Execution vs. Training under Endogenous Deadlines*

Yonggyun Kim† Curtis R. Taylor‡

September 12, 2025

Abstract

We study a principal-agent problem where the principal dynamically chooses between two methods for solving a problem: a direct execution route and a multi-stage training route with an observable milestone. To mitigate moral hazard, the principal commits to an endogenously determined deadline. The optimal contract is shaped by the interplay of three forces: the milestone effect from the training route's monitoring advantage, the deadline effect that favors the simpler execution route as time runs short, and the relative efficiency of each path.

JEL Classification: J41, L14, D86

Keywords: Dynamic Agency, Multiple Approaches, Poisson Arrival, Basic and Applied Research

*We thank Attila Ambrus, Arjada Bardhi, Christoph Carnehl, Doruk Cetemen, Yeon-Koo Che, Yingni Guo, Jinwoo Kim, Fei Li, Pino Lopomo, Francisco Poggi, Ludvig Sinander, Huseyin Yildirim, anonymous referees and seminar participants at Duke, Seoul National University, and the Southern Economic Association conference for comments and suggestions. All remaining errors are our own. This paper supersedes the previous working paper by Kim (2021) titled "The Direct vs. the Sequential Approach in Project Management."

†Department of Economics, Florida State University, ykim22@fsu.edu

‡Department of Economics, Duke University, crtaylor@duke.edu

1 Introduction

In managing complex projects, a central challenge is choosing the best methodology. A manager must often decide between pursuing a direct, single-stage approach using existing capabilities and adopting a more complex, multi-stage strategy that first requires developing a specialized skill. This tradeoff recurs in many settings. A software company might rely on its existing codebase to build a product, or it could first train its team on a new programming language tailored to the project's requirements. Similarly, a graduate student writing a dissertation might draw only on the training from their first- and second-year courses, or they might first build new competencies through additional coursework outside their department.

Perhaps most consequentially, this tradeoff was at the heart of the race to develop COVID-19 vaccines. Some ventures—Oxford-AstraZeneca, Johnson & Johnson, and Novavax—pursued established approaches such as viral-vector and protein-subunit methods. By contrast, Pfizer-BioNTech and Moderna-NIAID took the novel route of developing an mRNA platform, a strategy that required mastering the new technology before an actual vaccine could be produced.

To study this type of trade-off, we develop a model where a principal hires an agent to solve a problem. At any moment, the principal instructs the agent to pursue one of two methodologies. She can recommend the **execution route**, a direct path where the agent applies a basic skill to solve the problem, with success modeled as a Poisson arrival. Alternatively, she can recommend the **training route**, an indirect, two-stage path. On this route, the agent must first acquire an advanced skill before using it to solve the problem; both the initial skill acquisition and the final solution are also governed by Poisson processes.

A crucial feature of the training route is that the acquisition of the advanced skill is a *publicly observable* milestone, providing a clear signal of progress. This assumption is grounded in practice, where such preparatory stages often conclude with verifiable outcomes, such as obtaining a professional certification, completing a required training program, or a successful prototype demonstration.

The choice of methodology is complicated by a classic moral hazard problem. The solution or skill acquisition only occurs if the agent exerts effort, which is his private information. If the agent shirks, he enjoys a private benefit, and the principal's ability to counteract this is limited by the agent's limited liability, which restricts her to offering non-negative rewards. These frictions make deadlines essential for providing incentives; without one, the agent could shirk indefinitely. Therefore, the principal must impose a deadline, which is *endogenously* determined as part of the optimal contract.

The principal's optimal strategy is governed by the interplay of three competing economic forces. The first is the **milestone effect**, where the observable progress in the training route provides a monitoring advantage that allows for deadline extensions. This is weighed against the **deadline effect**, which makes the single-stage execution route more appealing as time runs short, since a single breakthrough is more probable than the two breakthroughs required for the training route. The final force is **efficiency**, defined by the expected time required to solve the problem. To isolate the core tensions, we first analyze a benchmark where both routes are equally efficient (i.e., have the same expected completion time) and later introduce the possibility that the execution route is more efficient (i.e., has a shorter expected completion time), creating a further trade-off against the training route's monitoring benefits.

Focusing on the case where the execution and training routes are equally efficient, we find that the optimal contract depends critically on the project return—the gross value of the solution relative to the operating cost (Theorem 2). When the project return is low, the optimal deadline is short, and the principal opts exclusively for the execution path. Conversely, when the project return is high, the benefit of monitoring the agent's progress becomes crucial, causing the milestone effect to dominate the deadline effect; thus, the principal always recommends the training route. For projects with intermediate returns, the optimal contract involves a switch in strategy: the principal initially recommends the training path but instructs the agent to switch to the execution path if the advanced skill is not acquired and the deadline becomes imminent.

We then introduce an efficiency loss from training, reflecting scenarios where its multi-stage nature may be slower than the more direct execution route. When this efficiency loss is small, our previous characterization remains robust; the optimal contract is still determined by low, intermediate, or high values of the project return (Proposition 2). However, when the efficiency loss is large, the trade-offs become sharper. The execution route's efficiency advantage makes it the preferred choice not only near the deadline but also at the beginning of a long project. Consequently, for most project returns, an execution-only contract is optimal. More interestingly, for very high-return projects, a novel two-switch contract emerges where all three economic forces shape the outcome: the principal begins with the execution route (due to efficiency), switches to the training route in the middle of the contract to leverage the milestone effect for monitoring, and finally reverts to the execution route as the deadline looms (Proposition 3).

Our findings align with observed differences in the structure of applied and basic scientific research. Applied research, such as developing a new drug or conducting clinical trials, is typically structured in stages with clear milestones. This corresponds to the training route in our model, which we find is optimal for high-return projects—a fitting description for applied research where the potential payoff is large relative to the operating cost. In contrast, basic research is often pursued "without thought of practical ends" (Bush, 1945), and thus has a lower immediate project return.[1] This prediction is consistent with funding mechanisms like the National Institutes of Health's R01 grant, which supports "a discrete, specified, circumscribed project" rather than a staged one.[2]

Related Literature There is a growing literature on contracting for multi-stage projects, e.g., Hu (2014); Green and Taylor (2016a); Wolf (2018); Moroni (2022). The most closely

[1] Bush argues that although broad and basic studies seem to be less important than applied ones, they are essential to combat diseases because progress in the treatment "will be made as the result of fundamental discoveries in subjects unrelated to those diseases, and perhaps entirely unexpected by the investigator." However, since this article does not consider externalities, we abstract from this possibility and focus on the principal's return from the completed project.

[2] https://grants.nih.gov/grants/funding/r01.htm

related study is Green and Taylor (2016a), who study a model in which multiple breakthroughs are needed to complete a project and in which an agent must be incentivized to exert unobservable effort. The training route considered here comprises the baseline model with the tangible breakthrough in the working paper version of their paper (Green and Taylor, 2016b). However, the option to complete the project directly, which is not considered in their setup, allows the principal to face a choice problem between the two routes. Moreover, this choice problem arises at every point in time. Therefore, the principal's problem is more complex from a dynamic perspective.

A related article is Carnehl and Schneider (2023), where they explore a two-armed bandit problem with one arm requiring a single breakthrough and the other needing multiple breakthroughs. The agent knows the arrival rates for the latter but must infer the feasibility of the former through experimentation—a key difference between their paper and this one. Unlike their focus on a single-agent decision problem with an exogenous deadline, this paper addresses principal-agent contracting with an endogenously determined deadline. Despite these differences, we share a common insight in that the chosen approaches may switch up to two times.

This investigation is also related to the literature on monitoring in dynamic contracts, e.g., Orlov (2022); Piskorski and Westerfield (2016); Dilmé and Garrett (2019); Marinovic and Szydlowski (2022); Varas et al. (2020); Marinovic and Szydlowski (2023); Chen et al. (2020); Wong (2023). In most of these papers, a monitoring process provides some information on the agent's current or past action. In this sense, the completion of the first breakthrough (advanced skill acquisition) in the training route can be considered as a monitoring device since it lets the principal know that the agent has worked. However, the skill acquisition gives more information than merely the agent's past actions. Before the skill acquisition, the success requires one relatively hard breakthrough or two easier breakthroughs. After skill acquisition, it requires only one relatively easy breakthrough. Thus, the skill acquisition is distinguished from standard monitoring processes since it also provides information about the subsequent likelihood of success.

This work is also relevant to the literature studying complementary innovations, e.g., Green and Scotchmer (1995); Gilbert and Katz (2011); Bryan and Lemus (2017); Poggi (2021). Two subprojects in the sequential approach can be considered as 'perfect' complements in the sense that completing one task in the training route does not create any value but completing both of them does. The most relevant paper in this line is Kim and Poggi (2025), which introduces an innovation race model with two R&D routes: one requiring a single breakthrough (direct development) and the other requiring two breakthroughs (research and development). However, to our knowledge, most studies in this literature focus on the problems involving competing firms or a single decision maker, whereas this article addresses an agency problem.

2 Model

A principal (she) hires an agent (he) to complete a project, specifically, to solve a problem. Problem-solving takes place in continuous time and can be performed in general over an infinite horizon: $t \in [0, \infty)$. When the problem is solved, the principal realizes a gross payoff $\Pi > 0$, and the game ends. The principal incurs an operating flow cost of $c > 0$ per unit of time until the problem is solved or the project is terminated. The principal is assumed to have an infinite amount of resources to fund the project, while the agent is protected by limited liability; that is, the principal can only transfer nonnegative rewards to the agent.[3] The principal and the agent are both risk-neutral and patient, i.e., they do not discount the future. Both players have outside options of zero.

At each point in time, the principal directs an agent to exert effort on one of two problem-solving methodologies, the *execution route* or the *training route*. The agent's hidden action is denoted by $\beta_t \in \{0, 1\}$, where $\beta_t = 1$ represents effort and $\beta_t = 0$ represents shirking. If

[3] See Remark 1 for further discussion of limited liability.

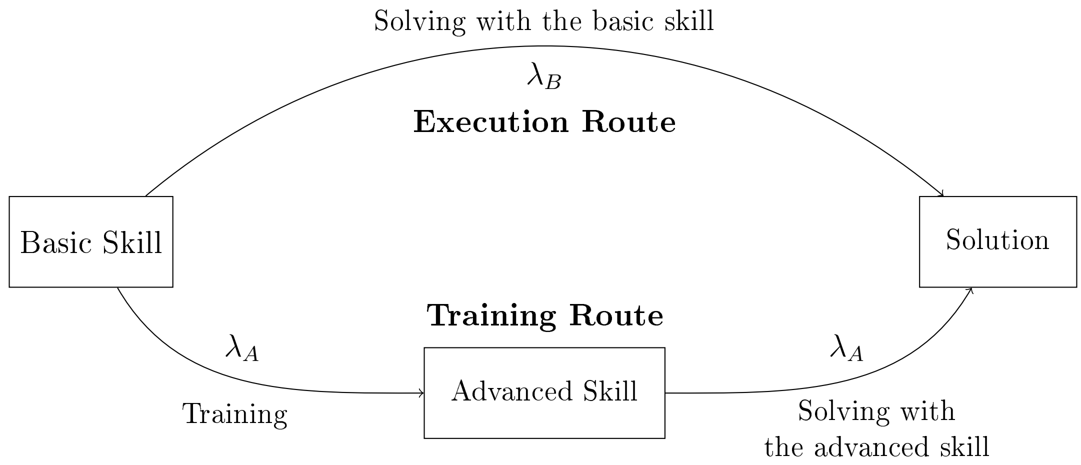

Figure 1: Problem-Solving Routes

the agent shirks, he receives a private flow benefit ϕ. We assume $c > \phi > 0$, i.e., the private benefit from shirking does not exceed the principal's operating cost.

The **execution** route is a direct path where the agent applies an existing *basic* skill to solve the problem. When this route is taken at time t, the problem is solved at a Poisson rate of $\lambda_B \beta_t$. The **training** route is a two-stage path where the agent must first acquire an *advanced* skill before using it to solve the problem. The acquisition of this skill is a publicly observable milestone. Under this route, both the initial skill acquisition and the subsequent problem-solving solution arrive at the same Poisson rate of $\lambda_A \beta_t$. These paths are illustrated in Figure 1.[4]

3 Benchmarks

Before deriving the optimal contract under asymmetric information, we explore two benchmark scenarios under which a social planner (it) chooses the best route for solving the problem without relying on an agent. We begin with the unconstrained first-best setting and then explore the presence of exogenous deadlines.

[4]The subscripts A and B stand for *advanced* and *basic* skills, respectively.

3.1 The First-best and Parametric Assumptions

Suppose the social planner is able to solve the problem on its own. If the planner employs the execution route, the expected amount of time until the problem is solved is $1/\lambda_B$. So its expected payoff from this route is $\Pi - c/\lambda_B$. On the other hand, the training route requires two breakthroughs, each with expected duration of $1/\lambda_A$. So the planner's expected payoff from this route is $\Pi - 2c/\lambda_A$. For the main analysis, we focus on the case where both routes are equally efficient: $2\lambda_B = \lambda_A$. In Section 5, we explore the case where the execution route is more efficient than the training route: $2\lambda_B > \lambda_A$.[5]

We also assume that the project is profitable enough to be undertaken under the execution route:

$$\Pi > \frac{c}{\lambda_B} \qquad \Longleftrightarrow \qquad \pi \equiv \frac{\lambda_B \Pi}{c} > 1, \tag{3.1}$$

where π denotes the project return.

3.2 Planner's Problem with Exogenous Deadlines

As an intermediate step toward characterizing the optimal contract under an infinite horizon, we next consider a setting where the social planner that operates the project itself faces exogenous deadlines. These deadlines generate inefficiencies, because if the project is worth starting (i.e., if (3.1) holds), then it should be run until the problem is solved. This analysis, which is somewhat involved, is important for understanding elements of the optimal contract in the presence of agency considerations and is also of independent interest, given the ubiquity of deadlines.

Assume that the project is exogenously terminated when time passes a deadline T. Additionally, the deadline is extended by $\Delta \geq 0$ if the advanced skill is acquired before T.[6] In

[5] If the training route were sufficiently more efficient than the execution route, the principal would rely exclusively on training. This would yield results similar to the tangible breakthrough case in Green and Taylor (2016b). To distinguish the analysis from that work, we focus on the parametric regions where there is tension between choosing two paths.

[6] Note that $\Delta = 0$ is permitted. Also, it is possible to consider the case of $\Delta < 0$, but we show that under

other words, the planner faces two deadlines: the original deadline T under no skill acquisition, and the extended deadline $T + \Delta$ upon acquisition of the advanced skill. Given these deadlines, the planner chooses which route (direct or training) to take at each point in time.

We begin by introducing a benchmark policy where it is optimal for the planner to initially employ the training route and later switch to the execution route if the advanced skill is not acquired.

Definition 1. A policy is called a *one-switch policy* if there exists an intermediate deadline $S \in [0, T]$ such that (i) the planner chooses training up to S ($a_t = 0$ for $t \leq S$), (ii) if the advanced skill is acquired before S, the planner tries to solve the problem with the advanced skill until the extended deadline $T + \Delta$, and (iii) if the advanced skill is not acquired by S, the planner switches to execution until the deadline T.

