-1}(1 - \lambda e_s^{\theta})\right) \Bigg[e_t^{\theta} \lambda \left(w_t^{\hat{\theta}}(\underline{c}) - \underline{c} q_S\right) + (1 - \lambda e_t^{\theta}) w_t^{\hat{\theta}}(c_{t+1}^{\hat{\theta}}) - \gamma e_t^{\theta} +$$

$$(1 - \lambda e_t^{\theta}) 1_{\{t=T^{\hat{\theta}}\}} c^{\theta}_{\sum_{s=1}^{T^{\hat{\theta}}} e_s^{\theta} + 1} q_F \Bigg].$$

where the indicator function $1_{\{t=T^{\hat{\theta}}\}}$ is used to denote the last period of experimentation.[19]

Conditional on the actual cost being low, which happens with probability β_0^{θ}, the probability of succeeding for the first time in period $t \leq T^{\hat{\theta}}$ is given by $\left(\prod_{s=1}^{t-1}(1 - \lambda e_s^{\theta})\right) e_t^{\theta} \lambda$. If the agent succeeds in period t, he will produce q_S and is paid $w_t^{\hat{\theta}}(\underline{c})$. If he fails in period $t < T^{\hat{\theta}}$, experimentation continues but we allow for the agent to be paid $w_t^{\hat{\theta}}(c_{t+1}^{\hat{\theta}})$. We will show that the agent never receives a positive payment after failure except in the final period. The agent fails in a period t either if the cost is high, which happens with probability $1 - \beta_0^{\theta}$, or, if he fails despite the cost being low, which happens with probability $\beta_0^{\theta}(1 - \lambda e_s^{\theta})$. If the agent fails $T^{\hat{\theta}}$ times despite the cost being low, which happens with probability $\beta_0^{\theta} \prod_{s=1}^{T^{\hat{\theta}}}(1 - \lambda e_s^{\theta})$, the agent produces q_F based on the expected cost at period $T + 1$.

[19] The updating occurs in the following period, thus the "+1" in $\sum_{s=1}^{T^{\hat{\theta}}} e_s^{\theta} + 1$.

We denote by $\vec{a}^{\theta}(\varpi^{\hat{\theta}}) \equiv argmax_{\vec{e}^{\theta}} U^{\theta}(\varpi^{\hat{\theta}}, \vec{e}^{\theta})$ the optimal action profile for type θ in all periods $t \leq T^{\theta}$ facing a contract $\varpi^{\hat{\theta}}$. Denoting the equilibrium effort profile by $\vec{e}^{\theta} = \vec{1}$ (with $e_t^{\theta} = 1$ for all $t \leq T^{\theta}$), the optimal contract must satisfy the following (global) moral hazard constraint:

$$(\boldsymbol{MH^{\theta}}) \qquad \vec{1} \in \vec{a}^{\theta}(\varpi^{\theta}).$$

The optimal contract will have to satisfy the following incentive compatibility constraint for all θ and $\hat{\theta}$:

$$(\boldsymbol{IC^{\theta}}) \qquad U^{\theta}(\varpi^{\theta}, \vec{1}) \geq U^{\theta}\left(\varpi^{\hat{\theta}}, \vec{a}^{\theta}(\varpi^{\hat{\theta}})\right).$$

We refer to *reward* after success and failure by the agent's payment in each event net of production cost. Thus, we denote by y_t^{θ} the reward after success in period t, and by x_t^{θ} the reward after failure until period t:

$$y_t^{\theta} \equiv w_t^{\theta}(\underline{c}) - \underline{c}q_S,$$

$$x_t^{\theta} \equiv w_t^{\theta}(c_{t+1}^{\theta}) - 1_{\{t=T^{\theta}\}} c_{T^{\theta}+1}^{\theta} q_F.$$

We denote the probability that an agent of type θ does not succeed in any of the first t periods of the experimentation stage given $e_j^{\theta} = 1$ for all $j \leq t$ by:

$$P_t^{\theta} \equiv 1 - \beta_0^{\theta} + \beta_0^{\theta}(1-\lambda)^t.$$

Finally, we assume the agent's reward net of expected production cost must be non-negative.[20] To account for this, we impose the following limited liability ($\boldsymbol{LL}$) constraints whether experimentation succeeds or fails:

$$(\boldsymbol{LLS_t^{\theta}}) \qquad y_t^{\theta} \geq 0 \text{ for } t \leq T^{\theta},$$

$$(\boldsymbol{LLF_t^{\theta}}) \qquad x_t^{\theta} \geq 0 \text{ for } t \leq T^{\theta}.$$

---

[20] Bankruptcy laws and minimum wage laws are well-known examples of legal restrictions on transfers that exemplify limited liability in contracts. See, e.g., Krähmer and Strausz (2015) for more examples. Note that the agent is not protected off the equilibrium path. Thus, this is not a constraint representing the agent's wealth. Without limited liability, the principal can receive first best profit since success during experimentation is a random event correlated with the agent's type (Crémer-McLean, 1985). To streamline the presentation, we assume the transfers must cover expected cost. This is reminiscent of the well-known cost-plus contracts in the procurement literature.

The principal's expected payoff from a contract ϖ^θ offered to an agent of type θ, that satisfies the above constraints, is given by

$$\pi^\theta(\varpi^\theta, \vec{1}) = \beta_0^\theta \sum_{t=1}^{T^\theta}(1-\lambda)^{t-1}\left[\lambda\left(V(q_S) - w_t^\theta(\underline{c})\right) - (1-\lambda)w_t^\theta(c_{t+1}^\theta)\right]$$

$$+\left(1-\beta_0^\theta+\beta_0^\theta(1-\lambda)^{T^\theta}\right)V(q_F) - (1-\beta_0^\theta)\sum_{t=1}^{T^\theta} w_t^\theta(c_{t+1}^\theta).$$

Principal's Problem: The principal maximizes the objective function:

$$E_\theta[\pi^\theta(\varpi^\theta, \vec{1})] = \nu\pi^H(\varpi^H, \vec{1}) + (1-\nu)\pi^L(\varpi^L, \vec{1}).$$

s.t. (MH^θ), (IC^θ), (LLS_t^θ) and (LLF_t^θ) for $\theta \in \{H, L\}$.

The timing is as follows:

1. The agent learns his type θ.
2. The principal offers a contract to the agent. If the agent rejects the contract, the game is over and both parties get payoffs normalized to zero; if the agent accepts the contract, the game proceeds to the experimentation stage with maximum duration as specified in the contract.
3. The experimentation stage begins.
4. If the agent learns that $c = \underline{c}$, the experimentation stage stops, and the production stage occurs with output and transfers as specified in the contract. In case no success is observed during the entire experimentation stage, the production stage occurs with output and transfers as specified in the contract.

Besides the fact that the degree of asymmetric information is endogenous depending on the termination date T^θ, the presence of a production stage adds novel aspects to the mixed model. First, there is a second moral hazard rent, which is consequential to the integration-separation decision and may justify the focus on pure moral hazard models in the literature on strategic experimentation. Second, both the (IC) constraints can be binding leaving the principal no option but to reward failure in order to screen. One reason why both the (IC) constraints can be binding is standard in a mixed model: a low type may have an incentive to misreport and shirk. Another reason is that experimentation leads to a common value problem because the agent's type β_0^θ directly enters the principal's objective function.[21]

3. Solution

3.1. Simplifying the (MH^θ) and (IC^θ) constraints

Consider the (MH^θ) constraints. We first follow the standard step of replacing without loss of generality the global moral hazard constraint (MH^θ) by a sequence of local one-period moral hazard constraints (MH_t^θ).[22] The one-period moral hazard constraints (MH_t^θ) ensure that the agent will not engage in a one-shot deviation and shirk at period $t \leq T^\theta$ (given that the agent has worked in all prior periods $j < t$ without success and will work in all subsequent periods $s > t$).

$$(MH_t^\theta) \qquad y_t^\theta - x_t^\theta \geq \frac{\gamma}{\lambda\beta_t^\theta} + \sum_{s=t+1}^{T^\theta}(1-\lambda)^{s-t-1}\left(\lambda y_s^\theta + (1-\lambda)x_s^\theta - \gamma\right) + \frac{(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F.$$

The principal can motivate the agent to work by paying a higher reward for success (y_t^θ) than the reward after failure (x_t^θ). The first two terms on the *RHS* of $\left(MH_t^\theta\right)$ capture a standard rent in a dynamic model of experimentation without production (see, e.g., Bergemann and Hege (1998)).[23] A new feature relative to a standard dynamic moral hazard problem of experimentation is due to the presence of a production stage in our model. Consequently, in addition to the standard moral hazard rent during experimentation, the agent receives a *second* moral hazard rent at the production stage $\frac{(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F$ because a shirking agent will have a lower expected cost compared to what the principal believes.[24]

[21] See, e.g., Laffont and Martimort (2002), page 53. In a common value setting under pure adverse selection, that we also solve in Appendix D, both upward and downward incentive compatibility constraints can be binding because of a conflict between the principal's preference for the high type to experiment longer for pure efficiency reasons and the monotonicity condition imposed by asymmetric information.

[22] The formal proof is in Appendix C.

[23] On top of the static moral hazard rent due to limited liability, which is the first term on the RHS of $\left(MH_t^\theta\right)$, there is an additional rent in the standard dynamic moral hazard problem, reflected in the second term: (i) if the agent secretly shirks at period t, he is *more* likely to get to any future period $s > t$ than what the principal anticipates, reflected in the different probabilities, $(1-\lambda)^{s-t-1}$, as opposed to $(1-\lambda)^{s-t}$ that the principal would anticipate.

[24] The principal's belief is based on one more period of working compared to that of a shirking agent. Thus, the shirking agent has a lower expected cost: $c_T^\theta = c_{T+1}^\theta - \left(\beta_T^\theta - \beta_{T+1}^\theta\right)\Delta c < c_{T+1}^\theta$, and he will receive an additional production stage rent as a result.

Before we present the (IC^θ) constraints, we discuss off-equilibrium effort on the RHS of these constraints since the moral hazard problem is also implicitly reflected in the (IC^θ) constraints.[25] In Lemma 1 in Appendix A, we show that the misreporting high type will work in all periods if he claims to be a low type as is true in a standard mixed model.[26] A payment scheme under which a low type works will be enough to induce a misreporting high type to work since he is less pessimistic and more likely to collect promised rewards after success.

$$(\boldsymbol{IC^H}) \qquad \beta_0^H \sum_{t=1}^{T^H} (1-\lambda)^{t-1}\lambda y_t^H + \sum_{t=1}^{T^H} P_t^H x_t^H - \sum_{t=1}^{T^H} P_{t-1}^H \gamma \geq$$

$$(1-\beta_0^H)\sum_{t=1}^{T^L}[x_t^L - \gamma] + \beta_0^H \sum_{t=1}^{T^L}(1-\lambda)^{t-1}[\lambda y_t^L + (1-\lambda)x_t^L - \gamma] + P_{T^L}^H \Delta c_{T^L+1} q_F.$$

The expression $P_{T^L}^H \Delta c_{T^L+1} q_F$ represents a familiar adverse selection rent at the production stage: since a low type must be paid at least his expected cost $c_{T^L+1}^L$ following T^L failures, a high type will have lower expected costs if he lies and experimentation fails in all periods, which is captured by $\Delta c_{T^L+1} = c_{T^L+1}^L - c_{T^L+1}^H > 0$ in the expression.

Unlike the mis-reporting high type, the low type may not work in every period when he lies. Again, this is a common feature in many mixed models. We denote by $t^{L,H}$ the number of periods a low type works when he misreports. The RHS of the (IC^L) simplifies since the low type's probability of success in any period and the expected cost after failure depend on the total number of periods worked (and failed) up to that period (not on when those failures occurred). Furthermore, for expositional convenience and without loss of generality, we write his off-the-equilibrium path effort as a stopping rule: the low type works up to period $t^{L,H} \leq T^H$ and shirks thereafter.[27]

[25] A similar characterization is not easily available in Halac et al. (2016) as the agent's private information is about λ the efficiency of learning parameter. In that case, the relative probability of success across the two types changes over time. As a result, the authors provide examples that it is possible to have multiple off-equilibrium paths for effort in the optimal contract.

[26] See, e.g., Laffont and Martimort (2002), or Chakraborty et al. (2021).

[27] The proofs in the Appendices do not rely on off the path effort of the low type being a stopping rule. But a stopping rule is without loss of generality in our model because the high type's rewards for success can be front loaded given that the relative likelihood of success $\frac{\beta_0^L(1-\lambda)^{t-1}\lambda}{\beta_0^H(1-\lambda)^{t-1}\lambda} = \frac{\beta_0^L}{\beta_0^H}$ is independent of t. Without the stopping rule, the second expression on the RHS of (IC^L) is replaced in the Appendix A with the expression $\beta_0^L \sum_{t=1}^{T^H}\left(\prod_{s=1}^{t-1}(1-\lambda e_s^{L,H})\right)\left[e_t^{L,H}\lambda y_t^H + \left(1-\lambda e_t^{L,H}\right)x_t^H - e_t^{L,H}\gamma\right]$, where $e_t^{L,H}$ is the effort chosen by the mis-reporting low type in period t.

$(\boldsymbol{IC^L})$ $$\beta_0^L \sum_{t=1}^{T^L} (1-\lambda)^{t-1}\lambda\, y_t^L + \sum_{t=1}^{T^L} P_t^L x_t^L - \sum_{t=1}^{T^L} P_{t-1}^L \gamma \geq$$

$$(1-\beta_0^L)\left(\sum_{t=1}^{T^H} x_t^H - \gamma t^{L,H}\right) + \beta_0^L \sum_{t=1}^{t^{L,H}} (1-\lambda)^{t-1}\left(\lambda y_t^H + (1-\lambda) x_t^H - \gamma\right)$$

$$\beta_0^L (T^H - t^{L,H}) x_t^H - P_{t^{L,H}}^L \left(c_{t^{L,H}+1}^L - c_{T^H+1}^H\right) q_F.$$

The first three expressions on the RHS describe the payoffs after success and failure in each period, while the fourth term captures the impact of asymmetric beliefs if production occurs after failure. As in the case of (IC^H), the difference in expected costs is due to what the principal must pay a truthful high type $c_{T^H+1}^H$, versus the expected cost for a misreporting low type who only works for $t^{L,H}$ periods, and thus has an expected cost $c_{t^{L,H}+1}^L$. If $t^{L,H}$ is much smaller than T^H, i.e., the misreporting low type shirks often, he could be more optimistic than a high type who has worked for T^H periods, resulting in $c_{T^H+1}^H - c_{t^{L,H}+1}^L > 0$. Thus, the low type planning to shirk after misreporting can lead to (IC^L) being binding since he may command a rent at the production stage if he is more optimistic after shirking.

However, if $t^{L,H}$ is close to T^H, i.e., the misreporting low type works often, he could be less optimistic than a high type who has worked for T^H periods, resulting in $c_{T^H+1}^H - c_{t^{L,H}+1}^L < 0$. Then, misreporting is a *gamble* for the low type with his payoff depending on the outcome of experimentation: the positive part comes from obtaining the high-type's rent if he succeeds, and the negative part comes from an expected loss in the production stage if he fails in experimentation. The (IC^L) could still be binding since the low-type's gamble can be positive due to a common value problem (see footnote 21).

3.2. Optimal contract – relative impact of moral hazard and adverse selection

The solution to the principal's problem depends on the relative importance of moral hazard and adverse selection.

We will begin discussing the optimal solution with a familiar case in the strategic experimentation literature, where experimentation is modeled primarily as a moral hazard problem – how to motivate the experimenter to work. In our setting, this will happen when λ is small and γ is large. Then, the moral hazard rents are high enough to induce truth-telling, which implies that the (IC) constraints are slack. The optimal contract is characterized by the binding moral hazard constraints, and we present this case in Section 3.2.1 below.

In contrast, if λ is large and γ is small, moral hazard rents are low and (IC) constraints are binding. Then, adverse selection rent may come into conflict with moral hazard incentives, which we analyze in detail. The most interesting case is when the adverse selection constraints induce the principal to reward failure to screen the two types. We present this case in Section 3.2.2.

The results for the remaining two intermediate cases when only one (IC) constraint is binding follow from the analysis of the two main cases, and we discuss these cases in Section 3.2.3.

3.2.1. Case 1. Strong moral hazard: Both IC constraints are slack

First, we consider a case where experimentation accuracy λ is low such that moral hazard is very strong relative to adverse selection, and none of the (IC) constraints binds. In other words, neither type has incentive to misreport because of the high moral hazard rent they receive from their own contracts. We present the main findings in this case along with sufficient conditions for this case to occur in Claim 1 below. They require λ to be small, γ sufficiently higher than Δc, and β_0^L to be small.

Claim 1. For any β_0^H there exist $\overline{\lambda}(\beta_0^H) > 0$, $\underline{A}(\beta_0^H) > 0$, and $\overline{\beta}_0^L(\beta_0^H) > 0$ such that all the moral hazard constraints are binding for each type, but neither (IC) constraint is binding if

$$\lambda < \overline{\lambda}, \gamma > \underline{A}\Delta c q_F, \text{ and } \beta_0^L < \overline{\beta}_0^L:$$

(i) *Both types of agents receive two moral hazard rents: a standard rent in the experimentation stage and a second rent in the production stage.*

(ii) *Both types of agents are rewarded only after success, and the optimal reward y_t^θ is constant, given by $y_t^\theta = \frac{\gamma}{\lambda \beta_{T^\theta}^\theta} + \frac{(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F$ for $t \le T^\theta$ and $\theta \in \{H, L\}$.*

(iii) *Both types of agents under-experiment relative to the first best with $T_{SB}^L < T_{SB}^H$.*

Proof: See Appendix A.

Each type of agent is rewarded only after success, and the optimal reward y_t^θ is constant for $t \le T^\theta$.[28] There is no reward after failure, i.e., $x_t^\theta = 0$ for all t. As explained above when

[28] The reward deters a one-step-deviation by the agent, and the agent's incentive to deviate does not depend on t when there is no discounting (Rodivilov, 2022). Then, the optimal contract is unique up to payoff-irrelevant

describing the one-period moral hazard (MH_t^{θ}) constraints, the term $\frac{(1-\beta_0^{\theta})}{P_{T^{\theta}}^{\theta}}\Delta c q_F$ is what we called the second moral hazard rent. It stems from the shirking agent having a lower expected cost of production after failure than the principal. The term $\frac{\gamma}{\lambda\beta_{T^{\theta}}^{\theta}}$ represents the standard moral hazard rent in a dynamic model of experimentation without production.[29]

While *under* experimentation is optimal for both types: $T_{SB}^{L} < T_{FB}^{L}$ and $T_{SB}^{H} < T_{FB}^{H}$, the second moral hazard rent leads to a greater degree of under-experimentation than in moral hazard models of experimentation without a production stage. Each moral hazard rent increases with the length of experimentation because the divergent beliefs due to shirking increases with T^{θ}. Furthermore, we also show that the high type experiments longer $T_{SB}^{L} < T_{SB}^{H}$ as is true under the first best.

We now discuss the intuition behind the sufficient conditions for neither (IC) to be binding. From the binding (MH_t^{θ}), we can see that the moral hazard rents get larger if λ is smaller and γ is larger, which are the first two sufficient conditions for this case. If λ is small, the outcome of an experiment is not informative about the effort making it costlier to incentivize the agent to work. We also need the cost of experimentation γ to be sufficiently higher than Δc, which again makes moral hazard more important than adverse selection. Truth telling is obtained "for free" in this case as neither type wants to misrepresent his private information about β_0^{θ}.[30] The third sufficient condition that β_0^{L} is small ensures that high moral hazard payment to the low type deters him from the high type's contract.

Next, we consider the other polar case where moral hazard is less strong and experimentation is very effective such that both (IC) constraints are binding. The sufficient conditions on λ and γ are a mirror image of the ones above, and inducing effort on the equilibrium path is not too costly. The agent's incentives are largely driven by the impact of asymmetric beliefs at the production stage due to adverse selection concerns, while the impact of

alteration. A similar reward structure holds in Halac et al. (2016) who argue that in the case of no discounting, the principal can be restricted to using constant bonus contracts.

[29] The standard rent has two parts, where $\frac{\gamma}{\lambda\beta_t^{\theta}}$ addresses the static gain, and $\gamma\sum_{s=1}^{T^{\theta}-t}\frac{(1-\beta_0^{\theta})}{\beta_0^{\theta}(1-\lambda)^{t+s-1}}$ is the rent coming from a higher probability of collecting future moral hazard rents (than the principal expects in equilibrium).

[30] The high type is not attracted by the low type's contract as the high type experiments longer and receives a moral hazard rent over more periods $(T_{SB}^{L} < T_{SB}^{H})$ when he tells the truth.

shirking becomes important mainly off the equilibrium path as noted above when discussing the (IC) constraints.

3.2.2. Case 2. Strong adverse selection: Both (IC^H) and (IC^L) are binding

When λ is high, experimentation is very effective, moral hazard payments are small, and both (IC) constraints are binding. Surprisingly, we find that the optimal contract requires rewarding the low type for failure, which conflicts with moral hazard constraints. While this forces the principal to increase rewards after success in each period to induce effort, a high λ makes it optimal to do so. In Claim 2 we present the main results for this case along with sufficient conditions. They require λ high, γ sufficiently smaller than Δc, and β_0^L small.

Claim 2. For any β_0^H there exist $\underline{\lambda}(\beta_0^H) < 1$, $\overline{A}(\beta_0^H) > 0$ and $\gamma^{2,MH} > 0$ such that both IC bind if

$$\lambda > \underline{\lambda}, \gamma < \min\{\overline{A}\Delta c q_F, \gamma^{2,MH}\}, \text{ and } \beta_0^L < \tfrac{1}{2}:$$

(i) *To address moral hazard, the principal must reward each type after success in every period ($y_t^\theta > 0$ for $t \leq T^\theta$).*

(ii) *The high type is not rewarded after failure ($x_t^H = 0$ for $t \leq T^H$), while the low type is, but only in the last period ($\boldsymbol{x_{T^L}^L > 0} = x_t^L$ for $t < T^L$).*

(iii) *Relative to the first best, the low type over-experiments ($T_{SB}^L > T_{FB}^L$), while the high type under-experiments ($T_{SB}^H < T_{FB}^H$).*

Proof: See Appendix A.

If λ is high, the impact of experimentation on the endogenous asymmetry of information is significant. Recall that the incentive to misreport depends on T^θ through the difference in expected cost Δc_{T^θ}. To understand the distortion in T^L, we need to examine the key positive element in the high-type's rent, $P_{T^L}^H \Delta c_{T^L+1} q_F$, which is decreasing in T^L. This leads to *over*-experimentation for the low type to reduce asymmetric information regarding the agent's type, which would not occur under pure moral hazard. To understand the distortion in T^H, recall that the low type's incentive to misreport is determined by the expected cost the principal must pay a truthful high type $c_{T^H+1}^H$ after failure, which is increasing in T^H. This leads to under-experimentation for the high type, reinforcing the impact of moral hazard. This also reduces the scope for low type to create asymmetric beliefs by shirking off the path and, as a result, his incentive to misreport. However, if experimentation is very accurate, this option is costly. The principal cannot distort T^H too much and (IC^L) is binding.[31]

One important result of our model is the optimality of rewarding failure if both (IC) are binding. Each type is rewarded for an event that is more likely to occur given the type, which is success for the high-type and failure for the low-type, respectively. Because the relative probability of failure $\frac{P_t^L}{P_t^H}$ is increasing in t, it is optimal to postpone the low-type's reward to the very last period of the relationship, $x_{T^L}^L > 0$, making it less likely for a (misreporting) high type to obtain it. In our mixed model, failure is not only an indicator of shirking but also a more likely event for the low type. If only one (IC) is binding, the principal does not reward failure as we will see in the two intermediate cases below.

A high λ also limits the moral hazard payments after success. Paying an additional screening rent after failure in the last period requires the principal to raise the reward after success (by the same amount) in each period. That is, the low-type agent must be given extra incentive to work in each period since he is now paid extra rent after failure in the last period. Ultimately, the low type is still paid more after success than failure, but he gets rent even if he fails, which is not the case for the high type.

Technically, to satisfy the (MH_t^L) constraints, the reward after success must increase not only in that last period, $y_{T^L}^L = \boldsymbol{x_{T^L}^L} + \frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F$, but also in the all the previous periods $t < T^L$:

$$y_t^L = \frac{\gamma}{\lambda\beta_t^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L - \gamma) + (1-\lambda)^{T^L-t}\boldsymbol{x_{T^L}^L},$$

which increases the payments y_t^L strictly above the optimal levels described in Claim 1.

Paying the high type a rent after a failure x_t^H makes it costlier to satisfy the binding (IC^H) since the high type is less likely to fail. Therefore, it is optimal to choose

$$x_t^H = 0 \text{ for } t \le T^H.$$

[31] For instance, if the high type is asked to produce without experimentation, the low type reports his type honestly.

The timing of payments after success y_t^H is ineffective as a screening instrument as the relative probability of success between the types $\left(\frac{\beta_0^L(1-\lambda)^{t-1}\lambda}{\beta_0^H(1-\lambda)^{t-1}\lambda} = \frac{\beta_0^L}{\beta_0^H} < 1\right)$ is constant across periods of experimentation. Consequently, there is no restriction on when to pay the screening rent to the high type via a combination of y_t^H. [32]

$$y_t^H \geq \frac{\gamma}{\lambda\beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F \textit{ for } t \leq T^H \text{ (strict inequality for some } t\text{).}$$

The principal can use the screening rent to induce effort in each period and satisfy the high type's MH constraints with no additional cost.

We can summarize the results of our two claims in the following proposition:

Proposition 1: *If moral hazard is important because γ is large and λ is small, then both types of agents are paid only after success and each type under experiments (Claim 1). If adverse selection plays a central role, with γ small and λ large, then the low-type agent over-experiments and is paid a rent after both success and failure, while the high-type under-experiments and is paid after success only (Claim 2).*

Next, we consider intermediate cases when only one (IC) constraint is binding.

3.2.3. Intermediate Cases 3 and 4: only one (IC) is binding

We outline how the analysis of the two main cases above provide the intuition for the key results and sufficient conditions, while the precise details are presented along with proofs in Appendix A as Claims 3 and 4.

In the two intermediate cases, either (IC^H) or (IC^L) is binding, but not both. There are common elements in the two cases. Both types are only paid rent after success and there is no rent after failure. If only one (IC) is binding, each type is paid their screening rent only after success without having to worry about increasing the cost of satisfying the other (IC).[33] For the type earning a screening rent, his moral hazard constraints are slack. Conversely, if a type does not earn a screening rent, all his moral hazard constraints are binding. Thus, if (IC^H) is binding,

[32] For example, it is without loss of generality to pay the extra rent to the high type after the very first success, i.e., front load the extra rent.

[33] Again, there is no restriction on when this rent is paid since the relative probability of success is independent of type.

the distortion in the low type's contract is only in the termination date T^L, while T^H is first best. Similarly, the distortion is in T^H when (IC^L) is binding while T^L is first best.

Again, moral hazard considerations tend to favor under-experimentation, while adverse selection has opposite implications on T^L and T^H. Increasing T^L reduces asymmetry of information rent by the high type, while reducing T^H decreases the payoff of a misreporting low type. The optimal distortion depends on the strength of the moral hazard versus adverse selection concerns. For example, when γ is very high, which would be one of the sufficient conditions for these intermediate cases, there is no over-experimentation. Thus, over-experimentation is largely driven by effective experimentation having the potential to reduce the impact of asymmetric information as may be intuitive but not readily found in the literature.

For only one (IC) to be binding, it is important to have γ high enough compared to Δc so that moral hazard payments are high to deter misreporting by the type whose IC is not binding. In addition, for only (IC^H) to be binding, it is sufficient to have β_0^H high and β_0^L small, while it is sufficient to have β_0^H close to β_0^L and a not too high λ for the case when only (IC^L) is binding.

4. Is integrating experimentation and production optimal?

In our model with experimentation and production, the interaction of adverse selection and moral hazard creates interdependent rents. In a pure moral hazard model, the principal would prefer to employ two different agents, one for experimenting and one for producing. This justifies the standard approach in the strategic experimentation literature, which studies a pure moral hazard experimentation stage model in isolation without a production stage. By separating the two stages, the principal saves what we have called the second moral hazard rent at the production stage. This begs the question of whether integrating the two tasks, as in our main model, can be optimal due to the presence of adverse selection.

A key benefit of integrating the two tasks is to use the adverse selection rent to induce effort, i.e., pay for the moral hazard rent. The rent needed to satisfy the (IC) constraints can be spread across time to satisfy the dynamic moral hazard constraints. Since the relative probability of success across types is time-invariant, the exact distribution of this rent does not impact the incentive to misreport. This benefit must be balanced against the cost of the second moral hazard rent when integrating. We find that integration is optimal if the adverse selection problem is severe enough relative to the moral hazard problem in experimentation (β_0^H far apart from β_0^L). We present below sufficient conditions for separation/integration to be optimal.

Proposition 2: ***Separation vs. integration*:**

(i) *Separation is optimal if the adverse selection problem is small enough: for any β_0^L there exists a value of β_0^H, called $\overline{\beta}_0^H(\beta_0^L)$, such that separation is optimal if $\beta_0^H < \overline{\beta}_0^H(\beta_0^L)$.*

(ii) *Integration is optimal if the adverse selection problem is severe enough (β_0^H is close to one and β_0^L sufficiently close to zero) and ν is high enough.*

Proof: See Appendix F.

To establish the above result, we can use a very simple extension of our model, where the principal outsources the experimentation task to a second agent (experimenter). The first agent, the 'in-house' agent, produces output based on what is learned publicly in the experimentation stage, and on his private information about the likelihood of low cost, β_0^θ. We discuss alternative models of separation at the end of this section. Before experimentation starts, the in-house producer is asked to publicly announce his type, based on which experimentation occurs. The principal pays an adverse selection rent to the producer to induce truthful reporting. The experimentation stage is a pure moral hazard problem, yielding only a standard moral hazard rent to the experimenter (based on a commonly known β_0^θ).

When she separates the two tasks, the principal saves the moral hazard rent at the production stage but pays an adverse selection rent to the in-house producer and a moral hazard rent to the experimenter. When this moral hazard rent is small relatively to the adverse selection rent, the principal's ability to use the adverse selection rent to satisfy moral hazard constraints dominates, and integration is optimal: for example, when the difference in cost Δc is small relative to the difference between β_0^H and β_0^L.[34]

A possible issue regarding the model of separation above is that we assume an in-house producer publicly pre-announces the type β_0^θ. We chose this benchmark for ease of comparison with the main model of integration. Instead, we could assume that the production agent is brought in after experimentation ends and therefore cannot announce his type before experimentation starts. Our key arguments regarding the optimality of integration would only

[34] As we show in our proofs, this basic intuition holds regardless of which (IC) is binding.

get stronger. This would also be the case if the in-house producer privately announced his type β_0^θ to the principal. We briefly discuss these two sub-extensions next.

Consider first that, under separation, experimentation occurs under a common prior, between the principal and the experimenter, that the cost of production is low with probability β_0. There is now an additional cost of separation as the length of experimentation can no longer be based on the private information about β_0^θ of the (integrated) agent. Next consider the case where the in-house producer privately announces its type to the principal, who contracts with an outside experimenter. In the interim, a principal's incentive constraint would also have to be satisfied, which will again reduce the benefit of separation.

5. Endogenous Output

In this section, we allow the principal to choose output optimally after success and after failure, and she can now use output as another screening variable. While our main findings continue to hold, output after failure can now be used as a screening device. Thus, the key new results occur if the experimentation stage fails: the low type is asked to under-produce relative to the first best, while the high type might over-produce. Just like over-experimentation, over-production can be used to increase the cost of lying.