This class of policies includes two extreme cases. A one-switch policy with $S = 0$ does not involve any training, and is referred to as the *execution-only policy.* Conversely, a one-switch policy with $S = T$ does not involve any execution (with the basic skill), and is referred to as the *training-only policy.*

The following theorem shows that the optimal policy takes the form of a one-switch policy.

Theorem 1. *Suppose that the two routes are equally efficient ($\lambda_A = 2\lambda_B$). When the planner faces the deadline T and the extension Δ resulting from the skill acquisition, the optimal policy is characterized as follows:*

- *(a)* ***(Long extension)*** *if $\Delta \geq \bar{\Delta} \equiv \frac{1}{\lambda_A} \log \left[\frac{2\pi-1}{\pi-1}\right]$, the training-only policy is optimal;*
- *(b)* ***(Short extension)*** *if $\Delta < \bar{\Delta}$, there exists $\hat{T} > 0$ such that*
  - *(i) when $T < \hat{T}$, the execution-only policy is optimal;*
  - *(ii) when $T > \hat{T}$, the one-switch policy with the intermediate deadline $T - \hat{T}$ is optimal.*

agency it is optimal to extend–not reduce–the deadline when the agent acquires the advanced skill.

This theorem implies that the basic skill is never used when the deadline extension is sufficiently long. On the other hand, when the deadline extension is relatively short, there exists a time $\hat{T}$ such that the basic skill begins to be used when fewer than $\hat{T}$ units of time remain.

Execution-only vs. training-only policies To provide intuition for Theorem 1, we compare the probability that the problem is solved by the deadline—namely, the solution probability—under the training-only policy (A) and the execution-only policy (B).

When the planner adopts the execution policy, the solution probability is

$$\mathcal{P}^B(T) \equiv \int_0^T \lambda_B \cdot e^{-\lambda_B \tau_m} d\tau = 1 - e^{-\lambda_B T}, \tag{3.2}$$

where τ is the date on which the problem is solved.

Next, when the planner employs the training-only policy, the solution probability is

$$\begin{aligned} \mathcal{P}^A(T, \Delta) \equiv \int_0^T \left[\int_{\tau_s}^{T+\Delta} \lambda_A e^{-\lambda_A(\tau - \tau_s)} d\tau \right] \cdot \lambda_A e^{-\lambda_A \tau_s} d\tau_s \\ = 1 - (1 + \lambda_A \cdot T \cdot e^{-\lambda_A \Delta}) \cdot e^{-\lambda_A T}, \end{aligned} \tag{3.3}$$

where τ_s is the date on which the advanced skill is acquired.

Next we compare the solution probabilities across policies when there is no deadline extension ($\Delta = 0$). The following lemma shows that $\mathcal{P}^B(T, 0)$ and $\mathcal{P}^A(T, 0)$ cross once as T increases. The proof is in Appendix A.2.

Lemma 1. *Suppose that $\lambda_A > \lambda_B$ and $\Delta = 0$. There exists $\check{T}$ such that $\mathcal{P}^B(T, 0) > \mathcal{P}^A(T, 0)$ for all $T < \check{T}$ and $\mathcal{P}^A(T, 0) > \mathcal{P}^B(T, 0)$ for all $T > \check{T}$.*

Notably, under the execution-only policy, the planner needs only one breakthrough, whereas the training-only policy requires two, which is challenging within a short time-frame. Therefore, when the deadline is short ($T < \check{T}$), the solution probability under the

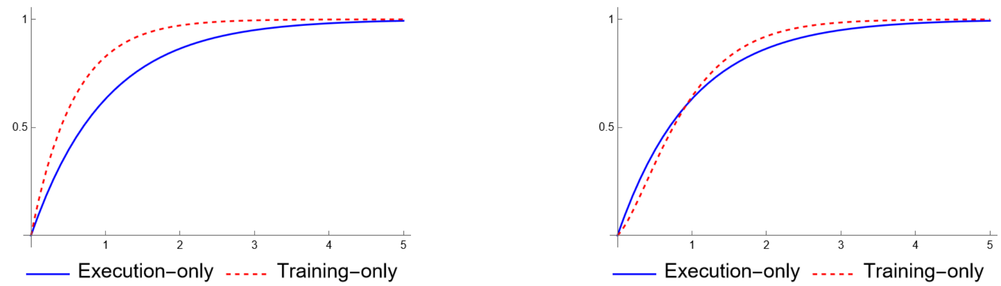

(a) Long Deadline Extension ($\Delta = 1$) (b) Short Deadline Extension ($\Delta = .1$)

Figure 2: Solving Probabilities ($\lambda_A = 2,\ \lambda_B = 1,\ \Pi = 3,\ c = 1$)

execution-only policy is higher than that under the training-only policy. On the other hand, when the deadline is relatively long ($T > \check{T}$), achieving two faster breakthroughs can be easier than achieving one slower breakthrough. In this case, the training-only policy has a higher chance to solve the problem than the execution-only policy. We call this dynamic the 'deadline effect.'

Next, observe that the deadline extension provides additional benefits to the training-only policy: $\mathcal{P}^A(T, \Delta)$ increases as Δ increases. We refer to this as the 'milestone effect' because this increase in the solution probability occurs due to the planner's ability to exploit intermediate progress under the training route.

Based on the deadline and milestone effects, we can infer that the execution-only policy has a higher solution probability than the training-only policy when both the original deadline and the deadline extension are short. Figure 2 illustrates this. The horizontal axis represents the deadline T. With the long deadline extension, as depicted in Figure 2a, the training-only policy always has higher solution probability—the milestone effect outweighs the deadline effect. When the extension is short, as depicted in Figure 2b, the execution-only policy has higher solution probability with short deadlines, whereas the training-only policy has higher solution probabilities with long deadlines.

These effects also play crucial roles in comparing expected surpluses between the two policies: the execution-only policy can achieve higher expected surplus only when both the deadline and the extension are short. This discussion suggests that the one-switch policy—employing (i) the training route when the deadline is distant and (ii) the execution route when the deadline is nigh—is likely to be optimal. Theorem 1 confirms that this is indeed the case.

Comparative statics of the extension cutoff Theorem 1 shows that there exists a cutoff for the deadline extension, $\overline{\Delta}$, such that the training-only policy is optimal if the extension is longer than the cutoff; otherwise, the basic skill is employed near the deadline.

Observe that

$$\frac{\partial \overline{\Delta}}{\partial \pi} = \frac{1}{\lambda_A}\left[\frac{2}{2\pi - 1} - \frac{1}{\pi - 1}\right] = -\frac{1}{\lambda_A(2\pi - 1)(\pi - 1)} < 0,$$

from (3.1). This establishes the following lemma.

Lemma 2. *Suppose that $\lambda_A = 2\lambda_B$ and $\pi > 1$. The extension cutoff $\overline{\Delta}$ decreases in the project return π.*

As the project return increases, the solution probability becomes more important because it is multiplied by π in the expected surplus. This favors the training route, which has a higher solution probability when the deadline is sufficiently long. Consequently, the cutoff $\overline{\Delta}$ decreases as π grows.

4 Agency

4.1 The Necessity of Deadlines

We now return to the contracting problem in which the principal must provide incentives for the agent to solve the problem for her. Note that the first-best outcome cannot be achieved under asymmetric information and limited liability because, without deadlines, the agent would shirk without detection forever and generate an infinite private benefit. Hence, an optimal contract must involve deadlines of some kind, trading off their inherent inefficiency with payment of agency rents.

Remark 1. As is well-known, the principal can generally implement the first-best outcome in moral-hazard settings unless the agent has limited liability or is risk-averse. This is accomplished in general by *selling the agent the project* upfront for a price equal to its first-best value; here $P = \Pi - \max\{2c/\lambda_A,\ c/\lambda_B\}$. Limited liability implies that the agent lacks the resources to make such an upfront payment, necessitating the use of contractual deadlines to control rents.

4.2 Contracts

At the beginning of the game, the principal offers a contract to the agent and fully commits to all contractual terms. If the agent rejects the offer, the principal and the agent receive payoffs of zero. Note that if the agent has neither solved the problem nor acquired the advanced skill, calendar time is the only relevant variable summarizing the public history. We focus on contracts where the agent is always recommended to work; this is without loss of generality, as the principal's operating cost exceeds the agent's private benefit from shirking.

A (deterministic) contract is denoted by $\Gamma \equiv \left\{T, \{a_t, R_t, \hat{\Gamma}^t\}_{0\le t\le T}\right\}$, where each variable is defined at calendar time t as follows:[7]

1. $T \in \mathbb{R}_+$: the deadline date at which the project is terminated absent the solution or skill acquisition.

2. $a_t \in \{0, 1\}$: the principal's recommendation of a route at t where $a_t = 1$ represents *execution* (i.e., solving with the basic skill) and $a_t = 0$ represents *training* (i.e., acquiring the advanced skill);[8]

[7] We show that deterministic contracts are optimal. See Remark 2 for a discussion.

[8] The agent will follow the recommended route because arrival of the skill at t when $a_t = 1$ or arrival of the solution at t when $a_t = 0$ results in termination without payment.

3. $R_t \geq 0$: the monetary payment from the principal to the agent for the solution at t;[9]

4. $\hat{\Gamma}^t \equiv \{T^t, \{R_s^t\}_{t \leq s \leq T^t}\}$: an updated contract when the advanced skill is acquired at t;

   (a) $T^t \geq t$: the deadline date at which the project is terminated;

   (b) $R_s^t \geq 0$: the monetary payment from the principal to the agent for the solution at time s.

4.3 The Optimal Contract

In this subsection, we characterize the optimal contract in the case where the routes are equally efficient. As in the tangible progress case in the mimeo Green and Taylor (2016b), the optimal contract can be implemented with three key properties: (i) the contract is terminated after a deadline; (ii) the reward for the project completion, R_t, linearly diminishes over time; and (iii) the deadline is extended by $1/\lambda_A$ upon advanced skill acquisition.[10]

Since the contract involves a deadline and an extension upon skill acquisition, Theorem 1 suggests that the optimal choice of approaches over time likely either involve one switch from the training route to the execution route—or no switch at all. In light of this conjecture, we define contracts involving the above characteristics as follows.

Definition 2. A contract is called a *one-switch contract* with a final deadline T and an intermediate deadline $S \in (0, T)$ if

(i) the agent is recommended to acquire the advanced skill (training route) by S,

(ii) when the advanced skill is acquired before S, the contract is extended by $1/\lambda_A$ and the reward upon solution at time t is $R_t^A \equiv \phi(T - t + 2/\lambda_A)$,

[9] Since both the principal and the agent are risk neutral and do not discount the future, without loss of generality, all monetary payments to the agent can be backloaded (see, e.g., Ray, 2002). The nonnegativity of R_t is due to limited liability.

[10] The details of these properties will be addressed in Section 4.5. What distinguishes this work from Green and Taylor (2016b) is the presence of the two routes for solving the problem and the principal's choice of which to recommend at each moment in time.

(iii) if the advanced skill is not acquired by S, the agent is recommended to switch and solve the problem using the basic skill (execution route) by T and the reward upon solution at time t is $R_t^B \equiv \phi(T - t + 1/\lambda_B)$, and

(iv) the contract is terminated if the problem is not solved by the deadlines ($T + 1/\lambda_A$ for (ii) and T for (iii)).

When $S = T$, we call the contract a *training-only contract*, and when $S = 0$, we call the contract a *execution-only contract.*

The following theorem shows that the optimal contract indeed takes one of the above forms.

Theorem 2. *Suppose that the execution and training routes are equally efficient ($\lambda_A = 2\lambda_B$). There exist thresholds π_F, π_A and π_B such that $\pi_A > \pi_B > \pi_F \equiv 1 + \phi/c$ and the optimal contract can be implemented as follows:*

(a) when $\pi > \pi_A$, a training-only contract is optimal;

(b) when $\pi_A > \pi > \pi_B$, there exists a one-switch contract that is optimal;

(c) when $\pi_B > \pi > \pi_F$, an execution-only contract is optimal; and

(d) when $\pi < \pi_F$, the project is infeasible.

As discussed in Lemma 2, the execution route is preferred when π is lower and the training route is preferred when π is higher. The above theorem aligns with that intuition. In the subsequent subsections, we provide the details of the derivation of this result.

4.4 Promised Utility and Incentive Compatibility

Following the standard approach of the dynamic contract literature, we consider the agent's promised utility as a state variable and write a contract recursively (e.g., Spear and Srivastava, 1987). For a contract Γ, let $P_0(\Gamma)$ and $U_0(\Gamma)$ be the expected payoffs of the principal and the agent at time 0 when the agent works.

The core of the analysis is the derivation of the principal's value function, denoted by $V(u)$, which represents her maximized expected payoff $P_0(\Gamma)$ subject to the promise-keeping constraint $U_0(\Gamma) = u$ and the incentive compatibility condition, which will be delineated later in this subsection. If a contract Γ satisfies $P_0(\Gamma) = V(u)$ and $U_0(\Gamma) = u$, Γ is said to *implement* a pair of expected payoffs $(V(u), u)$. Once the value function is characterized, the principal solves

$$\bar{u} \equiv \arg\max_{u \geq 0} V(u). \tag{MP}$$

Then, the optimal contract is the contract that implements $(V(\bar{u}), \bar{u})$. In the rest of this subsection, we describe how to derive the value function $V(u)$.

Promised utility upon skill acquisition We begin by considering the principal's problem, given that the advanced skill is acquired at time t. Let u_M^t denote the agent's promised utility, which will be considered as a state variable. Since this case requires only one more breakthrough, it is identical to the single-stage benchmark in Green and Taylor (2016a). They show that the optimal contract is to impose a deadline $t + u_M^t/\phi$ and a linearly diminishing reward schedule $\{\check{R}_s^t\}_{t \leq s \leq t+u_M^t/\phi}$ where

$$\check{R}_s^t = u_M^t + \frac{\phi}{\lambda_A} - \phi(s - t). \tag{4.1}$$

The intuition is that when the agent's promised utility is u_M^t, the principal can incentivize the agent to work at most u_M^t/ϕ units of time. If the principal grants more time without increasing the agent's expected utility, then he will prefer to shirk out the clock.

Incentive compatibility conditions Now consider the agent's problem when the advanced skill has not been acquired. Suppose that the promised utility is u_t at some time t. Under the execution route, if the agent works for a small interval of time $[t, t + dt)$, the breakthrough occurs and the agent receives the reward R_t with a probability $\lambda_B dt$. In this event, however, he loses the continuation utility, thus, the expected payoff from working is $\lambda_B(R_t - u_t)dt$. On the other hand, if he shirks, his payoff is ϕdt. From this, we can derive the incentive compatibility constraint under the execution route ($a_t = 1$):

$$R_t \geq u_t + \frac{\phi}{\lambda_B}. \tag{IC$_1$}$$

Next, under the training route, the agent is compensated in the form of a jump in promised utility upon acquiring the advanced skill. Thus, the expected payoff of working for $[t, t+dt)$ is $\lambda_A(u_M^t - u_t)dt$. Then, the incentive compatibility constraint under the training route ($a_t = 0$) is

$$u_M^t \geq u_t + \frac{\phi}{\lambda_A}. \tag{IC$_0$}$$

4.5 Value Function Characterization

In this subsection, we characterize the value function of the principal. A natural conjecture is that the principal's expected payoff is maximized when the incentive compatibility conditions bind.[11] We outline some key properties of contracts with binding IC conditions, and then characterize the value function.