When output is optimally chosen by the principal in the contract, the main change from the base model is that output after failure, which is denoted by q_F^θ, can vary depending on the expected cost. We can replace q_F by q_F^θ in the principal's problem.

We derive the formal output scheme in Appendix E but present the intuition here. When experimentation is successful, there is no asymmetric information and no reason to distort the output. Both types produce the first best output. When experimentation fails to reveal the cost, asymmetric information will induce the principal to distort the output to limit the rent.

When both (IC^L) and (IC^H) are slack, each type under-produces after failure to reduce the moral hazard rent. When the (IC^H) constraint is slack, but (IC^L) binding, the output for the low type, q_F^θ, does not affect information rents and, as a result, is not distorted. The high type, however, might be asked to over-produce whenever the misreporting low type is more pessimistic than the principal after failing in experimentation. Over-production is, therefore, used to increase the cost of the lying low type.[35]

When the (IC^H) binds and (IC^L) is slack, the low type is asked to under-produce in order to limit the rent of the high type. The output for the high type, q_F^θ, does not affect information rents and, as a result, is not distorted. When both (IC) constraints are binding, the low type under produces to limit the rent of the high type. Like the case when only (IC^H) binds, the high type might be asked to over produce to increase the cost of the lying low type.

6. Conclusions

We presented a dynamic model of strategic experimentation with both moral hazard and adverse selection. Technically, such a mixed model of experimentation can become quickly intractable with off the equilibrium path effort hard to characterize.[36] We offer a tractable model to provide an explanation for the co-existence of both high and low-powered incentive schemes, which is used in practice to spur innovative activity. We find that, while moral hazard always leads the principal to reward success in experimentation, the simultaneous presence of adverse selection may induce the principal to reward failure. The reason is that rewarding failure allows the principal to dynamically screen the agents, and it remains optimal even in the presence of moral hazard. We also characterize how the principal can use over and under-experimentation to provide incentives.

We also find that the principal may prefer to integrate experimentation and production by employing one agent for both. We show that the standard model of experimentation, where experimentation is studied in isolation without a production stage, is valid as long as adverse selection during experimentation is not a significant concern. Integration of experimentation and production allows the principal to use the adverse selection rent to incentivize the agent to work. By distributing the adverse selection rent optimally, the principal can alleviate the moral hazard

[35] Since the misreporting low type may shirk off the equilibrium path, his expected cost at the production stage does not necessarily have to be greater than the expected cost for a high type on the equilibrium path. Thus overproduction is optimal when off the equilibrium path effort of the misreporting low type involves very little shirking.

[36] If the adverse selection lies in the probability of success, as in Halac et al. (2016), the relative probability of success between the two types changes in ranking over time.

constraints. This is the case if the adverse selection problem is severe enough relative to moral hazard or, equivalently, if the adverse selection rent is high.

Appendix A

Outline: Proof of Claims 1 and 2: Parts (i), (ii), and (iii) of each claim are proved in Sections I and II. The sufficient conditions for each case are derived in Section III.

In Section I, we characterize the optimal payment structure for each of the 4 cases. In Section II, we characterize the optimal length of experimentation for each of the 4 cases. In Section III, we provide sufficient conditions for each of the 4 cases to occur in equilibrium.

Section I. The optimal payment structure

The principal's optimization problem is to choose contracts ϖ^θ for $\theta \in \{H, L\}$ to maximize

$$E_\theta \left\{ \Omega^\theta - \beta_0^\theta \sum_{t=1}^{T^\theta} (1-\lambda)^{t-1} \lambda y_t^\theta - \sum_{t=1}^{T^\theta} P_t^\theta x_t^\theta \right\} \text{ s.t.}$$

(IC^θ) $\quad U^\theta(\varpi^\theta, \vec{1}) \geq U^\theta\left(\varpi^{\hat{\theta}}, \vec{a}^\theta(\varpi^{\hat{\theta}})\right),$

(MH_t^θ) $\quad y_t^\theta - x_t^\theta \geq \frac{\gamma}{\lambda \beta_t^\theta} + \sum_{s=t+1}^{T^\theta} (1-\lambda)^{s-t-1} \left(\lambda y_s^\theta + (1-\lambda) x_s^\theta - \gamma\right) + \frac{(1-\beta_0^\theta)}{P_{T^\theta}^\theta} \Delta c q_F$ for $t \leq T^\theta$,

(LLS_t^θ) $\quad y_t^\theta \geq 0$ for $t \leq T^\theta$,

(LLF_t^θ) $\quad x_t^\theta \geq 0$ for $t \leq T^\theta$.

First, recall that we have replaced the global moral hazard constraint by a sequence of local (MH_t^θ) without loss of generality (see Appendix C for formal proof). The (MH_t^θ) constraints imply that all the (LLS_t^H) and (LLS_t^L) constraints are automatically satisfied and, therefore, can be ignored. When $\vec{e}^\theta = \vec{1}$, the notation for the expected costs in the (IC^θ) constraint is $c_{T^\theta+1}^\theta$. When the agent lies about his type and possibly shirks, we need to introduce new notation. On the *RHS* of (IC^H) we label the effort chosen by the agent in each period s as $e_s^{H,L} \in \{0,1\}$, and the expected cost is $c^H_{\sum_{s=1}^{T^L} e_s^{H,L}+1}$. Similarly, on the *RHS* of $(IC^{L,H})$ we label the effort chosen by the agent in each period s as $e_s^{L,H} \in \{0,1\}$, and the expected cost is $c^L_{\sum_{s=1}^{T^H} e_s^{L,H}+1}$.

Labeling ξ^H, ξ^L, $\{\mu_t^H\}_{t=1}^{T^H}$, $\{\mu_t^L\}_{t=1}^{T^L}$, $\{\eta_t^H\}_{t=1}^{T^H}$, $\{\eta_t^L\}_{t=1}^{T^L}$ as the Lagrange multipliers of the constraints associated with (IC^H), (IC^L), (MH_t^H), (MH_t^L), (LLF_t^H) and (LLF_t^L) respectively, the optimization problem has the following Lagrangian:

$$\mathcal{L} = \nu\left[\Omega^H - \beta_0^H \sum_{t=1}^{T^H} (1-\lambda)^{t-1} \lambda y_t^H - \sum_{t=1}^{T^H} P_t^H x_t^H\right]$$

$$+(1-\nu)\left[\Omega^L - \beta_0^L \sum_{t=1}^{T^L} (1-\lambda)^{t-1} \lambda y_t^L - \sum_{t=1}^{T^H} P_t^H x_t^L\right]$$

$$+\xi^H \left[\begin{array}{c} \beta_0^H \sum_{t=1}^{T^H} (1-\lambda)^{t-1} \lambda y_t^H + \sum_{t=1}^{T^H} P_t^H x_t^H - \sum_{t=1}^{T^H} P_{t-1}^H \gamma \\ -(1-\beta_0^H) \sum_{t=1}^{T^L} \left[x_t^L - \gamma e_t^{H,L}\right] - \beta_0^H \sum_{t=1}^{T^L} \left(\prod_{s=1}^{t-1} (1-\lambda e_s^{H,L})\right) \left[e_t^{H,L} \lambda y_t^L + \left(1-\lambda e_t^{H,L}\right) x_t^L - e_t^{H,L} \gamma\right] \\ -\left(1-\beta_0^H + \beta_0^H \left(\prod_{s=1}^{T^L} (1-\lambda e_s^{H,L})\right)\right) \left(c_{T^L+1}^L - c^H_{\sum_{s=1}^{T^L} e_s^{H,L}+1}\right) q_F \end{array}\right]$$

$$+\xi^L\left[\begin{array}{c}\beta_0^L\sum_{t=1}^{T^L}(1-\lambda)^{t-1}y_t^L+\sum_{t=1}^{T^L}P_t^L x_t^L-\sum_{t=1}^{T^L}P_{t-1}^L\gamma\\ -(1-\beta_0^L)\sum_{t=1}^{T^H}[x_t^H-\gamma e_t^{L,H}]-\beta_0^L\sum_{t=1}^{T^H}(\prod_{s=1}^{t-1}(1-\lambda e_s^{L,H}))[e_t^{L,H}\lambda y_t^H+(1-\lambda e_t^{L,H})x_t^H-e_t^{L,H}\gamma]\\ -\left(1-\beta_0^L+\beta_0^L\left(\prod_{s=1}^{T^H}(1-\lambda e_s^{L,H})\right)\right)\left(c_{T^H+1}^H-c_{\sum_{s=1}^{T^H}e_s^{L,H}+1}^L\right)q_F\end{array}\right]$$

$$+\sum_{t=1}^{T^H}\mu_t^H\left[y_t^H-x_t^H-\frac{\gamma}{\lambda\beta_t^H}-\sum_{s=t+1}^{T^H}(1-\lambda)^{s-t-1}(\lambda y_s^H+(1-\lambda)x_s^H-\gamma)-\frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F\right]$$

$$+\sum_{t=1}^{T^L}\mu_t^L\left[y_t^L-x_t^L-\frac{\gamma}{\lambda\beta_t^L}-\sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L+(1-\lambda)x_s^L-\gamma)-\frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F\right]$$

$$+\sum_{t=1}^{T^H}\eta_t^H x_t^H+\sum_{t=1}^{T^L}\eta_t^L x_t^L.$$

The relevant Kuhn-Tucker conditions for the optimization problem are:

(A1) $$\frac{\partial\mathcal{L}}{\partial y_t^H}=-\nu\beta_0^H(1-\lambda)^{t-1}\lambda+\xi^H\beta_0^H(1-\lambda)^{t-1}\lambda-\xi^L\beta_0^L(\prod_{s=1}^{t-1}(1-\lambda e_s^{L,H}))\lambda e_t^{L,H}$$
$$+\mu_t^H-\sum_{j=1}^{t-1}\mu_j^H(1-\lambda)^{t-j-1}\lambda=0;$$

(A2) $$\frac{\partial\mathcal{L}}{\partial y_t^L}=-(1-\nu)\beta_0^L(1-\lambda)^{t-1}\lambda-\xi^H\beta_0^H(\prod_{s=1}^{t-1}(1-\lambda e_s^{H,L}))\lambda e_t^{H,L}$$
$$+\xi^L\beta_0^L(1-\lambda)^{t-1}\lambda+\mu_t^L-\sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j-1}\lambda=0;$$

(A3) $$\frac{\partial\mathcal{L}}{\partial x_t^H}=-\nu P_t^H+\xi^H P_t^H-\xi^L(1-\beta_0^L+\beta_0^L\prod_{s=1}^{t}(1-\lambda e_s^{L,H}))$$
$$-\mu_t^H-\sum_{j=1}^{t-1}\mu_j^H(1-\lambda)^{t-j}+\eta_t^H=0;$$

(A4) $$\frac{\partial\mathcal{L}}{\partial x_t^L}=-(1-\nu)P_t^L-\xi^H(1-\beta_0^H+\beta_0^H\prod_{s=1}^{t}(1-\lambda e_s^{H,L}))+\xi^L P_t^L$$
$$-\mu_t^L-\sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j}+\eta_t^L=0.$$

We now characterize the optimal payment structure for each of the 4 cases. As noted above, the optimal length T^θ is derived in Part II below and the sufficient conditions for each of these cases in Part III.

Case 1: Strong moral hazard: Both IC constraints are slack.

If both the (IC) constraints are slack, we find that the moral hazard constraints for both types are binding in every period, and each type is rewarded with a constant payment after success. Neither type is rewarded after failure. This is Case 1 and the results are given in Claim A1.

Claim A1. $\xi^H=\xi^L=0\Rightarrow\eta_t^H,\eta_t^L,\mu_t^H,\mu_t^L>0$, and it is optimal to set $x_t^\theta=0$ and $y_t^\theta=\frac{\gamma}{\lambda\beta_{T^\theta}^\theta}+\frac{(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F$ for $t\le T^\theta$ and $\theta\in\{H,L\}$.

Proof:

$\boldsymbol{\mu_t^\theta>0}$. We first prove that if the two (IC) constraints are slack, then all the (MH_t^θ) constraints for $t\le T^\theta$ and $\theta\in\{H,L\}$ must be binding.

$\boldsymbol{\mu_t^H>0}$. Given that $\xi^H=\xi^L=0$, $(A1)$ at each period $t\le T^H$ can be rewritten as

$t=1$: $-\nu\beta_0^H\lambda+\mu_1^H=0\implies\mu_1^H=\nu\beta_0^H\lambda>0$;

$t = 2$: $-\nu\beta_0^H(1-\lambda)\lambda + \mu_2^H - \lambda\mu_1^H = 0 \implies \mu_2^H = \nu\beta_0^H\lambda > 0$;
Solving recursively for $t = 3, \dots, T^H$ we have $\mu_t^H = \nu\beta_0^H\lambda > 0$ for $t \le T^H$.

Thus, all the (MH_t^H) constraints are binding.

$\boldsymbol{\mu_t^L > 0}$**.** Given that $\xi^H = \xi^L = 0$, $(A2)$ at each period $t \le T^L$ can be rewritten as
$t = 1$: $-(1-\nu)\beta_0^L\lambda + \mu_1^L = 0 \implies \mu_1^L = (1-\nu)\beta_0^L\lambda > 0$;
$t = 2$: $-(1-\nu)\beta_0^L(1-\lambda)\lambda + \mu_2^L - \lambda\mu_1^L = 0 \implies \mu_2^L = (1-\nu)\beta_0^L\lambda > 0$;
Solving recursively for $t = 3, \dots, T^L$ we have $\mu_t^L = (1-\nu)\beta_0^L\lambda > 0$ for $t \le T^L$.

Thus, all the (MH_t^L) constraints are binding.

$\boldsymbol{x_t^\theta = 0}$**.** No rent after failure follows immediately from $(A3)$ and $(A4)$ as $\xi^H = \xi^L = 0$ implies $\eta_t^\theta > 0$ for all $t \le T^\theta$. Furthermore, given that $\mu_t^H = \nu\beta_0^H\lambda$, and $\mu_t^L = (1-\nu)\beta_0^L\lambda$, we can also show that $\eta_t^H = \nu(P_t^H + \beta_0^H\lambda) + \nu\beta_0^H\lambda\sum_{j=1}^{t-1}(1-\lambda)^{t-j} > 0$ for $t \le T^H$, and that $\eta_t^L = (1-\nu)(P_t^L + \beta_0^L\lambda) + (1-\nu)\beta_0^L\lambda\sum_{j=1}^{t-1}(1-\lambda)^{t-j} > 0$ for $t \le T^L$.