Deadline and extension With binding IC conditions, the agent's promised utilities should fall at the same rate as the benefit from shirking: $du/dt = \dot{u}_t = -\phi$, or equivalently, $u_t = u_0 - \phi t$. If the problem has not been solved by u_0/ϕ, the promised utility is equal to the agent's outside option 0, thus, the contract is terminated, or equivalently, the deadline of the contract is u_0/ϕ.

[11] See Remark 3 for a discussion of the binding IC condition.

When the training route is chosen, to make (IC_0) bind, we have

$$t + \frac{u_M^t}{\phi} = t + \frac{u_t}{\phi} + \frac{1}{\lambda_A} = \frac{u_0}{\phi} + \frac{1}{\lambda_A}.$$

This implies that upon skill acquisition the updated deadline $t + u_M^t/\phi$ extends the original deadline u_0/ϕ by $1/\lambda_A$.

Linearly diminishing rewards Let T denote the deadline u_0/ϕ. By using $u_t = u_0 - \phi t = \phi(T - t)$, to make ($\text{IC}_1$) bind, the reward for solving the problem at time t via the execution route is

$$R_t = u_t + \frac{\phi}{\lambda_B} = \phi\left(T - t + \frac{1}{\lambda_B}\right),$$

which corresponds to R_t^B in Definition 2.

Next, when the skill is acquired at $\check{t}$, to make (IC_0) bind, we have $u_A^{\check{t}} = u_{\check{t}} + \phi/\lambda_A$. Then, by (4.1), the reward for solving the problem at time $t \in [\check{t}, T + 1/\lambda_A]$ via the training route is

$$\check{R}_t^{\check{t}} = u_{\check{t}} + \frac{\phi}{\lambda_A} + \frac{\phi}{\lambda_A} - \phi(t - \check{t}) = \phi\left(T - t + \frac{2}{\lambda_A}\right),$$

which corresponds to R_t^A in Definition 2.

Value function Based on the above observations, we surmise that the principal's value function under agency is linked to the benchmark planner's problem with a deadline and its extension, as explored in the previous section. Let $\mathcal{W}^*(T, \Delta)$ denote the optimal expected social surplus under the deadline T and extension Δ, derived from the optimal policy in Theorem 1. Then, when the agent's promised utility is u, a conjecture for the principal's value function is the expected social surplus from the deadline u/ϕ and the extension $1/\lambda_A$, $\mathcal{W}^*(u/\phi, 1/\lambda_A)$, net of u.

The following proposition verifies this conjecture, with the proof provided in Appendix B.

Proposition 1. *The principal's value function $\mathcal{V}$ is characterized as follows:*

$$\mathcal{V}(u) = \mathcal{W}^*\left(u/\phi, 1/\lambda_A\right) - u. \tag{4.2}$$

Moreover, $\mathcal{V}$ is concave.

A key step in proving this proposition is finding a contract implementing the pair of the agent's promised utility, u, and the principal's expected payoff, $V(u)$. When choosing a path, the principal's incentives are perfectly aligned with those of the planner—who faces the same deadlines as the principal—in that both want to maximize the expected surplus. Since the planner's policy with at most one switch is optimal, we show that the principal can implement the pair using a contract that involves at most one switch with linearly diminishing rewards (Appendix B.3).

4.6 Proof of Theorem 2

Now that we have characterized the principal's value function, the next step is to pin down the optimal initial promised utility level for the agent, $\bar{u}$, which is the solution to (MP). This will establish the starting point of the contract in Figure 3 and determine the deadline length, $\overline{u}/\phi$. The key tradeoff derives from the fact that the rents needed to keep the agent from shirking grow linearly with the deadline, whereas the marginal benefit from extending the deadline falls because the problem is increasingly likely to be solved before the deadline is reached.

Recall that the basic skill is never employed when the extension is greater than $\overline{\Delta}$ (Lemma 1 (a)) and the cutoff is decreasing in π (Lemma 2). Let π_A be the solution of $\overline{\Delta} = 1/\lambda_A$. Then, for all $\pi > \pi_A$, the execution route will not be employed, even near the deadline. This establishes Theorem 2 (a) and is illustrated in Figure 3 (c).

When $\pi < \pi_A$, Theorem 1 (b) indicates that the optimal approach is switched from the training route to the execution route when $\hat{T}$ units of time remain. In terms of the promised

(a) Low π (b) Intermediate π (c) High π

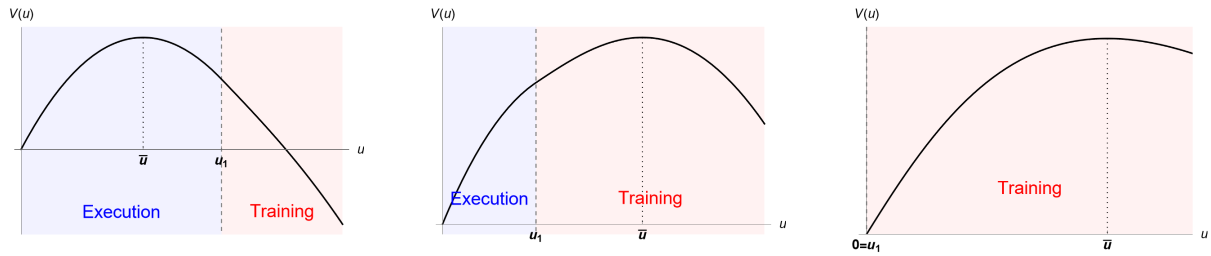

Figure 3: Value functions when routes are equally efficient

utility, the switch happens at $\hat{u}_1 \equiv \phi\hat{T}$. Then, the form of the optimal contract depends on whether $\bar{u}$ is greater than $\hat{u}_1$ or not. For example, the value functions in Figure 3a and 3b both involve a switching point $\hat{u}_1$, however, $\bar{u}$ is greater than $\hat{u}_1$ in Figure 3a and less than $\hat{u}_1$ in Figure 3b. Thus, the optimal contracts are an execution-only contract in Figure 3a and a contract with a switch from the training route to the execution route in Figure 3b. Let π_B be the threshold for the project return where the optimal promised utility $\bar{u}$ is equal to the switching point $\hat{u}_1$. Then, for any $\pi \in (\pi_B, \pi_A)$, the optimal contract will involve a switch, establishing Theorem 2 (b). Conversely, when $\pi < \pi_B$, only the execution route will appear in the optimal contract.

Last, the feasibility of the project depends on whether $\bar{u}$ is greater than 0 or not. When $\bar{u}$ is equal to zero, the principal's expected payoff is maximized at $u = 0$, meaning it is optimal for the principal not to initiate the contract—the project is infeasible. This occurs when the principal's flow profit is negative near the deadline T. Since the promised utility u is close to zero near the deadline, the reward R is approximately ϕ/λ_B. Then, the principal's flow profit in $[T - dt, T]$ is approximately

$$\lambda_B dt \cdot \left(\Pi - \frac{\phi}{\lambda_B}\right) - cdt = \lambda_B \left(\Pi - \frac{\phi + c}{\lambda_B}\right) dt.$$

Therefore, the project is feasible if and only if $\pi = \lambda_B \Pi / c$ is greater than $\pi_F \equiv 1 + \phi/c$. This makes sense because the principal must pay both the operating cost c and (because incentive compatibility binds) the shirking benefit ϕ for the potential duration of the contract. We show that $\pi_B \in (\pi_F, \pi_A)$ (Lemma 10). Then, when $\pi \in (\pi_F, \pi_B)$, the execution-only contract is optimal (Theorem 2 (c)); and when $\pi < \pi_F$, the project becomes infeasible (Theorem 2 (d)).

Remark 2. A mixture of contracts also generates another contract. For example, a contract with a soft deadline—randomly terminating the agent after reaching the soft deadline, as in Green and Taylor (2016a)—can be represented by a mixture of two contracts defined here. However, a mixed contract cannot improve upon the one characterized above. This follows because the value function $\mathcal{V}$ is concave (Lemma 8 (c)).

Consider a set of contracts $\{\Gamma_i\}_{1\leq i\leq n}$ where the agent's expected utility under Γ_i is u_i, and the weight is w_i with $\sum_{i=1}^{n} w_i = 1$ and $\sum_{i=1}^{n} w_i \cdot u_i = u$. The principal's expected payoff from this mixture is $\sum_{i=1}^{n} w_i \cdot P_0(\Gamma_i)$ and the agent's expected utility is u. By concavity, we have $V(u) \geq \sum_{i=1}^{n} w_i V(u_i)$. Additionally, $V(u_i) \geq P_0(\Gamma_i)$ holds for all $1 \leq i \leq n$ because $V(u_i)$ is the principal's maximized expected profit given that the agent's expected payoff is u_i. Thus, $V(u)$ is greater than or equal to the expected payoff of the mixed contract. Hence, any mixed contract cannot improve upon the deterministic contract characterized above.

Remark 3. In Green and Taylor (2016b), there is a parametric region where the incentive compatibility constraint does not bind. This occurs because the deadline extension set by binding IC is not long enough to make the probability of two breakthroughs sufficiently high. In contrast, in this model, when such a situation arises, the principal can switch to the execution path, which is appealing enough to replace the training path. This ensures that IC is always binding.

5 Extension: Efficiency Loss from Training

We last consider the case where training generates an efficiency loss, i.e., $\lambda_A < 2\lambda_B$. This introduces efficiency as another economic force, alongside milestone and deadline effects, that shape the optimal contract.

(a) Low π (b) Intermediate π (c) High π

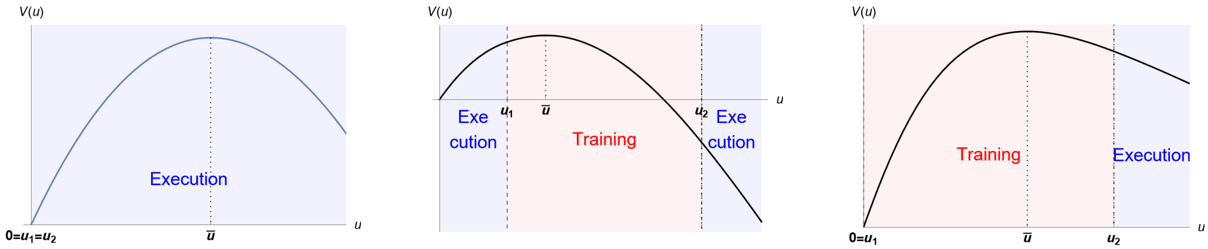

Figure 4: Value functions when the efficiency loss is small

We define a parameter $\eta \equiv \lambda_A/\lambda_B - 1$, which measures the relative efficiency of the training path. Note that $0 < \eta < 1$, and the efficiency loss increases as η decreases. In this section, we characterize the optimal contracts for two cases: (i) when the efficiency loss is small ($\eta > \overline{\eta} \equiv \max\{\sqrt{c/(c+\phi)}, 1/(e-1)\}$); and (ii) when the efficiency loss is large ($\eta < \underline{\eta} \equiv \min\{c/(c+\phi), 1/(e-1)\}$).[12]

In Figures 4 and 5, we illustrate the principal's value functions when there are efficiency losses from training. A key characteristic of these value functions is that the execution route is employed when the promised utility is high, indicating that the deadline is far off. To understand these dynamics, we compare the execution-only and training-only contracts again. As time horizons become longer, the sums of expected payoffs for both players from these contracts converge to the expected surpluses of the no-deadline benchmark: $\Pi - c/\lambda_B$ for the execution-only contract and $\Pi - 2c/\lambda_A$ for the training-only contract. Therefore, efficiency determines which approach should be chosen. Since we focus on the case where the training route is less efficient than the execution route, the principal would choose the execution route when the deadline is distant.

This observation, combined with milestone and deadline effects discussed in the previous sections, leads us to conjecture that there will be two switching points $\hat{u}_1$ and $\hat{u}_2$ in determining the value function. The execution route is chosen when $u > \hat{u}_2$ or $u < \hat{u}_1$, and the training route is chosen when $u \in (\hat{u}_1, \hat{u}_2)$.

[12]These do not cover cases where the efficiency loss is intermediate. In such cases, the form of the optimal contract depends heavily on the parameter values η and Π, resulting in many subcases to analyze. Thus, we focus on the extreme cases to provide results with clear economic implications.

Small efficiency loss When the efficiency loss is relatively small, we show that $\hat{u}_2$ is always greater than the optimal initial promised utility level $\overline{u}$ (Lemma OA.9). It implies that the switch occurs at most once in the optimal contract. Therefore, a result similar to the no-efficiency-loss case holds. In other words, Theorem 2 is robust to small efficiency losses.

Proposition 2. *Suppose that $\eta \in (\overline{\eta}, 1)$, i.e., the efficiency loss from training is small. There exist thresholds $\tilde{\pi}_A(\eta)$ and $\tilde{\pi}_B(\eta)$ with $\tilde{\pi}_A(\eta) > \tilde{\pi}_B(\eta) > \pi_F$ such that the optimal contract is determined as follows:*

(a) when $\pi > \tilde{\pi}_A(\eta)$, a training-only contract is optimal;

(b) when $\tilde{\pi}_A(\eta) > \pi > \tilde{\pi}_B(\eta)$, there exists a one-switch contract that is optimal;

(c) when $\tilde{\pi}_B(\eta) > \pi > \pi_F$, an execution-only contract is optimal.

Large efficiency loss Now suppose that the efficiency loss is large. Figure 5 illustrates that the training path is either not employed at all (for small π) or is employed in the middle of the contract (for large π). As π increases, the milestone effect becomes more significant, as it actuates the monitoring ability of the principal. However, the execution approach is preferred at the beginning of the contract due to its efficiency and at the end of the contract due to the deadline effect. Therefore, if training is ever employed, the contract will involve two switches. We formally define the two-switch contract and state the theorem for the case of large efficiency loss.