Thus, if the (IC) constraints are slack, both types are rewarded only for success and all the $\left(MH_t^\theta\right)$ constraints for $t \le T^\theta$ and $\theta \in \{H, L\}$ are binding:

$$y_t^\theta = \frac{\gamma}{\lambda\beta_t^\theta} + \sum_{s=t+1}^{T^\theta}(1-\lambda)^{s-t-1}\left(\lambda y_s^\theta - \gamma\right) + \frac{(1-\beta_0^\theta)}{P_T^\theta}\Delta c q_F,$$
$$x_t^\theta = 0 \text{ for } t \le T^\theta \text{ and } \theta \in \{H, L\}.$$

Finally, we prove that the unique sequence of y_t^θ that solves the system of binding $\left(MH_t^\theta\right)$ constraints is $y_t^\theta = \frac{\gamma}{\lambda\beta_T^\theta} + \frac{(1-\beta_0^\theta)}{P_T^\theta}\Delta c q_F$. Solving recursively binding $\left(MH_t^\theta\right)$ constraints for y_t^θ we obtain:

$$y_t^\theta = \frac{\gamma}{\lambda\beta_t^\theta} + \gamma\sum_{s=1}^{T-t}\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^{t+s-1}} + \frac{(1-\beta_0^\theta)}{P_T^\theta}\Delta c q_F.$$

We next prove that y_t^θ is constant:

$$y_t^\theta - y_{t+1}^\theta =$$
$$\left[\frac{\gamma}{\lambda\beta_t^\theta} + \gamma\sum_{s=1}^{T-t}\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^{t+s-1}} + \frac{(1-\beta_0^\theta)}{P_T^\theta}\Delta c q_F\right] - \left[\frac{\gamma}{\lambda\beta_{t+1}^\theta} + \gamma\sum_{s=1}^{T-t-1}\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^{t+s}} + \frac{(1-\beta_0^\theta)}{P_T^\theta}\Delta c q_F\right]$$
$$= \frac{\gamma}{\lambda\beta_t^\theta} - \frac{\gamma}{\lambda\beta_{t+1}^\theta} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^{t-1}}\sum_{s=1}^{T-t}\frac{1}{(1-\lambda)^s} - \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t}\sum_{s=1}^{T-t-1}\frac{1}{(1-\lambda)^s}.$$

Using the formula for geometric series, we can rewrite $\sum_{s=1}^{T-t}\frac{1}{(1-\lambda)^s}$ and $\sum_{s=1}^{T-t-1}\frac{1}{(1-\lambda)^s}$ as:

$$\sum_{s=1}^{T-t}\frac{1}{(1-\lambda)^s} = \frac{1}{(1-\lambda)}\left(\frac{1-\frac{1}{(1-\lambda)^{T-t}}}{1-\frac{1}{(1-\lambda)}}\right) = \frac{\left(\frac{1}{(1-\lambda)^{T-t}}-1\right)}{\lambda} = \frac{1-(1-\lambda)^{T-t}}{\lambda(1-\lambda)^{T-t}}$$

$$\text{and } \sum_{s=1}^{T-t-1}\frac{1}{(1-\lambda)^s} = \frac{\left(\frac{1}{(1-\lambda)^{T-t-1}}-1\right)}{\lambda} = \frac{1-(1-\lambda)^{T-t-1}}{\lambda(1-\lambda)^{T-t-1}}.$$

Thus, $y_t^\theta - y_{t+1}^\theta =$

$$= \frac{\gamma}{\lambda\beta_t^\theta} - \frac{\gamma}{\lambda\beta_{t+1}^\theta} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^{t-1}}\left(\frac{1-(1-\lambda)^{T-t}}{\lambda(1-\lambda)^{T-t}}\right) - \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t}\left(\frac{1-(1-\lambda)^{T-t-1}}{\lambda(1-\lambda)^{T-t-1}}\right)$$

$$= \frac{\gamma}{\lambda\beta_t^\theta} - \frac{\gamma(1-\lambda\beta_t^\theta)}{\lambda\beta_t^\theta(1-\lambda)} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta\lambda(1-\lambda)^{T-1}}(1-(1-\lambda)^{T-t} - 1 + (1-\lambda)^{T-t-1})$$

$$= \frac{\gamma}{\lambda\beta_t^\theta} - \frac{\gamma(1-\lambda\beta_t^\theta)}{\lambda\beta_t^\theta(1-\lambda)} + \gamma\frac{(1-\beta_0^\theta)(1-\lambda)^{T-t-1}}{\beta_0^\theta\lambda(1-\lambda)^{T-1}}\lambda = \gamma\frac{(1-\lambda)-(1-\lambda\beta_t^\theta)}{\lambda\beta_t^\theta(1-\lambda)} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t}$$

$$= -\gamma\frac{\lambda(1-\beta_0^\theta)}{\lambda\beta_0^\theta(1-\lambda)} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t} = -\gamma\frac{\lambda\left(\frac{1-\beta_0^\theta}{\beta_0^\theta(1-\lambda)^{t-1}+1-\beta_0^\theta}\right)}{\lambda\left(\frac{\beta_0^\theta(1-\lambda)^{t-1}}{\beta_0^\theta(1-\lambda)^{t-1}+1-\beta_0^\theta}\right)(1-\lambda)} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t}$$

$$= -\gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t} + \gamma\frac{(1-\beta_0^\theta)}{\beta_0^\theta(1-\lambda)^t} = 0.$$

Given that that y_t^θ is constant, we can rewrite y_t^θ by evaluating it at $t = T$:

$$y_t^\theta = \frac{\gamma}{\lambda\beta_T^\theta} + \frac{(1-\beta_0^\theta)}{P_T^\theta}\Delta c q_F.$$

This concludes the proof of Claim A1. *Q.E.D.*

We now prove a lemma characterizing off the path effort for the high type before solving Case 2.

Lemma 1. A lying high type works off the equilibrium path: $e_t^{H,L} = 1$ for $t \leq T^L$.

Proof: Consider the high type's incentives to engage in a one-shot deviation and shirk at period $t \leq T^L$ after accepting a contract designed for the low type. Upon deviating only at some period $t \leq T^L$, his continuation value from the relationship is

$$x_t^L + \beta_t^H \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}[\lambda y_s^L + (1-\lambda)x_s^L - \gamma] + (1-\beta_t^H)\sum_{s=t+1}^{T^L}(x_s^L - \gamma)$$
$$+\left(1-\beta_t^H + \beta_t^H(1-\lambda)^{T^L-t}\right)\left(c_{T^L+1}^L - c_{T^L}^H\right)q_F.$$

In contrast, if the lying high type decides to work at period t, his continuation value from the relationship becomes

$$-\gamma + \lambda\beta_t^H y_t^L + (1-\lambda\beta_t^H)x_t^L + \beta_t^H\sum_{s=t+1}^{T^L}(1-\lambda)^{s-t}[\lambda y_s^L + (1-\lambda)x_s^L - \gamma] +$$
$$(1-\beta_t^H)\sum_{s=t+1}^{T^L}(x_s^L - \gamma) + \left(1-\beta_t^H + \beta_t^H(1-\lambda)^{T^L-t+1}\right)\left(c_{T^L+1}^L - c_{T^L+1}^H\right)q_F.$$

By combining the two continuation values presented above, we can write a one-period moral hazard constraint at period t below for the lying high type, which we denote by $\left(MH_t^{H,L}\right)$:

$$\left(MH_t^{H,L}\right)\ y_t^L - x_t^L \geq \frac{\gamma}{\lambda\beta_t^H} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L + (1-\lambda)x_s^L - \gamma)$$
$$+\frac{(1-\beta_0^L)(1-\lambda)^{T^L-t}}{P_{T^L}^L}\Delta c q_F.$$

The low type's contract has to satisfy the sequence of local (MH_t^L) constraints:

$$(MH_t^L)\ y_t^L - x_t^L \geq \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L + (1-\lambda)x_s^L - \gamma)$$

$$+\frac{(1-\lambda)^{T^L-t}(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F.$$

Since $\beta_t^L < \beta_t^H$ for any t, we have $\frac{\gamma}{\lambda\beta_t^H} < \frac{\gamma}{\lambda\beta_t^L}$, and the RHS of $\left(MH_t^{H,L}\right)$ is smaller than the RHS of (MH_t^L). Thus, $\left(MH_t^{H,L}\right)$ is implied by (MH_t^L) for $t \leq T^L$.
This concludes the proof of Lemma 1. *Q.E.D.*

Case 2: Strong adverse selection: Both (IC^H) and (IC^L) are binding

We now characterize the payment structure for Case 2. If (IC^L) is binding, all the (MH_t^L) constraints are binding as well.[37]

Claim A2. $\xi^H > 0$, $\xi^L > 0 \Rightarrow \eta_t^H > 0 = \mu_t^H$ for $t \leq T^H$, $\mu_t^L > 0$ for $t \leq T^L$, $\eta_t^L > 0 = \eta_{T^L}^L$ for $t < T^L$. It is optimal to set

$$U^H\left(\varpi^L, \vec{1}\right) = \beta_0^H \sum_{t=1}^{T^H}(1-\lambda)^{t-1}(\lambda y_t^H - \gamma) - \gamma \sum_{t=1}^{T^H} P_t^H,$$

$$y_{T^L}^L = x_{T^L}^L + \frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F,$$

$$y_t^L = \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L - \gamma) + (1-\lambda)^{T^L-t}x_{T^L}^L + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \text{ for } t < T^L,$$

and

$$U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right) = \beta_0^L \sum_{t=1}^{T^L}(1-\lambda)^{t-1}(\lambda y_t^L - \gamma) - \gamma \sum_{t=1}^{T^L-1} P_t^L + P_{T^L}^L\left(x_{T^L}^L - \gamma\right).$$

Proof:
H-type. We first prove that the high type's rewards after success can be distributed such that all the (MH_t^H) constraints are satisfied at no additional cost, i.e., $\mu_t^H = 0$ for $t \leq T^H$.

$\boldsymbol{\mu_t^H = 0}$**.** There exists a solution to $(A1)$ and $(A3)$ for $t \leq T^H$ such that for all $t \leq T^H$:

$$\mu_t^H = 0;$$

$$\xi^L = \eta_t^H \frac{\beta_0^H(1-\lambda)^{t-1}\lambda}{\left(\left(1-\beta_0^L+\beta_0^L\prod_{s=1}^{t}\left(1-\lambda e_s^{L,H}\right)\right)\beta_0^H(1-\lambda)^{t-1}\lambda - \beta_0^L\left(\prod_{s=1}^{t-1}\left(1-\lambda e_s^{L,H}\right)\right)\lambda e_t^{L,H} P_t^H\right)} > 0;$$

$$\xi^H \beta_0^H(1-\lambda)^{t-1}\lambda P_t^H = \nu\beta_0^H(1-\lambda)^{t-1}\lambda P_t^H + \xi^L\beta_0^L\left(\prod_{s=1}^{t-1}\left(1-\lambda e_s^{L,H}\right)\right)\lambda e_t^{L,H} P_t^H > 0.$$

Therefore, the principal sets $x_t^H = 0$ and uses any combination of y_t^H such that

$$y_t^H \geq \frac{\gamma}{\lambda\beta_t^H} + \sum_{s=t+1}^{T^H}(1-\lambda)^{s-t-1}(\lambda y_s^H - \gamma) + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F \text{ for } t \leq T^H, \text{ and}$$

$$\beta_0^H \sum_{t=1}^{T^H}(1-\lambda)^{t-1}(\lambda y_t^H - \gamma) - \gamma \sum_{t=1}^{T^H} P_t^H = U^H\left(\varpi^L, \vec{1}\right).$$

L-type. We now describe the optimal contract for the low type.
$\boldsymbol{\mu_t^L > 0}$ **for** $\boldsymbol{t \leq T^L}$**.** Combining $(A2)$ and $(A4)$ we have:[38]

---

[37] We ignore the knife-edge case where both the (IC) constraints and all the (MH_t^H) are binding simultaneously. In that case, the adverse selection rent is exactly equal to the moral hazard rent and there is no extra rent to be paid.
[38] We multiply (A4) by $\beta_0^L(1-\lambda)^{t-1}\lambda$ and subtract it from (A2) multiplied by P_t^L.

$$\xi^H \lambda \begin{bmatrix} \beta_0^L(1-\beta_0^H)(1-\lambda)^{t-1} - \beta_0^H(1-\beta_0^L)e_t^{H,L}\prod_{s=1}^{t-1}(1-\lambda e_s^{H,L}) \\ +\beta_0^L\beta_0^H(1-\lambda)^{t-1}(1-e_t^{H,L})\prod_{s=1}^{t-1}(1-\lambda e_s^{H,L}) \end{bmatrix} + \mu_t^L P_{t-1}^L$$
$$= \beta_0^L(1-\lambda)^{t-1}\lambda\eta_t^L + (1-\beta_0^L)\sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j-1}\lambda \text{ for } t \le T^L.$$

Since that $e_t^{H,L} = 1$ for $t \le T^L$ by Lemma 1, the condition above simplifies to

$$(\boldsymbol{A5}) \qquad -\xi^H\lambda(1-\lambda)^{t-1}(\beta_0^H - \beta_0^L) + \mu_t^L P_{t-1}^L$$
$$= \beta_0^L(1-\lambda)^{t-1}\lambda\eta_t^L + (1-\beta_0^L)\sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j-1}\lambda \text{ for } t \le T^L.$$

Given that the *RHS* of $(A5)$ is non-negative, we have, $\mu_t^L > 0$ for every $t \le T^L$ and, as a result, the (MH_t^L) constraints must be binding:

$$y_t^L - x_t^L = \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L + (1-\lambda)x_s^L - \gamma) + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \text{ for } t \le T^L.$$

$\boldsymbol{x_{T^L}^L > x_t^L = 0}$ **for** $\boldsymbol{t < T^L}$**.** First, we prove by contradiction that the low type is rewarded for failure in only one period $s \le T^L$. Second, we prove that it is optimal to reward the low type for the very last failure, i.e., $s = T^L$.

Step1: Suppose that the low type is rewarded for failures in two distinct periods $z \le T^L$ and $k \le T^L$, such that $z < k$. Therefore, $x_z^L, x_k^L > 0$ $(\eta_z^L = 0 = \eta_k^L)$. Evaluating (A4) at $t = z$ and $t = k$, and using $e_t^{H,L} = 1$ for all t, we derive

$$-(1-\nu)P_z^L - \xi^H P_z^H + \xi^L P_z^L$$
$$-\mu_z^L - \sum_{j=1}^{z-1}\mu_j^L(1-\lambda)^{z-j} = 0 = -\mu_k^L - \sum_{j=1}^{k-1}\mu_j^L(1-\lambda)^{k-j}$$
$$-(1-\nu)P_k^L - \xi^H P_k^H + \xi^L P_k^L,$$

$$\frac{P_z^L[\xi^L - (1-\nu)] - \sum_{j=1}^{z}\mu_j^L(1-\lambda)^{z-j}}{P_z^H} = \xi^H = \frac{P_k^L[\xi^L - (1-\nu)] - \sum_{j=1}^{k}\mu_j^L(1-\lambda)^{k-j}}{P_k^H},$$

$$P_k^H P_z^L[\xi^L - (1-\nu)] - P_k^H\sum_{j=1}^{z}\mu_j^L(1-\lambda)^{z-j} = P_z^H P_k^L[\xi^L - (1-\nu)] - P_z^H\sum_{j=1}^{k}\mu_j^L(1-\lambda)^{k-j},$$

$$P_k^H\sum_{j=1}^{z}\mu_j^L(1-\lambda)^{z-j} = P_z^H\sum_{j=1}^{k}\mu_j^L(1-\lambda)^{k-j},$$

which leads to a contradiction since, $P_k^H < P_z^H$ and $\sum_{j=1}^{z}\mu_j^L(1-\lambda)^{z-j} < \sum_{j=1}^{k}\mu_j^L(1-\lambda)^{k-j}$ if $k > z$.[39] That is, the left-hand side of the inequality above is less than the right-hand side.

Step2: We now prove that the low type is rewarded for failure in the last period only.

Expressing $\sum_{t=1}^{T^H}(1-\lambda)^{t-1}\lambda y_t^H$ from the (IC^H):

$$\sum_{t=1}^{T^H}(1-\lambda)^{t-1}\lambda y_t^H = \frac{P_s^H x_s^L + P_{T^L}^H \Delta c_{T^L+1} q_F}{\beta_0^H} + \frac{\gamma\sum_{t=1}^{T^H}P_{t-1}^H - \gamma\sum_{t=1}^{T^L}P_{t-1}^H}{\beta_0^H} + \sum_{t=1}^{T^L}(1-\lambda)^{t-1}\lambda y_t^L,$$

it is immediate that the high type's rent is decreasing in s through $P_s^H x_s^L$ and, therefore, rewarding the last type's failure mitigates the high type's rent, $s = T^L$.

Similarly, from the (IC^L):

$$\beta_0^L\sum_{t=1}^{T^L}(1-\lambda)^{t-1}y_t^L + P_s^L x_s^L = \beta_0^L\sum_{t=1}^{T^H}\left(\prod_{s=1}^{t-1}(1-\lambda e_s^{L,H})\right)[e_t^{L,H}\lambda y_t^H - e_t^{L,H}\gamma] + \sum_{t=1}^{T^L}P_{t-1}^L\gamma$$
$$-\gamma(1-\beta_0^L)\sum_{t=1}^{T^H}e_t^{L,H} + \left(1 - \beta_0^L + \beta_0^L\left(\prod_{s=1}^{T^H}(1-\lambda e_s^{L,H})\right)\right)\left(c_{T^H+1}^H - c_{\sum_{s=1}^{T^H}e_s^{L,H}+1}^L\right)q_F,$$

---

[39] A similar contradiction would emerge if we assumed $k < z$.

the low type's rent is decreasing in s through $\beta_0^L \sum_{t=1}^{T^H}\left(\prod_{s=1}^{t-1}\left(1-\lambda e_s^{L,H}\right)\right)\left[e_t^{L,H}\lambda y_t^H - e_t^{L,H}\gamma\right]$ and, therefore, rewarding the last type's failure mitigates the low type's rent as well, $s = T^L$.

Therefore, the low type's payments are determined by

$$y_{T^L}^L = x_{T^L}^L + \frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F,$$

$$y_t^L = \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L - \gamma) + (1-\lambda)^{T^L-t}x_{T^L}^L + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \text{ for } t < T^L, \text{ and}$$

$$U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right) = \beta_0^L \sum_{t=1}^{T^L}(1-\lambda)^{t-1}(\lambda y_t^L - \gamma) - \gamma \sum_{t=1}^{T^L-1} P_t^L + P_{T^L}^L\left(x_{T^L}^L - \gamma\right).$$

$\boldsymbol{\vec{a}^L(\varpi^H)}$**.** It is without loss of generality to characterize off-the-equilibrium path effort as a stopping rule $\vec{a}^L(\varpi^H)$: the low type works up to period $t^{L,H} \leq T^H$ and shirks after.[40]
The reason is two-fold. First, it is without loss of generality to front load the high type's rewards for success because the relative likelihood of success $\frac{\beta_0^L(1-\lambda)^{t-1}\lambda}{\beta_0^H(1-\lambda)^{t-1}\lambda} = \frac{\beta_0^L}{\beta_0^H}$ is independent of t. Then, the principal can use any combination of y_t^H to pay rent to the high type without affecting the low type's incentives to work off-the-equilibrium path. Second, the low type's probability of success in any period and the expected cost after failure depends on the total number of failures up to that period (not on when those failures occurred).
Therefore, the low type's expected payoff when he pretends being high can be written as:

$$U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right) = -\gamma(1-\beta_0^L)t^{L,H} + \beta_0^L \sum_{t=1}^{t^{L,H}}(1-\lambda)^{t-1}\,(\lambda y_t^H - \gamma) + P_{t^{L,H}}^L\left(c_{T^H+1}^H - c_{t^{L,H}+1}^L\right)q_F,$$

where $t^{L,H}$ is the number of time periods that maximizes $U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right)$, for the optimal contract ϖ^H in Case 2:

$$t^{L,H} \coloneqq \arg\max_{0 \leq t^{L,H} \leq T^H} U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right).$$

This concludes the proof of Claim A2. *Q.E.D.*

Intermediate case 3: only (IC^H) constraint binds

If (IC^H) binds and (IC^L) is slack, the (MH_t^L) are binding in each period but (MH_t^H) are all slack.

Claim A3. $\xi^H > 0$, $\xi^L = 0 \Rightarrow \eta_t^L, \mu_t^L > 0$ and $\eta_t^H = \mu_t^H = 0$. It is optimal to set $x_t^L = 0$ and $y_t^L = \frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F$ for $t \leq T^L$ and any combination of x_t^H and y_t^H such that $y_t^H - x_t^H \geq \frac{\gamma}{\lambda\beta_t^H} + \sum_{s=t+1}^{T^H}(1-\lambda)^{s-t-1}(\lambda y_s^H + (1-\lambda)x_s^H - \gamma) + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F$ for $t \leq T^H$ and $U^H\left(\varpi^L, \vec{1}\right) = (1-\beta_0^H)\sum_{t=1}^{T^H}[x_t^H - \gamma] + \beta_0^H \sum_{t=1}^{T^H}(1-\lambda)^{t-1}[(\lambda y_t^H - \gamma) + (1-\lambda)x_t^H]$ is given by the binding (IC^H) constraint.
Proof:

[40] Alternatively, we could write that the agent worked for $t^{L,H}/T^H$ periods, but the notation would be cumbersome as we would need to indicate the periods he works in.

L-type. We first prove that the low type is rewarded only for success and all the (MH_t^L) constraints are binding for $t \le T^L$.

$\boldsymbol{\mu_t^L > 0}$**.** Given that $\xi^L = 0$ and $\xi^H > 0$, condition $(A2)$ at each period $t \le T^L$ can be rewritten as

$$\mu_t^L = (1-\nu)\beta_0^L(1-\lambda)^{t-1}\lambda + \xi^H\beta_0^H\left(\prod_{s=1}^{t-1}(1-\lambda e_s^{H,L})\right)\lambda e_t^{H,L}$$
$$+\sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j-1}\lambda > 0 \text{ for } t \le T^L.$$

Thus, all the (MH_t^L) constraints are binding.

We next prove that the low type is rewarded only for success, i.e., $x_t^L = 0$ for $t \le T^L$.

$\boldsymbol{x_t^L = 0}$**.** Given that $\xi^L = 0$ and $\xi^H > 0$ condition $(A4)$ at each period $t \le T^L$ can be rewritten as

$$\eta_t^L = (1-\nu)P_t^L + \xi^H\left(1-\beta_0^H + \beta_0^H\prod_{s=1}^{t}(1-\lambda e_s^{H,L})\right)$$
$$+\mu_t^L + \sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j} > 0 \text{ for } t \le T^L.$$

Therefore, $\eta_t^L > 0$ for every $t \le T^L$ and, as a result, the low type is not rewarded for failures:

$$x_t^L = 0 \text{ for } t \le T^L.$$

Thus, the low type is rewarded only for success with the rewards given by:

$$y_t^L = \frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \text{ for } t \le T^L.$$

H-type. We next characterize the optimal contract for the high type.

$\boldsymbol{\vec{a}^H(\varpi^L) = \vec{1}}$**.** In Lemma 1, we proved that the high type never shirks off-the-equilibrium path. We next prove that the principal can use any combination of x_t^H and y_t^H such that all (MH_t^H) constraints are satisfied and the high type' expected rent is $U^H(\varpi^L, \vec{1})$.

$\boldsymbol{\mu_t^H = \eta_t^H = 0}$**.** Given that $\xi^L = 0$, conditions $(A1)$ and $(A3)$ can be rewritten as

$$(A1') \quad \frac{\partial \mathcal{L}}{\partial y_t^H} = -\nu\beta_0^H(1-\lambda)^{t-1}\lambda + \xi^H\beta_0^H(1-\lambda)^{t-1}\lambda + \mu_t^H - \sum_{j=1}^{t-1}\mu_j^H(1-\lambda)^{t-j-1}\lambda = 0;$$

$$(A3') \quad \frac{\partial \mathcal{L}}{\partial x_t^H} = -\nu P_t^H + \xi^H P_t^H - \mu_t^H - \sum_{j=1}^{t-1}\mu_j^H(1-\lambda)^{t-j} + \eta_t^H = 0.$$

There exists a solution to $(A1')$ and $(A3')$ for $t \le T^H$ such that

$$\mu_t^H = \eta_t^H = 0 \text{ and } \xi^H = \nu \text{ for } t \le T^H.$$

Therefore, the principal can use any combination of x_t^H and y_t^H such that

$$y_t^H - x_t^H \ge \frac{\gamma}{\lambda\beta_t^H} + \sum_{s=t+1}^{T^H}(1-\lambda)^{s-t-1}(\lambda y_s^H + (1-\lambda)x_s^H - \gamma) + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F \text{ for } t \le T^H, \text{ and}$$

$$(1-\beta_0^H)\sum_{t=1}^{T^H}[x_t^H - \gamma] + \beta_0^H\sum_{t=1}^{T^H}(1-\lambda)^{t-1}[(\lambda y_t^H - \gamma) + (1-\lambda)x_t^H] = U^H(\varpi^L, \vec{1}).$$

This concludes the proof of Claim A3. *Q.E.D.*

Intermediate case 4: only (IC^L) constraint binds

If (IC^L) binding and (IC^H) is slack, the (MH_t^H) are all binding but the (MH_t^L) are all slack.

Claim A4. $\xi^H = 0$, $\xi^L > 0 \Rightarrow \eta_t^H, \mu_t^H > 0$ and $\mu_t^L = 0$, $\eta_t^L = 0$. It is optimal to set $x_t^H = 0$ and $y_t^H = \frac{\gamma}{\lambda\beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F$ for $t \le T^H$ and any combination of x_t^L and y_t^L such that $y_t^L - x_t^L \ge \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L + (1-\lambda)x_s^L - \gamma) + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F$ for $t \le T^L$ and

$U^L\big(\varpi^H,\vec{a}^L(\varpi^H)\big) = (1-\beta_0^L)\sum_{t=1}^{T^L}[x_t^L-\gamma] + \beta_0^L\sum_{t=1}^{T^L}(1-\lambda)^{t-1}[(\lambda y_t^L-\gamma)+(1-\lambda)x_t^L]$ is given by the binding (IC^L) constraint.

Proof:

H-type. We first prove that the high type is rewarded only for success and all the (MH_t^H) constraints are binding for $t \le T^H$.

$\boldsymbol{\mu_t^H > 0}$**.** Given that $\xi^H = 0$ and $\xi^L > 0$, condition $(A1)$ at each period $t \le T^H$ can be rewritten as

$$\mu_t^H = \nu\beta_0^H(1-\lambda)^{t-1}\lambda + \xi^L\beta_0^L\big(\textstyle\prod_{s=1}^{t-1}(1-\lambda e_s^{L,H})\big)\lambda e_t^{L,H}$$
$$+\textstyle\sum_{j=1}^{t-1}\mu_j^H(1-\lambda)^{t-j-1}\lambda > 0 \text{ for } t \le T^H.$$

Thus, all the (MH_t^H) constraints are binding.

We next prove that the high type is rewarded only for success, i.e., $x_t^H = 0$ for $t \le T^H$.

$\boldsymbol{x_t^H = 0}$**.** Given that $\xi^H = 0$ and $\xi^L > 0$ condition $(A3)$ at each period $t \le T^H$ can be rewritten as

$$\eta_t^H = \nu P_t^H + \xi^L\big(1-\beta_0^L+\beta_0^L\textstyle\prod_{s=1}^{t}(1-\lambda e_s^{L,H})\big)$$
$$+\mu_t^H + \textstyle\sum_{j=1}^{t-1}\mu_j^H(1-\lambda)^{t-j} > 0 \text{ for } t \le T^H.$$

Therefore, $\eta_t^H > 0$ for every $t \le T^H$ and, as a result, the high type is not rewarded for failures:

$$x_t^H = 0 \text{ for } t \le T^H.$$

Thus, the high type is rewarded only for success with the rewards given by:

$$y_t^H = \frac{\gamma}{\lambda\beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F \text{ for } t \le T^H.$$

L-type. We next characterize the optimal contract for the low type.

$\boldsymbol{\vec{a}^L(\varpi^H)}$**.** As in Case 2, it is again without loss of generality to characterize off-the-equilibrium path effort as a stopping rule $\vec{a}^L(\varpi^H)$: the low type works up to period $t^{L,H} \le T^H$ and shirks after. Note that $t^{L,H}$ is now given by the optimal contract ϖ^H in intermediate case 4:

$$t^{L,H} \coloneqq \arg\max_{0\le t^{L,H}\le T^H} U^L\big(\varpi^H,\vec{a}^L(\varpi^H)\big).$$

Then, the binding (IC^L) is given by,

$$\beta_0^L\sum_{t=1}^{T^L}(1-\lambda)^{t-1}y_t^L + \sum_{t=1}^{T^L}P_t^L x_t^L - \sum_{t=1}^{T^L}P_{t-1}^L\gamma =$$
$$\beta_0^L\sum_{t=1}^{t^{L,H}}(1-\lambda)^{t-1}\lambda y_t^H - \sum_{t=1}^{t^{L,H}}P_{t-1}^L\gamma + P_{t^{L,H}}^L\big(c_{T^H+1}^H - c_{t^{L,H}+1}^L\big)q_F.$$

Next, we prove that the principal can use any combination of x_t^L and y_t^L such that

$$y_t^L - x_t^L \ge \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L + (1-\lambda)x_s^L - \gamma) + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \text{ for } t \le T^L, \text{ and}$$

$$(1-\beta_0^L)\sum_{t=1}^{T^L}[x_t^L-\gamma] + \beta_0^L\sum_{t=1}^{T^L}(1-\lambda)^{t-1}[(\lambda y_t^L-\gamma)+(1-\lambda)x_t^L] = U^L\big(\varpi^H,\vec{a}^L(\varpi^H)\big).$$

$\boldsymbol{\mu_t^L = \eta_t^L = 0}$**.** Given that $\xi^H = 0$, conditions $(A2)$ and $(A4)$ can be rewritten as

(A2) $\frac{\partial\mathcal{L}}{\partial y_t^L} = -(1-\nu)\beta_0^L(1-\lambda)^{t-1}\lambda + \xi^L\beta_0^L(1-\lambda)^{t-1}\lambda + \mu_t^L - \sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j-1}\lambda = 0;$

(A4) $\frac{\partial\mathcal{L}}{\partial x_t^L} = -(1-\nu)P_t^L + \xi^L P_t^L - \mu_t^L - \sum_{j=1}^{t-1}\mu_j^L(1-\lambda)^{t-j} + \eta_t^L = 0.$

There exists a solution to $(A2)$ and $(A4)$ for $t \le T^L$ such that

$$\mu_t^L = \eta_t^L = 0 \text{ and } \xi^L = (1-\nu) \text{ for } t \le T^L.$$

Therefore, the principal can use any combination of x_t^L and y_t^L such that

$$y_t^L - x_t^L \geq \frac{\gamma}{\lambda\beta_t^L} + \sum_{s=t+1}^{T^L}(1-\lambda)^{s-t-1}(\lambda y_s^L + (1-\lambda)x_s^L - \gamma) + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \text{ for } t \leq T^L, \text{ and}$$

$$(1-\beta_0^L)\sum_{t=1}^{T^L}[x_t^L - \gamma] + \beta_0^L \sum_{t=1}^{T^L}(1-\lambda)^{t-1}[(\lambda y_t^L - \gamma) + (1-\lambda)x_t^L] = U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right).$$

This concludes the proof of Claim A4. *Q.E.D.*

For later use, we collect results we already derived in Case 2 and Intermediate case 3 for off the path effort $\boldsymbol{\vec{a}^L(\varpi^H)}$.

Section II Optimal length of experimentation

Case 1: Both the (IC^H) and (IC^L) constraints are slack (*under* experimentation for both types).

Information rents for both types are given by

$$U^\theta = \beta_0^\theta \sum_{t=1}^{T^\theta}(1-\lambda)^{t-1}\lambda y_t^\theta, \text{ where } y_t^\theta = \frac{\gamma}{\lambda\beta_{T^\theta}^\theta} + \frac{(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F \text{ for } t \leq T^\theta \text{ and } \theta \in \{H, L\}.$$

Since the agent's information rent is increasing in T^θ, there will be *under* experimentation for both types, that is, $T_{SB}^\theta < T_{FB}^\theta$ for $\theta \in \{H, L\}$.

Case 2: Both (IC^H) and (IC^L) bind ($T_{SB}^H < T_{FB}^H$; if γ is low enough, over-experimentation in T^L, $T_{SB}^L > T_{FB}^L$).