Definition 3. A contract is called a *two-switch contract* with a final deadline T and two intermediate deadlines $0 < S_1 < S_2 < T$ if

(i) the agent is recommended to solve the problem using the basic skill by S_1 and the reward upon project completion at time t is R_t^B,

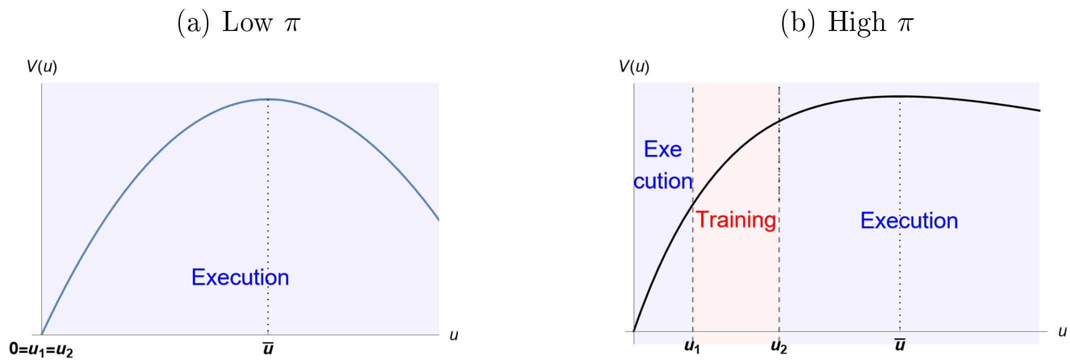

Figure 5: Value functions when the efficiency loss is large

(ii) if the problem is not solved by S_1, the agent is recommended to acquire the advanced skill by S_2, and if the advanced skill is acquired before S_2, the contract is extended by $1/\lambda_A$ with the reward upon solution at time t being R_t^A,

(iii) if the advanced skill is not acquired by S_2, the agent is recommended to solve the problem using the basic skill by T and the reward upon project completion at time t is R_t^B, and

(iv) the contract is terminated if the project is not completed by T.[13]

Proposition 3. *Suppose that η is less than $\underline{\eta}$, i.e., the efficiency loss from training is large. There exists a threshold $\tilde{\pi}_M(\eta)$ with $\tilde{\pi}_M(\eta) > \pi_F$ such that the optimal contract is determined as follows:*

(a) when $\pi > \tilde{\pi}_M(\eta)$, there exists a two-switch contract that is optimal;

(b) when $\tilde{\pi}_M(\eta) > \pi > \pi_F$, an execution-only contract is optimal.

Notably, the principal typically prefers the execution route, since training entails a substantial efficiency loss. When π is sufficiently large, however, she may exploit the monitoring benefit from the milestone effect by choosing the training path. If monitoring occurs, it is optimal to place it in the middle of the contract: efficiency dominates at the start, while the deadline effect dominates near the end.

[13]The rewards R_t^A and R_t^B are defined in the same way as in the one-switch contract.

For high-return projects, the theorem shows that all three forces—efficiency, the milestone effect, and the deadline effect—shape the contract. The principal begins with recommending execution with the basic skill (i.e., efficiency is the primary concern). If the problem remains unsolved by a certain time, she recommends acquiring the advanced skill to monitor the agent and bring them closer to the solution (i.e., the milestone effect becomes the primary concern). She extends the deadline if the agent acquires the advanced skill, but if he does not acquire it until the deadline is near, the principal reverts to execution with the basic skill in a "last-ditch" attempt to solve the problem (i.e., the deadline effect becomes the preeminent motivation).

6 Conclusion

In this article, we study the economic tradeoffs between execution and training in solving a problem in the presence of agency frictions. The optimal contract is determined by the interplay of three effects: efficiency, milestone, and endogenous deadlines. We show that the form of the optimal contract crucially depends on the project return. When the efficiency loss from training does not exist or is small, the optimal contract involves at most one switch. Specifically, if the project return is low, the principal always recommends the agent to solve the problem with a basic skill, whereas if the project return is high, the principal always recommends the agent to acquire the advanced skill, then extend the deadline upon skill acquisition. If the project return is intermediate, it is optimal to begin with training and then switch to execution upon lack of skill acquisition. When the efficiency loss is large, the principal generally recommends the agent to solve a problem with the more efficient basic skill route. However, if the project return is above a certain cutoff, for a short period of time in the middle of the contract, the principal recommends the agent to acquire the advanced skill to mitigate moral hazard (i.e., there may be two switches).

There are numerous avenues open for further research. For example, the principal may be able to design the approaches directly. In this article, we assume that the two approaches are exogenously given and the principal chooses between them. However, in practice, a project manager often designs how many milestones to partition the main project into and how difficult each subproject is. We could also introduce 'learning by doing' into the model. If we assume that the agent learns from early errors, the arrival rate of project completion would increase over time.[14] We leave these intriguing questions–and others–for future work.

References

Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.

Kevin A Bryan and Jorge Lemus. The direction of innovation. *Journal of Economic Theory*, 172:247–272, 2017.

Vannevar Bush. Science: The endless frontier. *Transactions of the Kansas Academy of Science (1903-)*, 48(3):231–264, 1945. URL https://www.jstor.org/stable/3625196.

Christoph Carnehl and Johannes Schneider. On risk and time pressure: When to think and when to do. *Journal of the European Economic Association*, 21(1):1–47, 2023.

Mingliu Chen, Peng Sun, and Yongbo Xiao. Optimal monitoring schedule in dynamic contracts. *Operations Research*, 68(5):1285–1314, 2020.

Francesc Dilmé and Daniel F Garrett. Residual deterrence. *Journal of the European Economic Association*, 17(5):1654–1686, 2019.

Richard J Gilbert and Michael L Katz. Efficient division of profits from complementary innovations. *International Journal of Industrial Organization*, 29(4):443–454, 2011.

[14] This possibility contrasts with the setting considered by Carnehl and Schneider (2023), where learning causes the expected arrival rate to fall.

Brett Green and Curtis R Taylor. Breakthroughs, deadlines, and self-reported progress: contracting for multistage projects. *American Economic Review*, 106(12):3660–99, 2016a.

Brett Green and Curtis R Taylor. On breakthroughs, deadlines, and the nature of progress: Contracting for multistage projects. 2016b. URL http://people.duke.edu/~crtaylor/breakthroughs.pdf.

Jerry R Green and Suzanne Scotchmer. On the division of profit in sequential innovation. *The Rand journal of economics*, pages 20–33, 1995.

Zehao Hu. Financing innovation with unobserved progress. 2014.

Yonggyun Kim and Francisco Poggi. Strategic concealment in innovation races. *Available at SSRN 4638677*, 2025.

Iván Marinovic and Martin Szydlowski. Monitoring with career concerns. *The RAND Journal of Economics*, 53(2):404–428, 2022.

Iván Marinovic and Martin Szydlowski. Monitor reputation and transparency. *American Economic Journal: Microeconomics*, 15(4):1–67, 2023.

Sofia Moroni. Experimentation in organizations. *Theoretical Economics*, 17(3):1403–1450, 2022.

Dmitry Orlov. Frequent monitoring in dynamic contracts. *Journal of Economic Theory*, 206: 105550, 2022.

Tomasz Piskorski and Mark M Westerfield. Optimal dynamic contracts with moral hazard and costly monitoring. *Journal of Economic Theory*, 166:242–281, 2016.

Francisco Poggi. The timing of complementary innovations. 2021.

Debraj Ray. The time structure of self-enforcing agreements. *Econometrica*, 70(2):547–582, 2002.

Stephen E Spear and Sanjay Srivastava. On repeated moral hazard with discounting. *The Review of Economic Studies*, 54(4):599–617, 1987.

Felipe Varas, Iván Marinovic, and Andrzej Skrzypacz. Random inspections and periodic reviews: Optimal dynamic monitoring. *The Review of Economic Studies*, 87(6):2893–2937, 2020.

Christoph Wolf. Informative milestones in experimentation. *University of Mannheim, Working Paper*, 2018. URL https://carnehl.github.io/IME.pdf.

Yu Fu Wong. Dynamic monitoring design. *Available at 4466562*, 2023.

Appendix

A Proofs for Section 3

A.1 Expected Surpluses of Benchmark Policies

In this section, we provide formal representations of expected social surpluses under the execution-only and the training-only policies.

Suppose that the principal employs the execution-only policy. Recall that the solution probability $\mathcal{P}^B(T)$ is derived in (3.2). Next, the expected duration of problem-solving is

$$\mathcal{D}^B(T) \equiv \int_0^T \tau \cdot \lambda_B \cdot e^{-\lambda_B \tau} d\tau + T \cdot e^{-\lambda_B T} = \frac{1}{\lambda_B}(1 - e^{-\lambda_B T}). \tag{A.1}$$

Then, the expected social surplus of the execution-only policy is

$$\mathcal{W}^B(T) \equiv \Pi \cdot \mathcal{P}^B(T) - c \cdot \mathcal{D}^B(T) = \left(\Pi - \frac{c}{\lambda_B}\right) \cdot (1 - e^{-\lambda_B T}). \tag{A.2}$$

Now suppose that the planner employs the training-only policy. Also recall that the solution probability $\mathcal{P}^B(T)$ is derived in (3.3). Conditional on advanced skill acquisition at τ_s, the expected duration is

$$\mathcal{D}^1(T, \Delta, \tau_s) \equiv \int_{\tau_s}^{T+\Delta} \tau \cdot \lambda_A e^{-\lambda_A(\tau - \tau_s)} d\tau + (T + \Delta) \cdot e^{-\lambda_A(T+\Delta-\tau_s)}.$$

Then, the expected duration of the project can also be derived as follows:

$$\begin{aligned}\mathcal{D}^A(T, \Delta) &\equiv \int_0^T \mathcal{D}^1(T, \Delta, \tau_s) \cdot \lambda_A e^{-\lambda_A \tau_s} d\tau_s + T \cdot e^{-\lambda_A T} \\ &= \frac{2}{\lambda_A}(1 - e^{-\lambda_A T}) - T \cdot e^{-\lambda_A(T+\Delta)}.\end{aligned} \tag{A.3}$$

Then, the expected social surplus of the training-only policy is

$$\begin{aligned} \mathcal{W}^A(T,\Delta) &\equiv \Pi \cdot \mathcal{P}^A(T,\Delta) - c \cdot \mathcal{D}^A(T,\Delta) \\ &= \left(\Pi - \frac{2c}{\lambda_A}\right) \cdot (1 - e^{-\lambda_A T}) - \lambda_A \left(\Pi - \frac{c}{\lambda_A}\right) \cdot T \cdot e^{-\lambda_A (T+\Delta)}. \end{aligned} \tag{A.4}$$

A.2 Proof of Lemma 1

Proof of Lemma 1. Observe that $\mathcal{P}^B(T,0) \geq \mathcal{P}^A(T,0)$ is equivalent to:

$$(1 + \lambda_A T) \geq e^{(\lambda_A - \lambda_B)T}.$$

Note that the equality holds at $T = 0$. While the left hand side linearly increases with the slope λ_A, the right-hand side exponentially increases and the slope at $T = 0$ is $\lambda_A - \lambda_B \in (0, \lambda_A)$. Therefore, for small enough T, $\mathcal{P}^B(T,0) > \mathcal{P}^A(T,0)$, but there exists $\check{T} > 0$, which makes the two sides equal. Then, we have $\mathcal{P}^B(T,0) > \mathcal{P}^A(T,0)$ for all $T < \check{T}$ and $\mathcal{P}^B(T,0) < \mathcal{P}^A(T,0)$ for all $T > \check{T}$. □

A.3 Proof of Theorem 1

Expected surplus upon skill acquisition Let W_x^1 represent the expected surplus when the advanced skill is acquired and the remaining time is x. By following steps similar to those used in the derivation of (A.2), we have

$$W_x^1 \equiv \left(\Pi - \frac{c}{\lambda_A}\right) \cdot (1 - e^{-\lambda_A x}). \tag{A.5}$$

Suppose that the advanced skill is acquired at calendar time $T - z$, meaning that z units of time remain until the original deadline. Then, the skill acquisition extends the deadline by Δ, giving the planner $z + \Delta$ units of time to solve the problem. Therefore, the expected surplus in this situation is $W^1_{z+\Delta}$.

Expected surplus without skill acquisition Now we consider the situation that neither the problem is solved nor the advanced skill is acquired by calendar time $T - z$. Then, the (optimal) expected surplus $W^{0,\Delta}_z$ can heuristically be written as follows:

$$W^{0,\Delta}_z = \max_{a_z \in \{0,1\}} \begin{array}{l} \Pi \cdot \lambda_B a_z \cdot dz + W^1_{z+\Delta} \cdot \lambda_A (1 - a_z) \cdot dz - c\, dz \\ + \{1 - \lambda_B a_z \cdot dz - \lambda_A (1 - a_z) \cdot dz\} \cdot W^{0,\Delta}_{z-dz}. \end{array}$$

By using a Taylor expansion, $W^{0,\Delta}_{z-dz} = W^{0,\Delta}_z - \dot{W}^0_z dz$, canceling out $W^{0,\Delta}_z$ on both sides, and taking the limit as $dz \to 0$, we obtain a Hamilton-Jacobi-Bellman (HJB) equation:

$$\dot{W}^{0,\Delta}_z = \max_{a_z \in \{0,1\}} \left[\lambda_B a_z \cdot (\Pi - W^{0,\Delta}_z) + \lambda_A (1 - a_z) \cdot (W^1_{z+\Delta} - W^{0,\Delta}_z) - c \right]. \tag{HJB$_W$}$$

Since the project is terminated at the deadline, $W^{0,\Delta}_0 = 0$. Then, by using standard verification arguments (e.g., Proposition 3.2.1 in Bertsekas (1995)), if a function $w : [0, T] \to \mathbb{R}$ is differentiable and satisfies (HJB$_W$) and $w(0) = 0$, then, $w(z) = W^{0,\Delta}_z$.

Expected surplus of the one-switch policy We now derive the expected surplus of the one-switch policy with an intermediate deadline S and a deadline T. Denote $Z \equiv T - S$. The one-switch policy implies that $a_z = 1$ for all $0 \le z < Z$, and $a_z = 0$ for all $Z \le z \le T$. Let $\mathcal{W}^{AB}_Z(z, \Delta)$ denote the expected surplus of this policy when the remaining time is z. The following differential equations then hold, where $\dot{W}^{AB}_Z = \frac{\partial W^{AB}_Z(z,\Delta)}{\partial z}$:

$$\dot{\mathcal{W}}^{AB}_Z(z, \Delta) = \lambda_A \cdot \left(W^1_{z+\Delta} - \mathcal{W}^{AB}_Z(z, \Delta) \right) - c \quad \text{for } z \ge Z, \tag{A.6}$$

$$\dot{\mathcal{W}}^{AB}_Z(z, \Delta) = \lambda_B \cdot \left(\Pi - \mathcal{W}^{AB}_Z(z, \Delta) \right) - c \quad \text{for } z < Z. \tag{A.7}$$

By solving this with $\mathcal{W}_Z^{AB}(0,\Delta)=0$, we derive

$$\mathcal{W}_Z^{AB}(z,\Delta)=\begin{cases}\mathcal{W}^B(z), & \text{if } z\le Z,\\ \left(\Pi-\frac{2c}{\lambda_A}\right)\cdot\left(1-e^{-\lambda_A(z-Z)}\right)+\mathcal{W}^B(Z)\cdot e^{-\lambda_A(z-Z)} & \\ \qquad -\lambda_A\left(\Pi-\frac{c}{\lambda_A}\right)\cdot(z-Z)\cdot e^{-\lambda_A(z+\Delta)}, & \text{if } z>Z.\end{cases} \tag{A.8}$$

Also note that $\mathcal{W}_0^{AB}(z,\Delta)=\mathcal{W}^A(z,\Delta)$ is the expected surplus of the training-only policy and $\mathcal{W}_T^{AB}(z,\Delta)=\mathcal{W}^B(z)$ is the expected surplus of the execution-only policy.

We will prove the theorem by showing that there exists $Z\in[0,T]$ such that $\mathcal{W}_Z^{AB}(z,\Delta)$ solves (HJB_W).