Information rents for both types are given by

$$U^H\left(\varpi^L, \vec{1}\right) = (1-\beta_0^H)\sum_{t=1}^{T^L}[x_t^L - \gamma] + \beta_0^H \sum_{t=1}^{T^L}(1-\lambda)^{t-1}[(\lambda y_t^L - \gamma) + (1-\lambda)x_t^L]$$
$$+P_{T^L}^H \Delta c_{T^L+1} q_F;$$

$$U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right) = -\gamma(1-\beta_0^L)t^{L,H} + \beta_0^L \sum_{t=1}^{t=t^{L,H}}(1-\lambda)^{t-1}[\lambda y_t^H - \gamma]$$
$$-P_{t^{L,H}}^H\left(c_{t^{L,H}+1}^L - c_{T^H+1}^H\right)q_F.$$

We next prove that, if γ is low enough, we have over-experimentation in T^L ($T_{SB}^L > T_{FB}^L$).

Consider the term $P_{T^L}^H \Delta c_{T^L+1} q_F$ which, as we prove next, is monotonically decreasing in T^L. Noting that $\Delta c_t = c_t^L - c_t^H = (\beta_t^H - \beta_t^L)(\overline{c} - \underline{c})$, and that $\beta_t^H - \beta_t^L = \frac{\beta_0^H(1-\lambda)^{t-1}}{\beta_0^H(1-\lambda)^{t-1}+1-\beta_0^H} - \frac{\beta_0^L(1-\lambda)^{t-1}}{\beta_0^L(1-\lambda)^{t-1}+1-\beta_0^L} = \frac{(1-\lambda)^{t-1}(\beta_0^H-\beta_0^L)}{P_{t-1}^H P_{t-1}^L}$, the difference in the expected cost can be rewritten as

$$\Delta c_{T^L+1} = \frac{(1-\lambda)^{T^L}(\beta_0^H-\beta_0^L)}{P_{T^L}^H P_{T^L}^L}\left(\overline{c} - \underline{c}\right).$$

Thus, $P_{T^L}^H \Delta c_{T^L+1} = \frac{(1-\lambda)^{T^L}(\beta_0^H-\beta_0^L)}{P_{T^L}^L}\left(\overline{c} - \underline{c}\right) = \frac{(1-\lambda)^{T^L}(\beta_0^H-\beta_0^L)}{\beta_0^L(1-\lambda)^{T^L}+1-\beta_0^L}\left(\overline{c} - \underline{c}\right) = \frac{(\beta_0^H-\beta_0^L)}{\beta_0^L+\frac{1-\beta_0^L}{(1-\lambda)^{T^L}}}\left(\overline{c} - \underline{c}\right),$

which is decreasing in T^L.

The first two terms of $U^H(\varpi^L, \vec{1})$, the 'moral hazard component' on the right-hand side of (IC^H) are increasing in T^L and is negatively proportional to γ. Therefore, the high-type's rent is lowered by increasing T^L if the effect of $P^H_{T^L}\Delta c_{T^L+1}$ dominates, which is the case if γ is low. We define a low enough value of γ, called $\gamma^{2,MH}$, such that the high-type's rent is decreasing in T^L if $\gamma < \gamma^{2,MH}$.

The term $P^H_{t^{L,H}}\left(c^L_{t^{L,H}+1} - c^H_{T^H+1}\right)q_F$ is monotonically decreasing in T^H. Since it appears in the low-type's rent with minus, the high-type's rent is lowered by decreasing T^H. Similarly, the moral hazard component on the right-hand side of (IC^L) is increasing in T^H. Thus, it is optimal to have under-experimentation in T^H ($T^H_{SB} < T^H_{FB}$).

Intermediate case 3: only (IC^H) constraint binds ($T^H_{SB} = T^H_{FB}$; if γ is high enough, over experimentation in T^L, $T^L_{SB} > T^L_{FB}$).

Information rents for both types are given by

$$\lambda y^\theta_t - \gamma = \frac{\gamma}{\beta^\theta_{T^\theta}} + \frac{\lambda\left(1-\beta^\theta_0\right)}{P^\theta_{T^\theta}}\Delta c q_F - \gamma = \frac{\lambda\left(1-\beta^\theta_0\right)}{P^\theta_{T^\theta}}\Delta c q_F + \gamma\left(\frac{1-\beta^\theta_0}{\beta^\theta_0(1-\lambda)^{T^\theta-1}}\right)$$

$$U^L = \beta^L_0 \textstyle\sum_{t=1}^{T^L}(1-\lambda)^{t-1}\lambda y^L_t, \text{ and}$$
$$U^H\left(\varpi^L, \vec{1}\right) = \beta^H_0 \textstyle\sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left(\lambda y^L_t - \gamma\right) - \gamma\left(1-\beta^H_0\right)T^L + P^H_{T^L}\Delta c_{T^L+1}q_F,$$

where $y^L_t = \frac{\gamma}{\lambda\beta^L_{T^L}} + \frac{(1-\beta^L_0)}{P^L_{T^L}}\Delta c q_F$ for $t \le T^L$.

The stopping time for the high type, T^H, does not affect information rents and, as a result, is not distorted: $T^H_{SB} = T^H_{FB}$.

We next prove that, if γ is high enough, we have over experimentation in T^L, $T^L_{SB} > T^L_{FB}$. Consider the RHS of (IC^H):

$$-\gamma\left(1-\beta^H_0\right)T^L + \beta^H_0 \textstyle\sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left[\lambda y^L_t - \gamma\right] + P^H_{T^L}\Delta c_{T^L+1}q_F$$
$$= -\gamma\left(1-\beta^H_0\right)T^L + \beta^H_0 \textstyle\sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left[\lambda\left[\frac{\gamma}{\lambda\beta^L_{T^L}} + \frac{(1-\beta^L_0)}{P^L_{T^L}}\Delta c q_F\right] - \gamma\right] + P^H_{T^L}\Delta c_{T^L+1}q_F$$
$$= -\gamma\left(1-\beta^H_0\right)T^L + \beta^H_0 \textstyle\sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left[\left[\frac{\gamma}{\beta^L_{T^L}} + \frac{\lambda(1-\beta^L_0)}{P^L_{T^L}}\Delta c q_F\right] - \gamma\right] + P^H_{T^L}\Delta c_{T^L+1}q_F.$$

Given that, $\frac{\gamma}{\beta^L_{T^L}} - \gamma = \gamma\left[\frac{1-\beta^L_0}{\beta^L_0(1-\lambda)^{T^L-1}}\right]$ and $\sum_{t=1}^{T^H}(1-\lambda)^{t-1} = \frac{1-(1-\lambda)^{T^H}}{\lambda}$, we rewrite the RHS of (IC^H) as follows:

$$-\gamma\left(1-\beta^H_0\right)T^L + \beta^H_0 \sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left[\left[\frac{\gamma}{\beta^L_{T^L}} + \frac{\lambda(1-\beta^L_0)}{P^L_{T^L}}\Delta c q_F\right] - \gamma\right] + P^H_{T^L}\Delta c_{T^L+1}q_F$$
$$= -\gamma\left(1-\beta^H_0\right)T^L + \beta^H_0 \textstyle\sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left[\gamma\left[\frac{1-\beta^L_0}{\beta^L_0(1-\lambda)^{T^L-1}}\right] + \frac{\lambda(1-\beta^L_0)}{P^L_{T^L}}\Delta c q_F\right] + P^H_{T^L}\Delta c_{T^L+1}q_F$$
$$-\gamma\left(1-\beta^H_0\right)T^L + \gamma\left[\frac{(1-\beta^L_0)\beta^H_0}{\beta^L_0\lambda(1-\lambda)^{T^L-1}}\right] - \gamma\left[\frac{(1-\beta^L_0)\beta^H_0}{\beta^L_0\lambda}\right] - \left[\frac{\beta^L_0(1-\lambda)^{T^L}(1-\beta^H_0)}{P^L_{T^L}}\right]\Delta c q_F.$$

Since the second term $\gamma\left[\frac{(1-\beta_0^L)\beta_0^H}{\beta_0^L\lambda(1-\lambda)^{T^L-1}}\right]$ is increasing in T^L and the last term is decreasing in T^L, there exist a high enough value of γ, such that the first effect dominates. Then, it is optimal to have over experimentation in T^L, $T_{SB}^L > T_{FB}^L$, to mitigate the high type's rent.

Intermediate case 4: only (IC^L) constraint binds ($T_{SB}^L = T_{FB}^L$, *under* experimentation for the high type, $T_{SB}^H < T_{FB}^H$).

Information rents for both types are given by

$$U^H = \beta_0^H \sum_{t=1}^{T^H}(1-\lambda)^{t-1}\lambda y_t^H, \text{ and}$$

$$U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right) = -\gamma(1-\beta_0^L)t^{L,H} + \beta_0^L \sum_{t=1}^{t^{L,H}}(1-\lambda)^{t-1}\left(\lambda y_t^H - \gamma\right) - P_{t^{L,H}}^H\left(c_{t^{L,H}+1}^L - c_{T^H+1}^H\right)q_F,$$

where $y_t^H = \frac{\gamma}{\lambda\beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F$ for $t \le T^H$.

The stopping time for the low type, T^L, does not affect information rents and, as a result, is not distorted: $T_{SB}^L = T_{FB}^L$.

The term $P_{t^{L,H}}^H\left(c_{t^{L,H}+1}^L - c_{T^H+1}^H\right)q_F$ is monotonically decreasing in T^H. Thus, under experimentation in T^H lowers the rent of the low type. Similarly, the moral hazard component on the right-hand side of (IC^L) is increasing in T^H. Thus, it is optimal to have under-experimentation in T^H ($T_{SB}^H < T_{FB}^H$).

Section III. Sufficient conditions for IC constraints to be binding/slack.

Case 1: For any β_0^H there exist $\bar{\lambda}(\beta_0^H) > 0$, $\underline{A}(\beta_0^H) > 0$, and $\overline{\beta}_0^L(\beta_0^H) > 0$ such that all the (MH_t^θ) are binding for each type, but neither IC is binding if $\lambda < \bar{\lambda}$, $\gamma > \underline{A}\Delta c q_F$, and $\beta_0^L < \overline{\beta}_0^L$.

Case 2: For any β_0^H there exist $\underline{\lambda}(\beta_0^H) < 1$ and $\overline{A}(\beta_0^H) > 0$ such that both *IC* bind if $\lambda > \underline{\lambda}$, $\gamma < \overline{A}\Delta c q_F$, and $\beta_0^L < \frac{1}{2}$.

Intermediate case 3: For any λ there exist $\underline{\beta_0^H}(\lambda) < 1$, $\overline{\beta}_0^L(\lambda) > 0$, and $\underline{\underline{A}}(\lambda) > 0$ such that (IC^H) and all the (MH_t^L) bind if $\beta_0^H > \underline{\beta_0^H}$, $\beta_0^L < \overline{\beta}_0^L$ and $\gamma > \underline{\underline{A}}\Delta c q_F$.

Intermediate case 4: For any β_0^L there exist $\beta_0^L < \underline{\beta_0^H}(\beta_0^L)$, $\tilde{A}(\beta_0^L) > 0$, and $\bar{\lambda}(\beta_0^L) > 0$ such that IC^L is binding and IC^H is not binding if $\beta_0^L < \beta_0^H < \underline{\beta_0^H}$, $\gamma > \tilde{A}\Delta c q_F$, and $\lambda < \bar{\lambda}$.

Proof: We prove the sufficient conditions in two steps. In *Step 1*, we prove that (i) (IC^H) is not binding if λ is small and γ is sufficiently higher than Δc and (ii) (IC^H) is binding: if either β_0^H is high enough or λ is high enough. In *Step 2*, we prove that (i) (IC^L) is slack if β_0^L is small and γ is high and (ii) (IC^L) is binding if either β_0^H close to β_0^L or λ *is high*, β_0^L *not too high, and* γ *is small*. Then, combining the corresponding sufficient conditions, we establish 4 cases.

Step 1, $(\boldsymbol{IC^H})$.

Step 1a. $(\boldsymbol{IC^H})$ ***is binding.*** We now prove that (IC^H) is binding if either β_0^H is high enough or λ is high enough. To characterize sufficient conditions for the (IC^H) to be binding, we establish the parameters under which the highest possible value of the LHS of the (IC^H) evaluated at $x_t^H = 0$ and $y_t^H = \frac{\gamma}{\lambda \beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F$ for $t \leq T^H$ is less than the lowest possible value of the RHS (what he can claim by misrepresenting his type) evaluated at $x_t^L = 0$ and $y_t^L = \frac{\gamma}{\lambda \beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F$ for $t \leq T^L$.

Therefore, recalling that we proved in Lemma 1 that the high type works in every period off-the-equilibrium, i.e., $e_t^{H,L} = 1$ for $t \leq T^L$, the (IC^H) is satisfied if and only if:

$$-\gamma(1-\beta_0^H)T^H + \beta_0^H \sum_{t=1}^{T^H}(1-\lambda)^{t-1}(\lambda y_t^H - \gamma)$$

$$\geq -\gamma(1-\beta_0^H)T^L + \beta_0^H \sum_{t=1}^{T^L}(1-\lambda)^{t-1}(\lambda y_t^L - \gamma) + P_{T^L}^H \Delta c_{T^L+1} q_F.$$

We next simplify (IC^H) using $P_{T^L}^H \Delta c_{T^L+1} q_F = (1-\lambda)^{T^L}\left(\frac{\beta_0^H - \beta_0^L}{P_{T^L}^L}\right)\Delta c q_F$, $\sum_{t=1}^{T^\theta}(1-\lambda)^{t-1} = \frac{1-(1-\lambda)^{T^\theta}}{\lambda}$, and $\lambda y_t^\theta - \gamma = \frac{\gamma}{\beta_{T^\theta}^\theta} + \frac{\lambda(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F - \gamma = \frac{\lambda(1-\beta_0^\theta)}{P_{T^\theta}^\theta}\Delta c q_F + \gamma\left(\frac{1-\beta_0^\theta}{\beta_0^\theta(1-\lambda)^{T^\theta - 1}}\right)$ to obtain:

$$(IC^H)\quad \frac{\beta_0^H \gamma}{\lambda}\left(\frac{(1-\beta_0^H)\left(1-(1-\lambda)^{T^H}\right)}{\beta_0^H(1-\lambda)^{T^H-1}} - \frac{(1-\beta_0^L)\left(1-(1-\lambda)^{T^L}\right)}{\beta_0^L(1-\lambda)^{T^L-1}}\right) - \gamma(1-\beta_0^H)(T^H - T^L)$$

$$\geq \left(\frac{\beta_0^H(1-\beta_0^L)+(1-\lambda)^{T^L}\beta_0^L(1-\beta_0^H)}{P_{T^L}^L} - \frac{\beta_0^H(1-\beta_0^H)\left(1-(1-\lambda)^{T^H}\right)}{P_{T^H}^H}\right)\Delta c q_F.$$

Finally, given that

$$\frac{(1-\beta_0^H)\left(1-(1-\lambda)^{T^H}\right)}{\beta_0^H(1-\lambda)^{T^H-1}} - \frac{(1-\beta_0^L)\left(1-(1-\lambda)^{T^L}\right)}{\beta_0^L(1-\lambda)^{T^L-1}} = \frac{(1-\beta_0^H)\beta_0^L(1-\lambda)^{T^L} - (1-\beta_0^L)\beta_0^H(1-\lambda)^{T^H} + (\beta_0^H - \beta_0^L)(1-\lambda)^{T^L+T^H}}{\beta_0^H \beta_0^L (1-\lambda)^{T^H+T^L-1}},$$

and $\frac{\beta_0^H(1-\beta_0^L)+(1-\lambda)^{T^L}\beta_0^L(1-\beta_0^H)}{P_{T^L}^L} - \frac{\beta_0^H(1-\beta_0^H)\left(1-(1-\lambda)^{T^H}\right)}{P_{T^H}^H} =$

$$\frac{\beta_0^H(1-\beta_0^L)(1-\lambda)^{T^H} + \beta_0^L(1-\beta_0^H)(1-2\beta_0^H)(1-\lambda)^{T^L} + 2\beta_0^L\beta_0^H(1-\beta_0^H)(1-\lambda)^{T^H+T^L}}{P_{T^L}^L P_{T^H}^H},$$