Optimal path at the deadline Note that $W_\Delta^1=\left(\Pi-\frac{c}{\lambda_A}\right)\cdot(1-e^{-\lambda_A\Delta})$ and $W_0^{0,\Delta}=0$. Then, at the deadline, the training path is preferred over the execution path if and only if

$$\begin{aligned} & \lambda_B(\Pi-W_0^{0,\Delta})\le\lambda_A(W_\Delta^1-W_0^{0,\Delta})\\ \iff\quad & \lambda_B\Pi\le(\lambda_A\Pi-c)\cdot(1-e^{-\lambda_A\Delta}).\end{aligned} \tag{A.9}$$

With $\lambda_A=2\lambda_B$ and simple algebra, we can derive that (A.9) is equivalent to $\Delta\ge\bar{\Delta}$.

Optimal Policy Derivation We introduce two crucial lemmas, then complete the proof of Theorem 1.

Lemma 3. *Suppose that $\lambda_A=2\lambda_B$ and $\Delta<\bar{\Delta}$. Then, there exists $\hat{T}$ such that (i) $\lambda_B(\Pi-\mathcal{W}^B(z))>\lambda_A(W_{z+\Delta}^1-\mathcal{W}^B(z))$ for all $z<\hat{T}$; and (ii) $\lambda_B(\Pi-\mathcal{W}^B(\hat{T}))=\lambda_A(W_{\hat{T}+\Delta}^1-\mathcal{W}^B(\hat{T}))$.*

Proof of Lemma 3. Define $H_z^1\equiv\lambda_A\left(W_{z+\Delta}^1-\mathcal{W}^B(z)\right)-\lambda_B\left(\Pi-\mathcal{W}^B(z)\right)$ and $x\equiv e^{-\lambda_B z}$. Then, with some algebra, H_z^1 is equivalent to

$$H_1(x)\equiv(\lambda_A\Pi-c)\cdot(1-e^{-\lambda_A\Delta}\cdot x^{\frac{\lambda_A}{\lambda_B}})-\lambda_B\Pi-(\lambda_A-\lambda_B)\cdot\left(\Pi-\frac{c}{\lambda_B}\right)\cdot(1-x). \tag{A.10}$$

By using $\lambda_A = 2\lambda_B$, with some algebra, $H_1(x)$ can be rewritten as follows:

$$H_1(x) = (\lambda_B \Pi - c) \cdot x - (\lambda_A \Pi - c) \cdot e^{-\lambda_A \Delta} \cdot x^2$$

Define

$$\hat{x} \equiv \frac{\lambda_B \Pi - c}{\lambda_A \Pi - c} \cdot e^{\lambda_A \Delta}. \tag{A.11}$$

Note that $\hat{x} < 1$ when $\Delta < \overline{\Delta}$. Additionally, observe that $H_1(\hat{x}) = 0$ and $H_1(x) < 0$ for all $\hat{x} < x \leq 1$.

Now set $\hat{T} \equiv -\frac{\log(\hat{x})}{\lambda_B}$. Then, $\hat{T} > z$ is equivalent to $x > \hat{x}$, which implies $H_1(x) < 0$. Therefore, for all $z < \hat{T}$, $\lambda_B(\Pi - \mathcal{W}^B(z)) > \lambda_A(W^1_{z+\Delta} - \mathcal{W}^B(z))$. In addition, $H_1(\hat{x}) = 0$ implies $\lambda_B(\Pi - \mathcal{W}^B(\hat{T})) = \lambda_A(W^1_{\hat{T}+\Delta} - \mathcal{W}^B(\hat{T}))$. □

Lemma 4. *Suppose that $\lambda_A = 2\lambda_B$ and $\lambda_A(W^1_{Z+\Delta} - \mathcal{W}^B(Z)) \geq \lambda_B(\Pi - \mathcal{W}^B(Z))$ for some $Z \geq 0$. Then, $\lambda_A(W^1_{z+\Delta} - \mathcal{W}^{AB}_Z(z, \Delta)) \geq \lambda_B(\Pi - \mathcal{W}^{AB}_Z(z, \Delta))$ for all $z > Z$.*

Proof of Lemma 4. Define $H^2_z \equiv \lambda_A(W^1_{z+\Delta} - \mathcal{W}^{AB}_Z(z, \Delta)) - \lambda_B(\Pi - \mathcal{W}^{AB}_Z(z, \Delta))$ and $y \equiv e^{-\lambda_A(z-Z)}$. Note that, for any $z > Z$,

$$\begin{aligned}\mathcal{W}^{AB}_Z(z, \Delta) = &\mathcal{W}^B(Z) + \left(\Pi - \frac{2c}{\lambda_A} - \mathcal{W}^B(Z)\right) \cdot (1 - y) \\ &+ \left(\Pi - \frac{c}{\lambda_A}\right) \cdot e^{-\lambda_A(Z+\Delta)} \cdot \log(y) \cdot y\end{aligned}$$

and

$$W^1_{z+\Delta} = W^1_{Z+\Delta} + \left(\Pi - \frac{c}{\lambda_A} - W^1_{Z+\Delta}\right) \cdot (1 - y)\,.$$

Then, with some algebra, H^2_z can be rewritten as follows:

$$H_2(y) \equiv H^2_Z + h_1 \cdot (1 - y) + h_2 \cdot \log(y) \cdot y \tag{A.12}$$

where

$$h_1 \equiv \lambda_A \cdot \left(\Pi - \frac{c}{\lambda_A} - W^1_{Z+\Delta}\right) - (\lambda_A - \lambda_B) \cdot \left(\Pi - \frac{2c}{\lambda_A} - \mathcal{W}^B(Z)\right),$$
$$h_2 \equiv \left(\Pi - \frac{c}{\lambda_A}\right) \cdot e^{-\lambda_A(Z+\Delta)} > 0.$$

Observe that

$$H_2''(y) = -\frac{h_2}{y} < 0,$$

i.e., H_2 is strictly concave. By the assumption, we have $H_2(1) = H_Z^2 \geq 0$. In addition, we have

$$\lim_{y\to 0} H_2(y) = \lambda_A\left(\Pi - \frac{c}{\lambda_A}\right) - \lambda_B\Pi - (\lambda_A - \lambda_B)\left(\Pi - \frac{2c}{\lambda_A}\right) = \frac{\lambda_A - 2\lambda_B}{\lambda_A}c = 0. \quad \text{(A.13)}$$

Then, by using the strict concavity of H_2, $H_2(1) \geq 0$ and $\lim_{y\to 0} H_2(y) = 0$, we have $H_2(y) \geq 0$ for all $y \in (0,1)$. Therefore, $\lambda_A(W^1_{z+\Delta} - \mathcal{W}^{AB}_Z(z,\Delta)) \geq \lambda_B(\Pi - \mathcal{W}^{AB}_Z(z,\Delta))$ for all $z > Z$. □

Proof of Theorem 1. (a) Suppose that $\Delta \geq \bar{\Delta}$. From (A.9) and $\mathcal{W}^B(0) = 0$, we have $\lambda_A(W^1_\Delta - \mathcal{W}^B(0)) \geq \lambda_B(\Pi - \mathcal{W}^B(0))$. Then, by Lemma 4, $\lambda_A(W^1_{z+\Delta} - \mathcal{W}^{AB}_0(z,\Delta)) \geq \lambda_B(\Pi - \mathcal{W}^{AB}_0(z,\Delta))$ for all $z > 0$. Then, by (A.6), $\mathcal{W}^A_{z,\Delta} = \mathcal{W}^{AB}_0(z,\Delta)$ solves (HJB_W) for all $z \in \mathbb{R}_+$, i.e., the training-only policy is optimal.

(b) Suppose that $\Delta < \bar{\Delta}$. Let $\hat{T}$ be the time defined in Lemma 3. If $T < \hat{T}$, $\lambda_B(\Pi - \mathcal{W}^B(z)) > \lambda_A(W^1_{z+\Delta} - \mathcal{W}^B(z))$ for all $z \in [0,T]$. Then, by (A.7), $\mathcal{W}^B(z)$ solves (HJB_W) for all $z \in [0,T]$, i.e., the execution-only policy is optimal.

Now consider the case with $T \geq \hat{T}$. Note that $\mathcal{W}^{AB}_{\hat{T}}(z,\Delta) = W^B_z$ for all $z \in [0,\hat{T}]$. From Lemma 3, we have $\lambda_B(\Pi - \mathcal{W}^{AB}_{\hat{T}}(z,\Delta)) > \lambda_A(W^1_{z+\Delta} - \mathcal{W}^{AB}_{\hat{T}}(z,\Delta))$ for all $z \in (\hat{T},T]$. Then, by (A.7), $\mathcal{W}^{AB}_{\hat{T}}(z,\Delta)$ solves (HJB_W) for all $z > \hat{T}$. In addition, we have $\lambda_B(\Pi - \mathcal{W}^{AB}_{\hat{T}}(\hat{T},\Delta)) = \lambda_A(W^1_{\hat{T}+\Delta} - \mathcal{W}^{AB}_{\hat{T}}(\hat{T},\Delta))$. By applying Lemma 4 for $Z = \hat{T}$, we have $\lambda_A(W^1_{z+\Delta} - \mathcal{W}^{AB}_{\hat{T}}(z,\Delta)) \geq \lambda_B(\Pi - \mathcal{W}^{AB}_{\hat{T}}(z,\Delta))$ for all $z > \hat{T}$. Then, by (A.6), $\mathcal{W}^{AB}_{\hat{T}}(z,\Delta)$ solves (HJB$_W$) for all $z > \hat{T}$. Therefore, $\mathcal{W}^{AB}_{\hat{T}}(z,\Delta)$ solves (HJB$_W$) for all $z \in [0,T]$, i.e., the one-switch policy with $T - \hat{T}$ is optimal.

□

B Proofs for Section 4

B.1 Expected Payoffs

In this section, we formally present the expected payoffs of the principal and the agent, conditional on the agent's effort schedule.

Post-skill-acquisition payoffs We begin with the subgame where the advanced skill is acquired at time t. Let u^t_M denote the agent's continuation utility—the expected payoff when the agent works until the problem is solved or the deadline T^t is reached.[15]

Since this subgame only requires one more breakthrough, it is identical to the single-stage benchmark of Green and Taylor (2016a). They show that the principal's value function V_M is characterized as follows:

$$V_M(u^t_M) = W^1_{u^t_M/\phi} - u^t_M = \left(\Pi - \frac{c}{\lambda_A}\right)\left(1 - e^{-\frac{\lambda_A}{\phi}u^t_M}\right) - u^t_M. \tag{B.1}$$

Expected payoffs at time 0 Now consider the problem at time 0. The agent's admissible effort schedule (prior to either solution or skill acquisition) is $\mathcal{B} \equiv \{\{\beta_t\}_{0\leq t\leq T} \ : \ \beta_t \in \{0,1\}\}$. In this case, any arrival depends not only on the agent's effort (β) but also the principal's path choice ($a = \{a_t\}_{0\leq t\leq T}$). Given (a,β), the probability that neither the solution nor the

[15] Specifically, the updated contract $\hat{\Gamma}^t$ determines

advanced skill has arrived by time τ is $f(a,\beta;0,\tau)$, where

$$f(a,\beta;t,\tau) \equiv e^{-\lambda_B \int_t^\tau a_s\beta_s ds - \lambda_A \int_t^\tau (1-a_s)\beta_s ds}.$$

Accordingly, the probability density of the solution at time τ is $\lambda_B a_\tau \beta_\tau \cdot f(a,\beta;0,\tau)$ and that of the skill acquisition at time τ is $\lambda_A(1-a_\tau)\beta_\tau \cdot f(a,\beta;0,\tau)$. If the agent follows the recommendation, the above expressions simplify to: $\tilde{f}(a;t,\tau) \equiv e^{-\lambda_B \int_t^\tau a_s ds - \lambda_A \int_t^\tau (1-a_s)ds}$, $\lambda_B a_\tau \cdot \tilde{f}(a;t,\tau)$, and $\lambda_A(1-a_\tau)\cdot\tilde{f}(a;t,\tau)$.

Given a contract Γ, an effort schedule β and under the assumption that the agent follows the recommendation in the updated contract, the principal's expected payoff at time 0 is:

$$P_0(\beta,\Gamma) \equiv \int_0^T \{(\Pi - R_\tau)\cdot\lambda_B a_\tau\beta_\tau + V_M(u_M^\tau)\cdot\lambda_A(1-a_\tau)\beta_\tau - c\}\cdot f(a,\beta;0,\tau)d\tau.$$

while the agent's expected payoff at time t is:

$$U_0(\beta,\Gamma) \equiv \int_0^T \{R_\tau\cdot\lambda_B a_\tau\beta_\tau + u_M^\tau\cdot\lambda_A(1-a_\tau)\beta_\tau + \phi(1-\beta_\tau)\}\cdot f(a,\beta;0,\tau)d\tau.$$

Similarly, the expected payoffs of the principal and the agent when the agent follows the recommendation can be obtained by evaluating the general expressions using $\tilde{f}(a;0,\tau)$ in place of $f(a,\hat{\beta}^t;0,\tau)$. We denote these by $\tilde{P}_0(\Gamma)$ and $\tilde{U}_0(\Gamma)$, respectively.

Incentive Compatibility Using the terms defined above, we define incentive compatibility (IC) of contracts as follows.

Definition 4. A contract $\Gamma = \left\{T, \{a_t, R_t, \hat{\Gamma}^t\}_{0\le t\le T}\right\}$ is *incentive compatible* if $\tilde{U}_0(\Gamma) \geq U_0(\beta,\Gamma)$ for all $\beta \in \mathcal{B}$.

The objective of the principal is to find a contract Γ that maximizes her ex ante expected payoff $P_0(\Gamma)$ subject to the incentive compatibility constraint and the limited liability constraints $R_\tau \geq 0$ for all $\tau \in [0,T]$ and $R_\tau^t \geq 0$ for all $t \in [0,T]$ and $\tau \in [t, T^t]$. Designate such a contract as an *optimal contract*.

B.2 Recursive Formulation

B.2.1 The Agent's Problem

Given a contract Γ, define the continuation utility that the advanced skill is acquired at t as $u_M^t = \tilde{U}^t(\hat{\Gamma}^t)$ and the continuation utility at t when neither the problem is solved nor the advanced skill is acquired as

$$u_t = \tilde{U}_t(\Gamma) \equiv \int_t^T \left\{ R_\tau \cdot \lambda_B a_\tau + u_M^t \cdot \lambda_A (1 - a_\tau) \right\} \cdot \tilde{f}(a; t, \tau) d\tau.$$

Observe that

$$0 = \dot{u}_t + (R_t - u_t)\lambda_B a_t + (u_M^t - u_t)\lambda_A (1 - a_t) \tag{HJB$_{PK}$}$$

where $\dot{u}_t \equiv \frac{du_t}{dt}$.

Also note that (IC$_0$) for $a_t = 0$ and (IC$_1$) for $a_t = 1$ can be written together as follows:

$$(R_t - u_t)\lambda_B a_t + (u_M^t - u_t)\lambda_A (1 - a_t) \geq \phi. \tag{IC}$$

The following lemma shows that this condition serves as a sufficient condition for incentive compatibility defined in Definition 4.