(IC^H) above simplifies to

$$(IC^H) \quad \left(\frac{(1-\beta_0^H)\beta_0^L(1-\lambda)^{T^L}-(1-\beta_0^L)\beta_0^H(1-\lambda)^{T^H}+(\beta_0^H-\beta_0^L)(1-\lambda)^{T^L+T^H}}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}-(1-\beta_0^H)(T^H-T^L)\right)\gamma$$

$$\geq \left(\frac{\beta_0^H(1-\beta_0^L)(1-\lambda)^{T^H}+\beta_0^L(1-\beta_0^H)(1-2\beta_0^H)(1-\lambda)^{T^L}+2\beta_0^L\beta_0^H(1-\beta_0^H)(1-\lambda)^{T^H+T^L}}{P_{T^L}^L P_{T^H}^H}\right)\Delta c q_F.$$

β_0^H *is high enough.* First, consider the LHS of (IC^H) if $\beta_0^H \to 1$. The coefficient in front γ is:

$$\frac{(1-\beta_0^H)\beta_0^L(1-\lambda)^{T^L}-(1-\beta_0^L)\beta_0^H(1-\lambda)^{T^H}+(\beta_0^H-\beta_0^L)(1-\lambda)^{T^L+T^H}}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}-(1-\beta_0^H)(T^H-T^L)=$$

$$\frac{(1-\beta_0^H)\beta_0^L(1-\lambda)^{T^L}}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}-\frac{(1-\beta_0^L)\beta_0^H(1-\lambda)^{T^H}}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}+\frac{(\beta_0^H-\beta_0^L)(1-\lambda)^{T^L+T^H}}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}-(1-\beta_0^H)(T^H-T^L).$$

Replacing β_0^H with the limit value of one in the expression above we obtain:

$$\frac{(1-\beta_0^L)\left[(1-\lambda)^{T^L+T^H}-(1-\lambda)^{T^H}\right]}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}.$$

Consider expression $\frac{(1-\beta_0^L)\left[(1-\lambda)^{T^L+T^H}-(1-\lambda)^{T^H}\right]}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}$ for any $\beta_0^L > 0$. If $T^L \geq 1$, then $(1-\lambda)^{T^L+T^H} < (1-\lambda)^{T^H}$ and, as a result, $\frac{(1-\beta_0^L)\left[(1-\lambda)^{T^L+T^H}-(1-\lambda)^{T^H}\right]}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}} < 0$ for any $\beta_0^L > 0$. If T^L becomes zero for small values of $\beta_0^L > 0$, then the numerator becomes zero while the denominator remains strictly positive. Therefore, for any $\beta_0^L > 0$, $\frac{(1-\beta_0^L)\left[(1-\lambda)^{T^L+T^H}-(1-\lambda)^{T^H}\right]}{\beta_0^L\lambda(1-\lambda)^{T^H+T^L-1}}$ is either negative or zero.

Thus, the LHS of (IC^H) if $\beta_0^H \to 1$ is either negative or zero for any $\beta_0^L > 0$.

Second, note that the RHS of (IC^H) is strictly positive if $\beta_0^H \to 1$.

As a result, (IC^H) is binding for high enough β_0^H. We define a high enough value of β_0^H, called $\underline{\beta}_0^H$, as:

$$\underline{\beta}_0^H: \quad \frac{\left(1-\underline{\beta}_0^H\right)\underline{\beta}_0^H(1-\lambda)^{T^L}-\left(1-\underline{\beta}_0^H\right)\underline{\beta}_0^H(1-\lambda)^{T^H}+\left(\underline{\beta}_0^H-\beta_0^L\right)(1-\lambda)^{T^L+T^H}}{\underline{\beta}_0^H\lambda(1-\lambda)^{T^H+T^L-1}}$$

$$=\left(1-\beta_{0,HLa}^H\right)(T^H-T^L).$$

Thus, (IC^H) is binding if $\beta_0^H > \underline{\beta}_0^H$.

λ *is high enough*. We prove that (IC^H) is binding if λ is sufficiently high. First, the coefficient in front of γ on the LHS in (IC^H) becomes negative if $\lambda \to 1$. Since $T^H \to 1$ and $T^L \to 1$:[41]

$$\lim_{\lambda\to 1}\left[\frac{(1-\beta_0^H)\beta_0^L(1-\lambda)-(1-\beta_0^L)\beta_0^H(1-\lambda)+(\beta_0^H-\beta_0^L)(1-\lambda)^2}{\beta_0^L\lambda(1-\lambda)}\right]$$

$$=\lim_{\lambda\to 1}\left[\frac{(1-\beta_0^H)\beta_0^L-(1-\beta_0^L)\beta_0^H+(\beta_0^H-\beta_0^L)(1-\lambda)}{\beta_0^L\lambda}\right]=\frac{(1-\beta_0^H)\beta_0^L-(1-\beta_0^L)\beta_0^H}{\beta_0^L}$$

$$=\frac{\beta_0^L-\beta_0^H\beta_0^L-(\beta_0^H-\beta_0^L\beta_0^H)}{\beta_0^L}=\frac{\beta_0^L-\beta_0^H}{\beta_0^L}<0.$$

We define a high enough value of λ, called λ_1, as:

$$\lambda_1: \qquad \frac{(1-\beta_0^H)\beta_0^L(1-\lambda_1)^{T^L}-(1-\beta_0^L)\beta_0^H(1-\lambda_1)^{T^H}+(\beta_0^H-\beta_0^L)(1-\lambda_1)^{T^L+T^H}}{\beta_0^L\lambda_1(1-\lambda_1)^{T^H+T^L-1}}=(1-\beta_0^H)(T^H-T^L).$$

Second, the coefficient in front of $\Delta c q_F$ on the RHS goes to zero if $\lambda \to 1$:

$$\lim_{\lambda\to 1}\frac{\beta_0^H(1-\beta_0^L)(1-\lambda)^{T^H}+\beta_0^L(1-\beta_0^H)(1-2\beta_0^H)(1-\lambda)^{T^L}+2\beta_0^L\beta_0^H(1-\beta_0^H)(1-\lambda)^{T^H+T^L}}{P_{T^L}^L P_{T^H}^H}$$

$$=\frac{\beta_0^H(1-\beta_0^L)(1-1)^1+\beta_0^L(1-\beta_0^H)(1-2\beta_0^H)(1-1)^1+2\beta_0^L\beta_0^H(1-\beta_0^H)(1-1)^{1+1}}{P_1^L P_1^H}=0.$$

Thus, the coefficient in front of $\Delta c q_F$ is positive if $\lambda > \lambda_2$.

Therefore, (IC^H) is binding if $\lambda > \lambda_3 = \max\{\lambda_1, \lambda_2\}$.

Step 1b. $(\boldsymbol{IC^H})$ ***is not binding.*** We now prove that (IC^H) is not binding if λ *is low enough and* γ *is high enough*. To characterize sufficient conditions for the (IC^H) not to be binding, we establish the parameters under which the lowest value of the LHS of (IC^H) evaluated at $x_t^H = 0$ and $y_t^H = \frac{\gamma}{\lambda\beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F$ for $t \le T^H$ is greater than the highest value of the RHS .

The (IC^H) can be rewritten as:

$$-\gamma(1-\beta_0^H)T^H + \beta_0^H\sum_{t=1}^{T^H}(1-\lambda)^{t-1}\left(\lambda\left[\frac{\gamma}{\lambda\beta_{T^H}^H}+\frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F\right]-\gamma\right)$$

$$\ge -\gamma(1-\beta_0^H)T^L + \beta_0^H\sum_{t=1}^{T^L}(1-\lambda)^{t-1}(\lambda y_t^L-\gamma) + P_{T^L}^H\Delta c_{T^L+1} q_F + P_{T^L}^H x_{T^L}^L.$$

Consider small values of λ and let us take a limit $\lambda \to 0$. As $\lambda \to 0$, there will be small enough values for which $T^L = t^{L,H} = 0$. We define a small enough value of λ, called $\bar{\lambda}$, suh that $T^L = t^{L,H} = 0$ if $\lambda < \bar{\lambda}$. As a result, for those small values of $\lambda < \bar{\lambda}$, $-\gamma(1-\beta_0^H)T^L + \beta_0^H\sum_{t=1}^{T^L}(1-\lambda)^{t-1}(\lambda y_t^L-\gamma) = 0$ and $x_{T^L}^L = (c_{T^H}^H - c_1^L)q_F > 0$ if $c_{T^H}^H > c_1^L$ and $x_{T^L}^L \le 0$ if

[41] Because T^H and T^L are discrete, they converge to 1 faster than λ.

$c^H_{T^H} \leq c^L_1$.[42] And $\Delta c_1 = c^L_1 - c^H_1$. Since off-the-equilibrium rent of the high type is smaller in the latter case ($c^H_{T^H} \leq c^L_1$), it is sufficient to characterize sufficient conditions for (IC^H) to be satisfied in the former ($c^H_{T^H} > c^L_1$) case only. Consider the former case ($c^H_{T^H} > c^L_1$). Given that $P^H_{T^L} = 1$, the (IC^H) constraint simplifies to:

$$-\gamma(1-\beta_0^H)T^H + \beta_0^H \sum_{t=1}^{T^H}(1-\lambda)^{t-1}\left(\lambda\left[\frac{\gamma}{\lambda\beta^H_{T^H}} + \frac{(1-\beta_0^H)}{P^H_{T^H}}\Delta c q_F\right] - \gamma\right) \geq (c^H_{T^H} - c^H_1)q_F,$$

$$\gamma\left(\beta_0^H\left(\frac{1}{\beta^H_{T^H}} - 1\right)\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right) - (1-\beta_0^H)T^H\right) + \beta_0^H\frac{(1-\beta_0^H)}{P^H_{T^H}}\Delta c q_F\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right) \geq (c^H_{T^H} - c^H_1)q_F.$$

Since $\Delta c > c^H_{T^H} - c^H_1$, it is sufficient to prove the following condition (we replace $c^H_{T^H} - c^H_1$ with Δc on the RHS of the (IC^H) above):

$$\gamma\left(\beta_0^H\left(\frac{1}{\beta^H_{T^H}} - 1\right)\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right) - (1-\beta_0^H)T^H\right) + \beta_0^H\frac{(1-\beta_0^H)}{P^H_{T^H}}\Delta c q_F\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right) \geq \Delta c q_F,$$

$$\gamma\left(\beta_0^H\left(\frac{1}{\beta^H_{T^H}} - 1\right)\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right) - (1-\beta_0^H)T^H\right) \geq \left(1 - \beta_0^H\frac{(1-\beta_0^H)}{P^H_{T^H}}\Delta c q_F\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right)\right)\Delta c q_F,$$

Since the coefficient in front of γ is positive (the moral hazard rent in a model of experimentation without production is positive), the (IC^H) is satisfied if

$$\gamma \geq \frac{\left(1-\beta_0^H\frac{(1-\beta_0^H)}{P^H_{T^H}}\Delta c q_F\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right)\right)}{\left(\beta_0^H\left(\frac{1}{\beta^H_{T^H}}-1\right)\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right)-(1-\beta_0^H)T^H\right)}\,\Delta c q_F,$$

Therefore, (IC^H) is not binding if $\lambda < \bar{\lambda}$ and $\gamma > \tilde{A}\Delta c q_F$, where

$$\tilde{A} \equiv \max_{T^H,T^L}\frac{\left(1-\beta_0^H\frac{(1-\beta_0^H)}{P^H_{T^H}}\Delta c q_F\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right)\right)}{\left(\beta_0^H\left(\frac{1}{\beta^H_{T^H}}-1\right)\left(\frac{1-(1-\lambda)^{T^H}}{\lambda}\right)-(1-\beta_0^H)T^H\right)}.$$

Step 2, $(\boldsymbol{IC^L})$.

Step 2a. $(\boldsymbol{IC^L})$ ***is binding.*** We now prove that (IC^L) is binding if either λ *is high,* β_0^L *not too high, and* γ *is small* or β_0^H *close to* β_0^L. To characterize sufficient conditions for the (IC^L) to be binding, we establish the parameters under which the highest possible value of the LHS of (IC^L)

[42] The term c^L_t is the expected cost for the L type after $t-1$ failure (so c^L_1 is the expected cost after no experimentation).

evaluated at $x_t^L = 0$ and $y_t^L = \frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F$ for $t \le T^L$ is less than the lowest possible value of the RHS evaluated at $x_t^H = 0$ and $y_t^H = \frac{\gamma}{\lambda\beta_{T^H}^H} + \frac{(1-\beta_0^H)}{P_{T^H}^H}\Delta c q_F$ for $t \le T^H$.

The (IC^L) can be rewritten:

$$-\gamma(1-\beta_0^L)T^L + \beta_0^L \sum_{t=1}^{T^L}(1-\lambda)^{t-1}(\lambda y_t^L - \gamma)$$
$$\ge -\gamma(1-\beta_0^L)t^{L,H} + \beta_0^L \sum_{t=1}^{t^{L,H}}(1-\lambda)^{t-1}\,(\lambda y_t^H - \gamma) + P_{t^{L,H}}^L\left(c_{T^H+1}^H - c_{t^{L,H}+1}^L\right)q_F,$$

which can be simplified to:

$$(IC^L)\ \left(\frac{(1-\beta_0^L)(1-(1-\lambda)^{T^L})}{\lambda(1-\lambda)^{T^L-1}} - \frac{(1-\beta_0^H)\beta_0^L\left(1-(1-\lambda)^{t^{L,H}}\right)}{\beta_0^H\lambda(1-\lambda)^{T^H-1}} - (1-\beta_0^L)(T^L - t^{L,H})\right)\gamma$$
$$\ge \left(\frac{\beta_0^L(1-\beta_0^H)(1-\lambda)^{T^L} + \beta_0^H(1-\beta_0^L)(1-2\beta_0^L)(1-\lambda)^{T^H} + 2\beta_0^L\beta_0^H(1-\beta_0^L)(1-\lambda)^{T^H+T^L}}{P_{T^L}^L P_{T^H}^H}\right)\Delta c q_F.$$

λ is high, β_0^L not too high, and γ is small. First, if λ is high enough, then the coefficient in front of γ is positive as well. First, note that $T^H \to 1$, $T^L \to 1$, and $t^{L,H} \to 1$.[43] As a result, this coefficient becomes:

$$\lim_{\lambda\to 1}\left[\frac{(1-\beta_0^L)}{(1-\lambda)^0} - \frac{(1-\beta_0^H)\beta_0^L}{\beta_0^H(1-\lambda)^0}\right] = \left(\frac{\beta_0^H(1-\beta_0^L)-(1-\beta_0^H)\beta_0^L}{\beta_0^H}\right)\frac{1}{(1-\lambda)^0} = \left(\frac{\beta_0^H-\beta_0^L}{\beta_0^H}\right)\lim_{\lambda\to 1}\left[\frac{1}{(1-\lambda)^0}\right] > 0.$$

Therefore, there exists a high enough value of λ called λ_4 such that $\lambda_4 < 1$:

$$\lambda_4:\ \frac{(1-\beta_0^L)(1-(1-\lambda_4)^{T^L})}{\lambda_4(1-\lambda_4)^{T^L-1}} - \frac{(1-\beta_0^H)\beta_0^L\left(1-(1-\lambda_4)^{t^{L,H}}\right)}{\beta_0^H\lambda_4(1-\lambda_4)^{T^H-1}} - (1-\beta_0^L)(T^L - t^{L,H}) = 0.$$