Lemma 5. *Given a contract Γ, suppose that there exists a continuous and differentiable process $\{u_t\}_{0 \leq t \leq T}$ satisfying (HJB$_{PK}$) and $u_T = 0$, and (IC) holds for $0 \leq t \leq T$. Then, Γ is incentive compatible.*

Proof of Lemma 5. The proof is inspired by Proposition 3.2.1 in Bertsekas (1995).

Consider an arbitrary admissible action $\beta \in \mathcal{B}$. Using (HJB$_{PK}$), (IC$_0$) and (IC$_1$), we can derive that

$$0 \geq \dot{u}_t + (R_t - u_t)\lambda_B a_t \beta_t + (u_M^t - u_t)\lambda_A(1 - a_t)\beta_t + \phi(1 - \beta_t),$$

or equivalently,

$$-\dot{u}_t + (\lambda_B a_t + \lambda_A(1 - a_t))\beta_t \cdot u_t \geq R_t \cdot \lambda_B a_t \beta_t + u_M^t \cdot \lambda_A(1 - a_t)\beta_t + \phi(1 - \beta_t).$$

It is further equivalent to

$$\frac{d}{dt}\left[-u_t \cdot f(a, \beta; 0, t)\right] \geq \left[\left(R_t \lambda_B a_t + u_M^t \lambda_A(1 - a_t)\right)\beta_t + \phi(1 - \beta_t)\right] \cdot f(a, \beta; 0, t).$$

By integrating the above inequality from 0 to T and using $u_T = 0$, we can derive that

$$u_0 \geq \int_0^T \left[\left(R_t \lambda_B a_t + u_M^t \lambda_A(1 - a_t)\right)\beta_t + \phi(1 - \beta_t)\right] \cdot f(a, \beta; 0, t)dt = U_0(\beta, \Gamma)$$

for all $\beta \in \mathcal{B}$. Furthermore, the equality holds when $\beta_t = 1$ for all $t \in [0, T]$ from (HJB$_{PK}$), i.e., u_0 is equal to $\tilde{U}_0(\Gamma)$. Thus, $\tilde{U}_0(\Gamma) \geq U_0(\beta, \Gamma)$ for all $\beta \in \mathcal{B}$, which implies incentive compatibility. □

B.2.2 The Principal's Problem

We now explore how the principal's value function $V(u_t)$ evolves. Note that $V(0) = 0$ since the agent will not participate in the contract when the continuation utility is zero. This will serve as a boundary condition. The value function $V(u_t)$ can be heuristically written as follows:

$$V(u_t) = \max_{R_t, u_M^t, a_t} \begin{array}{l} - cdt + (\Pi - R_t)\lambda_B a_t dt + V_M(u_M^t)\lambda_A(1 - a_t)dt \\ + \{1 - \lambda_B a_t dt - \lambda_A(1 - a_t)dt\} \cdot V(u_{t+dt}) \end{array}$$

subject to (IC).

By using $V(u_{t+dt}) = V(u_t) + V'(u_t)\dot{u}_t dt + o(dt)$, canceling $V(u_t)$ on both sides, taking the limit as $dt \to 0$ and plugging (HJB_{PK}) in, we obtain an HJB equation:

$$0 = \max_{R,u_M,a} \mathcal{J}(V(\cdot), R, u_M, a). \tag{HJB$_V$}$$

where

$$\begin{aligned}\mathcal{J}(V(\cdot), R, u_M, a) \equiv & -c + (\Pi - R - V(u))\lambda_B a + (V_M(u_M) - V(u))\lambda_A(1-a) \\ & - \{(R-u)\lambda_B a + (u_M - u)\lambda_A(1-a)\} \cdot V'(u)\end{aligned} \tag{B.2}$$

Then, the principal's problem is to solve (HJB_V) subject to (IC) with the boundary condition $V(0) = 0$. The following lemma shows that the solution of the problem maximizes the principal's expected payoff subject to a promise keeping constraint $U_0(\Gamma) = u$.

Lemma 6 (Verification Lemma)**.** *Suppose that a differentiable and concave function $\tilde{V}$ solves* (HJB_V) *subject to* (IC) *with the boundary condition $\tilde{V}(0) = 0$. Then, for any incentive-compatible contract Γ with $U_0(\Gamma) = u$,*

$$\tilde{V}(u) \geq \tilde{P}_0(\Gamma).$$

Proof of Lemma 6. Consider an arbitrary (deterministic) incentive-compatible contract Γ where the agent's expected payoff is given by u_t. The payoff to the principal under Γ is

$$\begin{aligned}\tilde{P}_0(\Gamma) = & \int_0^T (\Pi - R_t - c \cdot t) \cdot \lambda_B a_t \tilde{f}(a;0,t)dt \\ & + \int_0^T \big(V_M(u_M^t) - c \cdot t\big) \cdot \lambda_A(1-a_t)\tilde{f}(a;0,t)dt - c \cdot T \cdot \tilde{f}(a;0,T) \\ = & \int_0^T \big((\Pi - R_t) \cdot \lambda_B a_t + V_M(u_M^t) \cdot \lambda_A(1-a_t) - c\big)\,\tilde{f}(a;0,t)dt.\end{aligned}$$

Since $\tilde{V}$ solves the HJB equation, we have

$$
\begin{aligned}
0 \geq & -c + (\Pi - R_t - \tilde{V}(u_t))\lambda_B a_t + (V_M(u_M^t) - \tilde{V}(u_t))\lambda_A(1-a_t) \\
& - \left\{(R_t - u_t) \cdot \lambda_B a_t + (u_M^t - u_t) \cdot \lambda_A(1-a_t)\right\} \tilde{V}'(u_t).
\end{aligned}
$$

By using (HJB$_{PK}$), rearranging, and multiplying by $\tilde{f}(a;0,t)$, we can obtain that

$$
\begin{aligned}
& (\lambda_B a_t + \lambda_A(1-a_t))\tilde{f}(a;0,t) \cdot \tilde{V}(u_t) - \tilde{f}(a;0,t) \cdot \tilde{V}'(u_t)\dot{u}_t \\
\geq & \tilde{f}(a;0,t) \left((\Pi - R_t) \cdot \lambda_B a_t + V_M(u_M^t) \cdot \lambda_A(1-a_t) - c\right)
\end{aligned}
\tag{B.3}
$$

Note that

$$
\frac{d}{dt}\left(-\tilde{f}(a;0,t) \cdot \tilde{V}(u_t)\right) = (\lambda_B a_t + \lambda_A(1-a_t))\tilde{f}(a;0,t) \cdot \tilde{V}(u_t) - \tilde{f}(a;0,t) \cdot \tilde{V}'(u_t)\dot{u}_t.
$$

Then, by integrating (B.3) over $[0,T]$ and noting that $\tilde{f}(a;0,0) = 1$, $u_T = 0$ and $\tilde{V}(0) = 0$, we have

$$
\begin{aligned}
& \tilde{V}(u_0) = \tilde{V}(u_0) - \tilde{f}(a;0,T)\tilde{V}(u_T) \\
\geq & \int_0^T \tilde{f}(a;0,t) \cdot \left((\Pi - R_t) \cdot \lambda_B a_t + V_M(u_M^t) \cdot \lambda_A(1-a_t) - c\right) dt = \tilde{P}_0(\Gamma).
\end{aligned}
$$

Therefore, $\tilde{V}(u_0)$ is greater than or equal to any deterministic contract where the agent's expected payoff is equal to u_0. Since $\tilde{V}$ is assumed to be concave, it is greater than or equal to any randomized contract. □

B.3 Value Function Candidates and Implementation

Lemma 7. *The following statements hold.*

(a) An execution-only contract with the deadline u/ϕ implements a pair of expected payoffs

of the principal and the agent $(V^B(u), u)$ *where*

$$V^B(u) \equiv \mathcal{W}^B\left(\frac{u}{\phi}\right) - u. \tag{B.4}$$

(b) *When* $0 < u_1 < u$*, a one-switch contract with the intermediate deadline* $(u - u_1)/\phi$ *and the final deadline* u/ϕ *implements* $(V^{AB}(u|u_1), u)$ *where*

$$V^{AB}(u|u_1) \equiv \mathcal{W}^{AB}_{u_1/\phi}\left(\frac{u}{\phi}, \frac{1}{\lambda_A}\right) - u. \tag{B.5}$$

(c) *A training-only contract with the deadline* u/ϕ *implements* $(V^{AB}(u|0), u)$.

(d) *The following differential equations hold:*

$$\phi V^{B\prime}(u) = \lambda_B\left(\Pi - \frac{\phi}{\lambda_B} - u - V^B(u)\right) - c, \tag{B.6}$$

$$\phi V^{AB\prime}(u|u_1) = \lambda_A\left(V_M(u + \tfrac{\phi}{\lambda_A}) - V^{AB}(u|u_1)\right) - c. \tag{B.7}$$

Together with Theorem 1, this lemma implies that $(\mathcal{V}(u), u)$—defined in (4.2)—can be implemented by one of the above three contracts. Moreover, there exists $\hat{u}_1 \geq 0$ such that $\mathcal{V}$ can be rewritten as follows:

$$\mathcal{V}(u) = \begin{cases} V^B(u), & \text{if } u < \hat{u}_1, \\ V^{AB}(u|\hat{u}_1), & \text{if } u \geq \hat{u}_1. \end{cases} \tag{B.8}$$

Specifically, $\hat{u}_1$ is chosen to be equal to $\phi\hat{T}$ if $1/\lambda_A < \overline{\Delta}$, and 0 if $1/\lambda_A \geq \overline{\Delta}$. The following lemma provides useful properties of $\mathcal{V}$ and $\hat{u}_1$.

Lemma 8. *Suppose that* $\lambda_A = 2\lambda_B$*. The following statements hold.*

(a) *if* $\hat{u}_1 > 0$*,* $V^{AB\prime}(\hat{u}_1|\hat{u}_1) = V^{B\prime}(\hat{u}_1)$ *and* $V^{AB\prime}(u|u) < V^{B\prime}(u)$ *for all* $u < u_1$*, and if* $\hat{u}_1 = 0$*,* $V^{AB\prime}(0|0) \geq V^{B\prime}(0)$.

(b) $\mathcal{V}'(u) \geq -1$ *for all* $u \geq 0$.

(c) $\mathcal{V}$ *is concave.*

B.3.1 Proof of Lemmas

Proof of Lemma 7. (a) Let $\Gamma_B(T)$ denote a execution-only contract with the deadline T. The agent's expected payoff is

$$U_0(\Gamma_B(T)) = \int_0^T R_\tau \lambda_B e^{-\lambda_B \tau} d\tau = \int_0^T \phi \left(T - \tau + \frac{1}{\lambda_B} \right) \lambda_B e^{-\lambda_B \tau} d\tau$$
$$= -\phi (T - \tau) e^{-\lambda_B \tau} \Big|_0^T = \phi T.$$

Therefore, $U_0(\Gamma_B(u/\phi)) = u$.

Also note that the sum of the expected payoffs of the principal and the agent should equal to the expected surplus from the execution-only policy with a deadline of T:

$$P_0(\Gamma_B(T)) + U_0(\Gamma_B(T)) = \mathcal{W}^B(T).$$

Therefore,

$$P_0 \left(\Gamma_B \left(u/\phi \right) \right) = \mathcal{W}^B \left(u/\phi \right) - u = V^B(u).$$

(b) Let $\Gamma_{AB}(T_1, T)$ denote a contract with a switch from the training path to the execution path at T_1 and the deadline T. The subcontract at time $t \leq T_1$ is denoted by $\hat{\Gamma}_{AB}(t|T_1, T)$. Then, the agent's expected payoff for the subcontract $\hat{\Gamma}_{AB}(t|T_1, T)$ at time t is

$$U_t(\hat{\Gamma}_{AB}(t|T_1,T)) = \int_t^{T+\frac{1}{\lambda_A}} \phi\left(T+\frac{1}{\lambda_A}-\tau+\frac{1}{\lambda_A}\right)\lambda_A e^{-\lambda_A(\tau-t)}d\tau$$
$$= -\phi\left(T+\frac{1}{\lambda_A}-\tau\right)e^{-\lambda_A(\tau-t)}\Big|_t^{T+\frac{1}{\lambda_A}} = \phi\left(T+\frac{1}{\lambda_A}-t\right).$$

Also note that

$$\int_0^{T_1} U_\tau(\hat{\Gamma}_{AB}(\tau|T_1,T))\lambda_A e^{-\lambda_A\tau}d\tau_s = \int_0^{T_1}\phi(T+\frac{1}{\lambda_A}-\tau)\lambda_A e^{-\lambda_A\tau}d\tau$$
$$= -\phi(T-\tau)e^{-\lambda_A\tau}\Big|_0^{T_1} = \phi T - \phi(T-T_1)e^{-\lambda_A T_1}.$$

Then, the agent's expected payoff at time 0 is

$$U_0(\Gamma_{AB}(T_1,T)) = \int_0^{T_1} U_\tau(\hat{\Gamma}_{AB}(\tau|T_1,T))\lambda_A e^{-\lambda_A\tau}d\tau$$
$$+ e^{-\lambda_A T_1}\int_{T_1}^{T}\phi(T+\frac{1}{\lambda_B}-\tau)\lambda_B e^{-\lambda_B(\tau-T_1)}d\tau$$
$$=\phi T - \phi(T-T_1)e^{-\lambda_A T_1} - e^{-\lambda_A T_1}\left[\phi\left(T-\tau\right)e^{-\lambda_B(\tau-T_1)}\Big|_{T_1}^{T}\right] = \phi T.$$

Thus, $U_0(\Gamma_{AB}(T_1,u/\phi)) = u$.

As in the previous case, the sum of the expected payoffs of the principal and the agent is equal to the one-switch policy with the intermediate deadline T_1, the deadline T, and the extension $1/\lambda_A$:

$$P_0(\Gamma_{AB}(T_1,T)) + U_0(\Gamma_{AB}(T_1,T)) = \mathcal{W}^{AB}_{T-T_1}(T,1/\lambda_A).$$

By plugging in $T = u/\phi$ and $T_1 = (u-u_1)/\phi$, (B.5) holds.

(c) Note that a training-only contract with a deadline T is equivalent to a contract with a switch from the training path to the execution path at $T_1 = T$ and a deadline

T. Therefore, by the previous result, a training-only contract with the deadline u/ϕ implements $(V^{AB}(u|0), u)$.

(d) By the construction of $\mathcal{W}^B$,

$$\dot{\mathcal{W}}^B(T) = \lambda_B(\Pi - \mathcal{W}^B(T)) - c, \tag{B.9}$$

for all $T \geq 0$. Similarly,

$$\dot{\mathcal{W}}^{AB}_{\hat{T}}(T, 1/\lambda_A) = \lambda_A(W^1_{T+1/\lambda_A} - \mathcal{W}^{AB}_{\hat{T}}(T, 1/\lambda_A)) - c, \tag{B.10}$$

for all $T \geq \hat{T}$.

Using the definitions of V^B, V^{AB} and V_M, (B.6) and (B.7) follow.

$\square$

Proof of Lemma 8. (a) Suppose that $\hat{u}_1 > 0$, which implies that $1/\lambda_A < \overline{\Delta}$. Now set $\Delta = 1/\lambda_A$. In Lemma 3, $\hat{T}$ is chosen to satisfy $\lambda_B(\Pi - \mathcal{W}^B(\hat{T})) = \lambda_A(W^1_{\hat{T}+1/\lambda_A} - \mathcal{W}^B(\hat{T}))$ and $\lambda_B(\Pi - \mathcal{W}^B(z)) > \lambda_A(W^1_{z+1/\lambda_A} - \mathcal{W}^B(z))$ for all $z < \hat{T}$.

Using Lemma 7 (d) and $\mathcal{W}^B(\hat{T}) = \mathcal{W}^{AB}_{\hat{T}}(\hat{T}, 1/\lambda_A)$, we can derive that $V^{AB\prime}(\hat{u}_1|\hat{u}_1) = V^{B\prime}(\hat{u}_1)$ and $V^{AB\prime}(u|u) < V^{B\prime}(u)$ for all $u < \hat{u}_1$.

When $\hat{u}_1 = 0$, and thereby $1/\lambda_A \geq \overline{\Delta}$, it follows from (A.9) that $V^{AB\prime}(\hat{u}_1|\hat{u}_1) \geq V^{B\prime}(\hat{u}_1)$.

(b) Since a longer the deadline increases the expected social surplus, $\mathcal{W}^*(T, \Delta)$ is increasing in T. Therefore,

$$\mathcal{V}'(u) = \frac{\dot{\mathcal{W}}^*(u/\phi, 1/\lambda_A)}{\phi} - 1 \geq -1.$$

(c) If $u \leq \hat{u}_1$,

$$\mathcal{V}''(u) = V^{B\prime\prime}(u) = -\left(\Pi - \frac{c}{\lambda_B}\right)\frac{\lambda_B^2}{\phi^2} e^{-\frac{\lambda_B}{\phi}u} < 0.$$

Now assume that $u > \hat{u}_1$. By using (B.9), (B.10) and $V^{B\prime}(\hat{u}_1) \leq V^{AB\prime}(\hat{u}_1|\hat{u}_1)$, we can derive that

$$\mathcal{W}^B(\hat{u}_1/\phi) \leq \Pi - \frac{c}{\lambda_A - \lambda_B} - \frac{\lambda_A}{\lambda_A - \lambda_B}\left(\Pi - \frac{c}{\lambda_A}\right) e^{-\frac{\lambda_A}{\phi}\hat{u}_1 - 1}. \tag{B.11}$$

Using (A.8), we can derive the followings though some algebra:

$$\begin{aligned} V^{AB\prime\prime}(u|\hat{u}_1) = &\left(\frac{\lambda_A}{\phi}\right)^2 e^{\frac{\lambda_A}{\phi}(\hat{u}_1 - u)} \left[\mathcal{W}^B(\tfrac{\hat{u}_1}{\phi}) - \left(\Pi - \frac{2c}{\lambda_A}\right) + 2\left(\Pi - \frac{c}{\lambda_A}\right) e^{-\frac{\lambda_A}{\phi}\hat{u}_1 - 1}\right] \\ &- \left(\Pi - \frac{c}{\lambda_A}\right)\left(\frac{\lambda_A}{\phi}\right)^3 (u - \hat{u}_1) e^{-\frac{\lambda_A}{\phi}(u + \frac{1}{\lambda_A})} \end{aligned}$$

By plugging (B.11) in, we have

$$\begin{aligned} V^{AB\prime\prime}(u|\hat{u}_1) \leq &\left(\frac{\lambda_A}{\phi}\right)^2 e^{\frac{\lambda_A}{\phi}(\hat{u}_1 - u)} \left[\frac{\lambda_A - 2\lambda_B}{\lambda_A(\lambda_A - \lambda_B)} c + \frac{\lambda_A - 2\lambda_B}{\lambda_A - \lambda_B}\left(\Pi - \frac{c}{\lambda_A}\right) e^{-\frac{\lambda_A}{\phi}\hat{u}_1 - 1}\right] \\ &- \left(\Pi - \frac{c}{\lambda_A}\right)\left(\frac{\lambda_A}{\phi}\right)^3 (u - \hat{u}_1) e^{-\frac{\lambda_A}{\phi}(u + \frac{1}{\lambda_A})}. \end{aligned} \tag{B.12}$$

Then, from $\lambda_A = 2\lambda_B$, $\mathcal{V}''(u) = V^{AB\prime\prime}(u|\hat{u}_1) \leq 0$ for all $u \geq \hat{u}_1$.

□

B.4 Value Function Verification (Proposition 1)

The goal of this subsection is to prove Proposition 1. Specifically, we show that the value function defined in the previous section solves (HJB_V) subject to (IC). To achieve this, we introduce functions that specify potential deviations and then establish useful properties as a lemma, followed by the proof for Proposition 1.

First, define

$$L^B(u, R) \equiv \mathcal{J}(\mathcal{V}(\cdot), R, \cdot, 1) = \lambda_B(\Pi - R - \mathcal{V}(u)) - c - \lambda_B(R - u)\mathcal{V}'(u). \tag{B.13}$$

Given u, maximizing this function with respect to $R \geq u + \phi/\lambda_B$ is equivalent to maximizing the right hand side of (HJB$_V$) under (IC) with $a = 1$.

Similarly, define

$$L^A(u, w) \equiv \mathcal{J}(\mathcal{V}(\cdot), \cdot, w, 0) = \lambda_A(V_M(w) - \mathcal{V}(u)) - c - \lambda_A(w - u)\mathcal{V}'(u). \tag{B.14}$$

Given u, maximizing this function with respect to $w \geq u + \phi/\lambda_A$ is equivalent to maximizing the right hand side of (HJB$_V$) under (IC) with $a = 0$.

Lemma 9. *Suppose that $\Pi > c/\lambda_B$ and $\lambda_A = 2\lambda_B$. Then, for any $u \geq 0$, $L^B(u, R) \leq 0$ for all $R \geq u + \phi/\lambda_B$, and $L^A(u, w) \leq 0$ for all $w \geq u + \phi/\lambda_A$.*

Proof of Lemma 9. We begin by showing $L^B(u, R) \leq 0$ for all $R \geq u + \phi/\lambda_B$. Observe that

$$\frac{\partial L^B}{\partial R} = -\lambda_B(1 + \mathcal{V}'(u)) \leq 0.$$

from Lemma 8 (b). Also note that

$$\begin{aligned} L^B(u, u + \phi/\lambda_B) =& \lambda_B(\Pi - u - \mathcal{V}(u)) - c - \phi(\mathcal{V}'(u) + 1) \\ =& \lambda_B(\Pi - \mathcal{W}^*(u/\phi, 1/\lambda_A)) - c - \dot{\mathcal{W}}^*(u/\phi, 1/\lambda_A) \leq 0. \end{aligned}$$

The last inequality is due to (HJB$_W$). Therefore, $L^B(u, R) \leq 0$ for all $R \geq u + \phi/\lambda_B$.

For L^A, observe that

$$\frac{\partial L^A}{\partial u} = -\lambda_A(w - u)\mathcal{V}''(u) \geq 0$$

by the concavity of $\mathcal{V}$. Therefore, it is sufficient to check whether $L^A(u, u + \phi/\lambda_A) \leq 0$ for all $u \geq 0$.

Note that for all $u \geq \hat{u}_1$, $L^A(u, u + \phi/\lambda_A) = 0$ holds by (B.7). Now, suppose that $u < \hat{u}_1$, thereby $\mathcal{V}(u) = V^B(u)$. Using (B.7) and Lemma 8 (a), we have $L^A(u, u + \phi/\lambda_A) = \phi(V^{AB\prime}(u|u) - V^{B\prime}(u)) < 0$ for all $u < \hat{u}_1$. □

Now We prove Proposition 1.

Proof of Proposition 1. We begin by showing that $\mathcal{J}$ becomes zero when the value function $\mathcal{V}$ is utilized alongside contractual terms with binding ICs.

When $u \geq \hat{u}_1$, $\mathcal{V}(u) = V^{AB}(u|\hat{u}_1)$. Then, by (B.7),

$$\mathcal{J}\left(V^{AB}(u|\hat{u}_1),\ \cdot,\ u + \phi/\lambda_A,\ 0\right) = 0.$$

Likewise, when $u < \hat{u}_1$, $\mathcal{V}(u) = V^B(u|)$, and by (B.6),

$$\mathcal{J}\left(V^B(u),\ u + \phi/\lambda_B,\ \cdot,,\ 1\right) = 0.$$

Lemma 9 shows that $\mathcal{J}$ is nonpositive for any feasible deviations. Therefore, $\mathcal{V}$ solves (HJB_V) subject to (IC).

The concavity of $\mathcal{V}$ is shown in Lemma 8 (c). If $\hat{u}_1 = 0$, $\mathcal{V}(u) = V^{AB}(u|0)$ is differentiable for all $u \geq 0$. If $\hat{u}_1 > 0$, $V^B(u)$ is differentiable for all $u < \hat{u}_1$ and $V^{AB}(u|\hat{u}_1)$ is differentiable for all $u > \hat{u}$. By Lemma 8 (a), $\mathcal{V}$ is differentiable at $\hat{u}_1$ as well. Also note that $\mathcal{V}(0) = 0$. Therefore, by Lemma 6, for any incentive compatible contract promising the agent u units of utility, the principal's expected payoff is lower than or equal to $\mathcal{V}(u)$.

Last, by Lemma 7, there exists a contract implementing $(\mathcal{V}(u), u)$. Therefore, $\mathcal{V}(u)$ is the principal's maximized expected payoff subject to the promise-keeping constraint $U_0(\Gamma) = u$ and the incentive compatibility constraints. $\square$

B.5 Thresholds

In this section, we explain how to pin down the thresholds π_B and π_A and provide some properties of them.

First, recall that π_A is the solution of $1/\lambda_A = \overline{\Delta}$. This gives us

$$\frac{1}{\lambda_A} = \frac{1}{\lambda_A} \log\left[\frac{2\pi - 1}{\pi - 1}\right] \qquad \Leftrightarrow \qquad \pi_A = \frac{e-1}{e-2} \approx 2.392. \tag{B.15}$$

Also recall that the threshold is relevant to whether the switching point $\hat{u}_1$ is greater than $\overline{u}$ or not. Since $\mathcal{V}$ is concave, $\hat{u}_1 \leq \overline{u}$ if and only if $\mathcal{V}'(\hat{u}_1) \geq 0$. Observe that by using the formula of V^B and (A.11), we can derive that

$$\mathcal{V}'(\hat{u}_1) = V^{B\prime}(\hat{u}_1) = \frac{(\lambda_B \Pi - c)^2}{\phi \cdot (\lambda_A \Pi - c)} e - 1.$$

By solving the equation making the above equal to zero, it follows that $\mathcal{V}'(\hat{u}_1) \geq 0$ if and only if

$$\pi \leq \pi_B \equiv \frac{c \cdot e + \phi + \sqrt{\phi(c \cdot e + \phi)}}{c \cdot e}. \tag{B.16}$$

We conclude the section by showing that π_B lies between π_F and π_A.

Lemma 10. *When* $\lambda_A = 2\lambda_B$, $\pi_B \in (\pi_F, \pi_A)$.

Proof of Lemma 10. From $\phi \leq c$, we have

$$\pi_B \leq \frac{(e+1) + \sqrt{e+1}}{e} \approx 2.077.$$

Therefore, by (B.15), $\pi_B < \pi_A$.

Next, observe that

$$\pi_B - \pi_F = \frac{-\phi(e-1) + \sqrt{\phi(c \cdot e + \phi)}}{c \cdot e} \geq \frac{\phi(\sqrt{e+1} - (e-1))}{c \cdot e}.$$

Since $\sqrt{e+1} - (e-1) \approx .21 > 0$, $\pi_B > \pi_F$. $\square$

Online Appendix for

"Execution vs. Training under Endogenous Deadlines"

Yonggyun Kim Curtis Taylor

In this online appendix, we provide the proofs for Section 5, specifically the value function and optimal contract characterizations when there is an efficiency loss from training (Proposition 4, Proposition 2, and Proposition 3). Some results in Appendix B can still be utilized (e.g., Lemma 5, Lemma 6, and Lemma 7) as they do not use the parametric assumption $\lambda_A = 2\lambda_B$.

OA.1 Value Function Characterization

We begin by specifying a value function that can be implemented by a two-switch contract defined in Definition 3.

Lemma OA.1. *The following statements hold.*

(a) When $0 < u_1 < u_2 < u$, a two-switch contract with the intermediate deadlines $(u-u_2)/\phi$, $(u-u_1)/\phi$ and the final deadline u/ϕ implements $(V^{BAB}(u|u_1,u_2),u)$ where

$$V^{BAB}(u|u_1,u_2) \equiv \left(\Pi - \frac{c}{\lambda_B}\right)\left(1 - e^{\frac{\lambda_B}{\phi}(u_2-u)}\right) + (V^{AB}(u_2|u_1)+u_2)e^{\frac{\lambda_B}{\phi}(u_2-u)} - u. \tag{OA.1.1}$$

(b) The following differential equation holds:

$$\phi V^{BAB\prime}(u|u_1,u_2) = \lambda_B\left(\Pi - \frac{\phi}{\lambda_B} - u - V^{BAB}(u|u_1,u_2)\right) - c. \tag{OA.1.2}$$

Proof of Lemma OA.1. (a) Let $\Gamma_{BAB}(T_1, T_2, T)$ denote a contract with two switches at $T_1 = (u - u_2)/\phi$ and $T_2 = (u - u_1)/\phi$ and a deadline $T = u/\phi$. Note that at time T_1 (if the problem has not been solved), the remaining contract is equivalent to $\Gamma_{AB}(T_2 - T_1, T - T_1)$. Recall that $U_0(\Gamma_{AB}(T_2 - T_1, T - T_1)) = \phi(T - T_1)$. Then, the agent's expected payoff at time 0 is

$$\begin{aligned}
U_0(\Gamma_{BAB}(T_1, T_2, T)) =& \int_0^{T_1} \phi(T + 1/\lambda_B - \tau)\lambda_B e^{-\lambda_B \tau} d\tau \\
&+ e^{-\lambda_A T_1} U_0(\Gamma_{AB}(T_2 - T_1, T - T_1)) \\
=& \phi T - \phi(T - T_1)e^{-\lambda_B T_1} + e^{-\lambda_B T_1}\phi(T - T_1) = \phi T = u.
\end{aligned}$$

Also note that

$$\begin{aligned}
&P_0(\Gamma_{BAB}(T_1, T_2, T)) + U_0(\Gamma_{BAB}(T_1, T_2, T)) \\
=& \int_0^{T_1} (\Pi - c\tau)\lambda_B e^{-\lambda_B \tau} d\tau - cT_1 e^{-\lambda_B T_1} \\
&+ e^{-\lambda_B T_1}(P_0(\Gamma_{AB}(T_2 - T_1, T - T_1)) + U_0(\Gamma_{AB}(T_2 - T_1, T - T_1))).
\end{aligned}$$

Recall that $U_0(\Gamma_{AB}(T_2', T_1')) + P_0(\Gamma_{AB}(T_2', T_1')) = V^{AB}(\phi T_1' | \phi(T_1' - T_2')) + \phi(T_1' - T_2')$. By plugging in $T_1' = T - T_1$, $T_2' = T_2 - T_1$, $\phi T_1 = u - u_2$ and $\phi T_2 = u - u_1$, the right hand side of the above equation is equal to:

$$\left(\Pi - \frac{c}{\lambda_B}\right)(1 - e^{-\lambda_B T_1}) + e^{-\lambda_B T_1} \cdot \left\{V^{AB}(u_2|u_1) + u_2\right\} = V^{BAB}(u) + u,$$

thus, $P_0(\Gamma_{BAB}(T_1, T_2, T)) = V^{BAB}(u|u_1, u_2)$.