Second, if β_0^L is not too high (sufficient $\beta_0^L < \frac{1}{2}$), then the coefficient in front of Δc is positive. Thus, (IC^L) is binding if $1 > \lambda > \lambda_4$, $\beta_0^L < \frac{1}{2}$, and $\gamma < \overline{A}\Delta c q_F$, where

$$\overline{A} \equiv \min_{T^H,T^L} \frac{\left(\frac{\beta_0^L(1-\beta_0^H)(1-\lambda)^{T^L} + \beta_0^H(1-\beta_0^L)(1-2\beta_0^L)(1-\lambda)^{T^H} + 2\beta_0^L\beta_0^H(1-\beta_0^L)(1-\lambda)^{T^H+T^L}}{P_{T^L}^L P_{T^H}^H}\right)}{\left(\frac{(1-\beta_0^L)(1-(1-\lambda)^{T^L})}{\lambda(1-\lambda)^{T^L-1}} - \frac{(1-\beta_0^H)\beta_0^L\left(1-(1-\lambda)^{t^{L,H}}\right)}{\beta_0^H\lambda(1-\lambda)^{T^H-1}} - (1-\beta_0^L)(T^L - t^{L,H})\right)}.$$

β_0^H close to β_0^L. If the coefficients in front of γ in (IC^L) is zero, the (IC^L) constraint is binding. We next prove that this is the case if β_0^H is not too high as the coefficient in front of γ in (IC^L) is going to zero. In particular, note that if $\beta_0^L \to \beta_0^H$, then $T^L \to t^{L,H} \to T^H \to T$, and, as a result,

[43] Because T^H and T^L are discrete, they converge to 1 faster than λ.

$$\lim_{\beta_0^L \to \beta_0^H} \left[\frac{(1-\beta_0^L)\left(1-(1-\lambda)^{T^L}\right)}{\lambda(1-\lambda)^{T^L-1}} - \frac{(1-\beta_0^H)\beta_0^L\left(1-(1-\lambda)^{t^{L,H}}\right)}{\beta_0^H \lambda(1-\lambda)^{T^H-1}} - (1-\beta_0^L)(T^L - t^{L,H}) \right]$$

$$= \lim_{\beta_0^L \to \beta_0^H} \left[\left(\frac{(1-(1-\lambda)^T)}{\lambda(1-\lambda)^{T-1}} \right) \left(\frac{\beta_0^H(1-\beta_0^H)-(1-\beta_0^H)\beta_0^H}{\beta_0^H} \right) \right] = 0.$$

Therefore, there exists a small enough value of β_0^H, called $\underline{\beta}_0^H$, such that $\underline{\beta}_0^H > \beta_0^L$:

$$\underline{\beta}_0^H: \ \frac{(1-\beta_0^L)(1-(1-\lambda)^{T^L})}{\lambda(1-\lambda)^{T^L-1}} - \frac{\left(1-\underline{\beta}_0^H\right)\beta_0^L\left(1-(1-\lambda)^{t^{L,H}}\right)}{\underline{\beta}_0^H \lambda(1-\lambda)^{T^H-1}} = (1-\beta_0^L)(T^L - t^{L,H}).$$

Thus, (IC^L) is binding if $\beta_0^L < \beta_0^H < \underline{\beta}_0^H$.

Step 2b. $(\boldsymbol{IC^L})$ ***is not binding.*** We now prove that (IC^L) is not binding if β_0^L *is small and* γ *is high*. To characterize sufficient conditions for the (IC^L) not to be binding, we establish the parameters under which the lowest value of the LHS of (IC^L) evaluated at $x_t^L = 0$ and $y_t^L = \frac{\gamma}{\lambda \beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F$ for $t \le T^L$ is greater than the highest value of the RHS. To characterize this highest value, consider small values of β_0^L and take the limit $\beta_0^L \to 0$. As $\beta_0^L \to 0$, there will be small enough values for which $T^L = t^{L,H} = 0$. We define a small enough value of β_0^L, called $\overline{\beta}_0^L$, such that $T^L = t^{L,H} = 0$ if $\beta_0^L < \overline{\beta}_0^L$. As a result, for those small values of $\beta_0^L < \overline{\beta}_0^L$, off-the-equilibrium RHS for the low type is $U^L\left(\varpi^H, \vec{a}^L(\varpi^H)\right) = \left(c_{T^H}^H - c_1^L\right)q_F$. it is sufficient to consider the case $(c_{T^H}^H > c_1^L)$. Using similar steps as in step 1b, the (IC^L) constraint is satisfied if the following condition holds:

$$-\gamma(1-\beta_0^L)T^L + \beta_0^L \sum_{t=1}^{T^L}(1-\lambda)^{t-1}\left(\lambda\left[\frac{\gamma}{\lambda\beta_{T^L}^L} + \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F\right] - \gamma\right) \ge \left(c_{T^H}^H - c_1^L\right)q_F.$$

$$\gamma\left(\beta_0^L\left(\frac{1}{\beta_{T^L}^L} - 1\right)\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right) - (1-\beta_0^L)T^L\right) + \beta_0^L \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right) \ge \left(c_{T^H}^H - c_1^H\right)q_F.$$

Since $\Delta c > c_{T^H}^H - c_1^H$, it is sufficient to prove the following condition (we replace $c_{T^H}^H - c_1^H$ with Δc on the RHS of the constraint above):

$$\gamma\left(\beta_0^L\left(\frac{1}{\beta_{T^L}^L} - 1\right)\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right) - (1-\beta_0^L)T^L\right) + \beta_0^L \frac{(1-\beta_0^L)}{P_{T^L}^L}\Delta c q_F \left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right) \ge \Delta c q_F,$$

$$\gamma\left(\beta_0^L\left(\frac{1}{\beta_{T^L}^L} - 1\right)\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right) - (1-\beta_0^L)T^L\right) \ge \left(1 - \beta_0^L \frac{(1-\beta_0^L)}{P_{T^L}^L}\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right)\right)\Delta c q_F,$$

Since the coefficient in front of γ is positive (the moral hazard rent in a model of experimentation without production is positive), the (IC^L) is satisfied if

$$\gamma \geq \frac{\left(1-\beta_0^L\frac{\left(1-\beta_0^L\right)}{P_{T^L}^L}\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right)\right)}{\left(\beta_0^L\left(\frac{1}{\beta_{T^L}^L}-1\right)\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right)-\left(1-\beta_0^L\right)T^L\right)}\,\Delta c q_F,$$

Therefore, (IC^L) is not binding if $\beta_0^L < \overline{\beta}_0^L$ and $\gamma > \underline{\underline{A}}\Delta c q_F$, where

$$\underline{\underline{A}} \equiv \max_{T^L} \frac{\left(1-\beta_0^L\frac{\left(1-\beta_0^L\right)}{P_{T^L}^L}\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right)\right)}{\left(\beta_0^L\left(\frac{1}{\beta_{T^L}^L}-1\right)\left(\frac{1-(1-\lambda)^{T^L}}{\lambda}\right)-\left(1-\beta_0^L\right)T^L\right)}.$$

We now state the sufficient conditions for each of the 4 cases formally.

***Case 1**:* We established that there exist $\overline{\lambda}$ and $\tilde{A}$ such that (IC^H) is automatically satisfied if $\lambda < \overline{\lambda}$ and $\gamma > \tilde{A}\Delta c q_F$. In addition, we established that there exist $\overline{\beta}_0^L$ and $\underline{\underline{A}}$ such that (IC^L) is automatically satisfied if $\beta_0^L < \overline{\beta}_0^L$ and $\gamma > \underline{\underline{A}}\Delta c q_F$. Combining the two results, we obtain:

Both (IC^H) and (IC^L) are not binding if $\lambda < \overline{\lambda}$, $\beta_0^L < \overline{\beta}_0^L$, and $\gamma > \underline{A}\Delta c q_F$,

where $\underline{A} = \max\{\tilde{A}, \underline{\underline{A}}\}$.

Case 2: We established that there exist λ_3 such that (IC^H) is binding if $\lambda > \lambda_3$. In addition, we established that there exist λ_4 and A_3 such that (IC^L) is binding if $1 > \lambda > \lambda_4$, $\beta_0^L < \frac{1}{2}$, and $\gamma < \overline{A}\Delta c q_F$. Combining the two results, we obtain:

Both IC constraints are binding if $\lambda > \underline{\lambda}$, $\gamma < \overline{A}\Delta c q_F$, and $\beta_0^L < \frac{1}{2}$,

where $\underline{\lambda} = \max\{\lambda_3, \lambda_4\}$.

Intermediate Case 3*:* We established that there exist $\underline{\beta}_0^H$ such that (IC^H) is binding if $\beta_0^H > \underline{\beta}_0^H$. In addition, we established that there exist $\overline{\beta}_0^L$ and A_4 such that (IC^L) is automatically satisfied if $\beta_0^L < \overline{\beta}_0^L$ and $\gamma > \underline{\underline{A}}\Delta c q_F$. Combining the two results, we obtain:

(IC^H) is binding and (IC^L) is not binding if $\beta_0^H > \underline{\beta}_0^H$, $\beta_0^L < \overline{\beta}_0^L$, and $\gamma > \underline{\underline{A}}\Delta c q_F$.

Intermediate Case 4*:* We established that there exist $\overline{\lambda}$ and $\tilde{A}$ such that (IC^H) is not binding if $\lambda < \overline{\lambda}$ and $\gamma > \tilde{A}\Delta c q_F$. In addition, we established that there exist $\overline{\beta}_0^L$ such that (IC^L) is binding if $\beta_0^L < \beta_0^H < \underline{\beta}_0^H$. Combining the two results, we obtain:

(IC^H) is not binding and (IC^L) is binding if $\beta_0^L < \beta_0^H < \underline{\beta}_0^H$, $\gamma > \tilde{A}\Delta c q_F$, and $\lambda < \overline{\lambda}$.

Q.E.D.

References

Anand, S., Wittes, J., and Yusuf, S. (2011). What information should a sponsor of a randomized trial receive during its conduct? *Clin. Trials* 8 (6), 716-719.

Benabou R. and Tirole J. (2003). Intrinsic and Extrinsic Motivation. *Review of Economic Studies* 70, 489-520.

Bergemann D. and Hege U. (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking and Finance* 22 (6-8), 703-735.

Bergemann D. and Hege U. (2005). The Financing of innovation: learning and stopping. *Rand Journal of Economics* 36 (4), 719-752.

Bergemann, D., and Välimäki, J. (2008). Bandit problems. In: Durlauf, S., Blume, L. (Eds.), The New Palgrave Dictionary of Economics, 2nd edition. Palgrave Macmillan Ltd, Basingstoke, New York.

Bhaskar V. and Mailath G. J. (2019). The curse of long horizons. *Journal of Mathematical Economics* 82, 74-89.

Bhaskar, V. and Roketskiy, N. (2023). The ratchet effect: A learning perspective. *Journal of Economic Theory* 214, 105762.

Bolton P. and Harris C. (1999). Strategic Experimentation. *Econometrica* 67(2), 349-374.

Castro-Pires, H., Chade H., and Swinkels J. (2024). Disentangling Moral Hazard and Adverse Selection. American Economic Review 114 (1), 1-37.

Castro-Pires, H. and Moreira H. (2021) "Limited Liability and Non-responsiveness in Agency Models" *Games and Economic Behavior,* 128, 73-103.

Chade H. and Kovrijnykh N. (2016). Delegated information acquisition with moral hazard. Journal of Economic Theory 162, 55-92.

Chakraborty I., Khalil F. and Lawarree J. (2021). Competitive procurement with ex post moral hazard. *Rand Journal of Economics* 52(1), 179-206.

Crémer J. and Khalil F. (1992). Gathering Information before Signing a Contract. *American Economic Review* 82(3), 566-578.

Crémer J. and Mclean R. P. (1985). Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent. *Econometrica* 53, 345-361.

Crémer J., Khalil F. and Rochet J.-C. (1992). Contracts and productive information gathering. *Games and Economic Behavior* 25(2), 174-193.

Downs J. (2021). Information gathering by overconfident agents. *Journal of Economics and Management Strategy* 30, 554-568.

Ederer F. and Manso G. (2013). Is pay for performance detrimental to innovation? *Management Science* 59(7), 1496-1513.

Edmans A., Gabaix X. and Jenter D. (2017) Executive Compensation: A Survey of Theory and Evidence, in Hermalin B. and Weisbach M. *The Handbook of the Economics of Corporate Governance,* Vol.1, 383-539.

Escobar, J. F. and Zhang Q. (2021). Delegating learning. *Theoretical Economics* 16(2), 571-603.

Foarta D. and Sugaya T. (2021). The management of talent: Optimal contracting for selection and incentives. *Rand Journal of Economics* 52(1), 49-77.

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

Gomes R., Gottlieb D. and Maestri L. (2016). Experimentation and project selection: screening and learning. *Games and Economic Behavior* 96, 145-169.

Gottlieb, D. and Moreira H. (2022). Simple contracts with adverse selection and moral hazard. *Theoretical Economics* 17(3): 1357-1401.

Guo, Y. (2016). Dynamic Delegation of Experimentation. *American Economic Review* 106(8),1969-2008.

Häfner S. and Taylor C. (2022). On young Turks and yes men: optimal contracting for advice. *Rand Journal of Economics* 53(1): 63-94.

Halac M., Kartik N. and Liu Q. (2016). Optimal contracts for experimentation. *Review of Economic Studies* 83(3), 1040-1091.

Hoppe, E. I., and P. W. Schmitz. (2013). Public-private partnerships versus traditional procurement: Innovation incentives and information gathering. *The RAND Journal of Economics* 44, 56-74.

Hoppe, E.I., and Schmitz, P.W. (2021). How (not) to foster innovations in public infrastructure projects. *Scandinavian Journal of Economics* 123, 238-266.

Horner J. and Samuelson L. (2013). Incentives for experimenting agents. *RAND Journal of Economics* 44(4), 632-663.

Iossa, E., and D. Martimort. (2012). Risk allocation and the costs and benefits of public-private partnerships. *The RAND Journal of Economics* 43, 442-474.

Iossa, E and D. Martimort. (2015). Pessimistic information gathering. *Games and Economic Behavior* 91, 75-96.

Khalil, F., J. Lawarree, and A. Rodivilov. (2020). Learning from failures: Optimal contracts for experimentation and production. *Journal of Economic Theory* 190, 105107.

Khalil F., D. Kim, and D. Shin. (2006). Optimal task design: to integrate or separate planning and implementation? *Journal of Economics and Management Strategy* 14, 269-291.

Khrämer D. and Strausz R. (2011). Optimal Procurement Contracts with Pre-Project Planning. *Review of Economic Studies* 78, 1015-1041.

Laffont, J.-J., and Martimort, D. (2002). Theory of Incentives. Princeton University Press.

Lerner J. and Wulf J. (2007). Innovation and Incentives: Evidence from Corporate R&D. *Review of Economics and Statistics* 89, 634-644.

Lewis T. and Sappington D. (1997). Information Management in Incentive Problems. *Journal of Political Economy* 105(4), 796-821.

Li, S., Sun, H., Yan, J., and Yu, J. (2015). Bundling decisions in procurement auctions with sequential tasks. *Journal of Public Economics* 128, 96-106.

Lindblad, A.S., Manukyan, Z., and Purohit-Sheth, T. (2014). Central site monitoring: results from a test of accuracy in identifying trials and sites failing food and drug administration inspection. *Clin. Trials* 11 (2), 205-217.

Manso G. (2011). Motivating innovation. *Journal of Finance* 66, 1823-1860.

Martimort D., Poudou J.-C. and Thomas L. (2023). Optimal contracts with adverse selection and moral hazard: Are incentives high- or low-powered? *Mimeo*.

Moroni, S. (2022). Experimentation in Organizations. *Theoretical Economics* 17(3), 1403-1450.

Ollier S., and Thomas L. (2013). Ex post participation constraint in a principal-agent model with adverse selection and moral hazard. *Journal of Economic Theory* 148, 2383-2403.

Pogue, J.M., Devereaux, P.J., Thorlund, K., and Yusuf, S. (2013). Central statistical monitoring: detecting fraud in clinical trials. *Clin. Trials* 10 (2), 225-235.

Rodivilov, A. (2022). Monitoring innovation. *Games and Economic Behavior* 135, 297-326.

Rodivilov, A. (2021). A note on robust procurement contracts. *Economics Letters* 201.

Sadler, E. (2021). Dead Ends. *Journal of Economic Theory* 191, 105167.