(b) Last, by taking the derivative of (OA.1.1) and multiplying by ϕ, we have

$$\begin{aligned}\phi V^{BAB\prime}(u|u_1,u_2) =& \lambda_B\left(\Pi-\frac{c}{\lambda_B}\right)e^{\frac{\lambda_B}{\phi}(u_2-u)}-\lambda_B\left(V^{AB}(u_2|u_1)+u_2\right)e^{\frac{\lambda_B}{\phi}(u_2-u)}-\phi\\ =& \lambda_B\left(\Pi-\frac{\phi}{\lambda_B}-u-V^{BAB}(u|u_1,u_2)\right)-c,\end{aligned}$$

thus, (OA.1.2) holds.

□

Based on the intuition presented in the main text, we conjecture the value function defined as follows.

$$\mathcal{V}(u)=\begin{cases}V^B(u), & \text{if } 0\le u\le \hat{u}_1,\\ V^{AB}(u|\hat{u}_1), & \text{if } \hat{u}_1<u\le \hat{u}_2,\\ V^{BAB}(u|\hat{u}_1,\hat{u}_2), & \text{if } \hat{u}_2<u.\end{cases} \tag{OA.1.3}$$

The following proposition shows that there exist $\hat{u}_1$ and $\hat{u}_2$ such that the above three value functions are smoothly pasted, and the resulting function is the principal's value function.

Proposition 4. *Suppose that η is less than 1.*

(a) ***(Smooth Pasting)*** *There exist $\hat{u}_2\ge\hat{u}_1\ge 0$ such that*

i. $V^{B\prime}(u)>V^{AB\prime}(u|u)$ *for all* $0\le u<\hat{u}_1$;

ii. $V^{AB\prime}(\hat{u}_1|\hat{u}_1)\ge V^{B\prime}(\hat{u}_1)$, *and if the equality holds,* $V^{AB\prime\prime}(\hat{u}_1|\hat{u}_1)>V^{B\prime\prime}(\hat{u}_1)$;

iii. $V^{AB\prime}(u|\hat{u}_1)>V^{BAB\prime}(u|\hat{u}_1,u)$ *for all* $\hat{u}_1<u<\hat{u}_2$;

iv. $V^{BAB\prime}(\hat{u}_2|\hat{u}_1,\hat{u}_2)=V^{AB\prime}(\hat{u}_2|\hat{u}_1)$ *and* $V^{BAB\prime\prime}(\hat{u}_2|\hat{u}_1,\hat{u}_2)>V^{AB\prime\prime}(\hat{u}_2|\hat{u}_1)$;

v. $V^{BAB\prime}(u|\hat{u}_1,\hat{u}_2)>\frac{1}{\phi}\left[\lambda_A(V_M(u+\phi/\lambda_A)-V^{BAB}(u|\hat{u}_1,\hat{u}_2))-c\right]$ *for all* $u>\hat{u}_2$.

(b) ***(Large Loss)*** *If* $\eta\le\frac{1}{e-1}$, *there exists* $\tilde{\pi}_M(\eta)$ *such that*

i. $\hat{u}_2 = \hat{u}_1 = 0$ *if* $\tilde{\pi}_M(\eta) \geq \pi > \frac{c}{\lambda_B}$*;*

ii. $\hat{u}_2 > \hat{u}_1 > 0$ *if* $\pi > \tilde{\pi}_M(\eta)$*.*

(c) ***(Small Loss)*** *If* $\frac{1}{e-1} < \eta < 1$*, there exist* $\tilde{\pi}_A(\eta) > \tilde{\pi}_B(\eta)$ *such that*

i. $\hat{u}_2 = \hat{u}_1 = 0$ *if* $\tilde{\pi}_B(\eta) \geq \pi > 1$*;*

ii. $\hat{u}_2 > \hat{u}_1 > 0$ *if* $\tilde{\pi}_A(\eta) \geq \pi > \tilde{\pi}_B(\eta)$*;*

iii. $\hat{u}_2 > \hat{u}_1 = 0$ *if* $\pi \geq \tilde{\pi}_B(\eta)$*.*

(d) The function $\mathcal{V}$ *defined in* (OA.1.3)*, with* $\hat{u}_1$ *and* $\hat{u}_2$ *derived in (a), serves as the principal's value function.*

OA.1.1 Proof of Proposition 4

We begin by identifying which approach will be chosen at the deadline. Note that the execution path is chosen at the deadline if and only if $V^{B\prime}(0) > V^{AB\prime}(0|0)$. The following lemma provides the parametric condition for this.

Lemma OA.2. *If* $\eta \leq 1/(e-1)$*, the inequality* $V^{B\prime}(0) > V^{AB\prime}(0|0)$ *always holds. If* $\eta > 1/(e-1)$*,* $V^{B\prime}(0) > V^{AB\prime}(0|0)$ *is equivalent to*

$$\pi < \tilde{\pi}_A(\eta) \equiv \frac{e-1}{\eta(e-1)-1}.$$

Moreover, if $\pi = \tilde{\pi}_A(\eta)$*, then* $V^{B\prime}(0) = V^{AB\prime}(0|0)$ *and* $V^{B\prime\prime}(0) < V^{AB\prime\prime}(0|0)$*.*

Proof of Lemma OA.2. By $V^B(0) = V^{AB}(0|0) = 0$ and Lemma 7 (d), we have

$$\phi V^{B\prime}(0) = \lambda_B \Pi - \phi - c,$$
$$\phi V^{AB\prime}(0|0) = \lambda_A V_M(\frac{\phi}{\lambda_A}) - c = \lambda_A \left(\Pi - \frac{c}{\lambda_A}\right)(1 - e^{-1}) - \phi - c.$$

Therefore, $V^{B\prime}(0) > V^{AB\prime}(0|0)$ is equivalent to:

$$(\eta(e-1)-1)\,\lambda_B\Pi < c(e-1).$$

Therefore, when $\eta \leq \frac{1}{e-1}$, $V^{B\prime}(0) > V^{AB\prime}(0|0)$ always holds, and when $\eta > \frac{1}{e-1}$, $V^{B\prime}(0) > V^{AB\prime}(0|0)$ is equivalent to $\pi < \tilde{\pi}_A(\eta)$.

Next, assume that $\eta > 1/(e-1)$ and $\pi = \tilde{\pi}_A(\eta)$. With some algebra, it follows that

$$\phi^2 V^{AB\prime\prime}(0|0) - \phi^2 V^{B\prime\prime}(0) = \lambda_B c\left[\frac{(e-1)\eta^2}{(e-1)\eta - 1}\right].$$

The right hand side is positive from $\eta > 1/(e-1)$, thus, $V^{AB\prime\prime}(0|0) > V^{B\prime\prime}(0)$. □

Next, we establish a condition under which the execution route is always employed. When this condition does not hold, a switch from execution to training occurs. We show the existence of the switching point $\hat{u}_1$.

Lemma OA.3. *There exists $\tilde{\pi}_M(\eta) \geq 2\lambda_B/\lambda_A = 2/(\eta+1)$ with $\tilde{\pi}_M(1) = 1$ such that the following statements hold.*

(a) If $1 \leq \pi < \tilde{\pi}_M(\eta)$, $V^{B\prime}(u) > V^{AB\prime}(u|u)$ for all $u \geq 0$.

(b) Suppose that one of the following statements hold: (i) $\eta \leq 1/(e-1)$ and $\pi > \tilde{\pi}_M(\eta)$; (ii) $\eta > 1/(e-1)$ and $\tilde{\pi}_A(\eta) \geq \pi > \tilde{\pi}_M(\eta)$. Then, there exists $\hat{u}_1 > 0$ such that $V^{B\prime}(\hat{u}_1) = V^{AB\prime}(\hat{u}_1|\hat{u}_1)$, $V^{B\prime\prime}(\hat{u}_1) < V^{AB\prime\prime}(\hat{u}_1|\hat{u}_1)$ and $V^{B\prime}(u) > V^{AB\prime}(u|u)$ for all $u \in [0, \hat{u}_1)$;

Proof of Lemma OA.3. Define $x \equiv e^{-\frac{\lambda_B}{\phi}u}$. Using (B.6) and (B.7), $\phi V^{AB\prime}(u|u) - \phi V^{B\prime}(u)$ can be expressed in the form of $H_1(x)$ with $\Delta = 1/\lambda_A$, as defined in (A.10) in the proof of Lemma 3. We also consider this as a function of π. With the definition of η, it can be rewritten as follows:

$$\tilde{H}_1(x;\pi) \equiv -\{(\eta+1)\pi - 1\} \cdot c \cdot e^{-1} x^{\eta+1} + \eta(\pi - 1) \cdot c \cdot x - (1-\eta) \cdot c. \tag{OA.1.4}$$

Observe that

$$\frac{\partial^2 \tilde{H}_1}{\partial x^2}(x;\pi) = -(\eta+1)\eta\{(\eta+1)\pi - 1\} \cdot c \cdot e^{-1} x^{\eta-1},$$

thus $\tilde{H}_1$ is a strict concave function in x when $\pi \geq 1$. Let $x^*(\pi)$ be the solution of $\max_x H_1(x;\pi)$ subject to $0 \leq x \leq 1$. Then, when $\pi \geq 1$, from the first order condition, we can derive that

$$x^*(\Pi) = \left[\frac{\eta(\pi-1)}{(\eta+1)\{(\eta+1)\pi - 1\}e^{-1}}\right]^{\frac{1}{\eta}}. \tag{OA.1.5}$$[16]

Now define

$$h(\pi) \equiv \tilde{H}_1(x^*(\pi);\pi) = K\left(\frac{\pi - 1}{(\eta+1)\pi - 1}\right)^{\frac{1}{\eta}}(\pi - 1)\cdot c - (1-\eta)\cdot c$$

where $K = \frac{\eta^2}{\eta+1}\left(\frac{\eta e}{\eta+1}\right)^{\frac{1}{\eta}}$. Observe that

$$h\left(\frac{2}{\eta+1}\right) = (1-\eta)c\left[\frac{\eta^2}{(\eta+1)^2}\left(\frac{\eta(1-\eta)e}{(\eta+1)^2}\right)^{\frac{1}{\eta}} - 1\right] < 0$$

from $\eta < 1$ and $\eta(1-\eta)e \leq \frac{e}{4} < 1 \leq (\eta+1)^2$. In addition, $\lim_{\pi\to\infty} h(\pi) = \infty$ and

$$h'(\pi) = \frac{K\pi(1+\eta)c}{\pi - 1} \cdot \left(\frac{\pi - 1}{(\eta+1)\pi - 1}\right)^{\frac{1}{\eta}+1} > 0.$$

Therefore, there exists a unique π such that $h(\pi) = 0$ and $\pi \geq 2/(\eta+1)$, and we denote this

[16] When $\pi \leq \frac{(\eta+1)e^{-1}-\eta}{(\eta+1)^2 e^{-1}-\eta}$, the solution of the maximization problem $\max_{0\leq x\leq 1} H_1(x;\pi)$ is $x^*(\pi) = 1$. However, we can show that $\frac{(\eta+1)e^{-1}-\eta}{(\eta+1)^2 e^{-1}-\eta} < 1$ for any $0 < \eta$, which implies that we can focus on the interior solution when $\pi \geq 1$.

solution by $\tilde{\pi}_M(\eta)$. Also note that when $\eta = 1$, $h(2/(\eta+1)) = h(1) = 0$ thus $\tilde{\pi}_M(1) = 1$.

(a) Suppose that $1 \leq \pi < \tilde{\pi}_M(\eta)$. We have $0 > h(\pi) = \tilde{H}_1(x^*(\pi);\pi) \geq \tilde{H}_1(x;\pi)$ for all $0 \leq x \leq 1$. It is equivalent to $V^{B'}(u) > V^{AB'}(u|u)$ for all $u \geq 0$ in this case.

(b) First, suppose that $\eta \leq 1/(e-1)$ and $\pi > \tilde{\pi}_M(\eta)$. Then, we have $0 < h(\pi) = \tilde{H}_1(x^*(\pi);\pi)$. In addition, by Lemma OA.2, we have $\tilde{H}_1(1;\pi) = \phi(V^{AB'}(0|0) - V^{B'}(0)) < 0$. Then, by concavity of $\tilde{H}_1$ w.r.t. x and $\frac{\partial \tilde{H}_1}{\partial x}(x^*(\pi);\pi) = 0$, there exists $x_1 \in (x^*(\pi), 1]$ such that $\tilde{H}_1(x_1;\pi) = 0$, $\frac{\partial \tilde{H}_1}{\partial x}(x_1;\pi) < 0$ and $\tilde{H}_1(x;\pi) < 0$ for all $x \in (x_1, 1]$. By defining $\hat{u}_1 \equiv -(\phi/\lambda_B)\log x_1$, the above conditions can be translated into: $V^{B'}(\hat{u}_1) = V^{AB'}(\hat{u}_1|\hat{u}_1)$, $V^{B''}(\hat{u}_1) < V^{AB''}(\hat{u}_1|\hat{u}_1)$ and $V^{B'}(u) > V^{AB'}(u|u)$ for all $u \in [0, \hat{u}_1)$.

Next, suppose that $\eta > 1/(e-1)$. Note that by the definition of $\tilde{\pi}_A(\eta)$, if $\pi \geq \tilde{\pi}_A(\eta)$, $\tilde{H}_1(1;\pi) \geq 0$. It implies that $h(\pi) \geq \tilde{H}_1(1;\pi) \geq 0$ and $\pi \geq \tilde{\pi}_M(\eta)$. Therefore, we can see that $\tilde{\pi}_A(\eta) \geq \tilde{\pi}_M(\eta)$. If $\tilde{\pi}_A(\eta) \geq \pi > \tilde{\pi}_M(\eta)$, we also have $\tilde{H}_1(x^*(\pi);\pi) > 0 > \tilde{H}_1(1;\pi)$. By using the same arguments as above, we can show that there exists $\hat{u}_1 > 0$ such that $V^{B'}(\hat{u}_1) = V^{AB'}(\hat{u}_1|\hat{u}_1)$, $V^{B''}(\hat{u}_1) < V^{AB''}(\hat{u}_1|\hat{u}_1)$ and $V^{B'}(u) > V^{AB'}(u|u)$ for all $u \in [0, \hat{u}_1)$.

□

The following lemma shows that when there is an efficiency loss from training and the training path is employed, there will be an additional switching point, $\hat{u}_2$.

Lemma OA.4. *Suppose that $\eta < 1$, $\Pi > c/\lambda_B$ and one of the followings hold: (i) $V^{B'}(\hat{u}_1) < V^{AB'}(\hat{u}_1|\hat{u}_1)$; (ii) $V^{B'}(\hat{u}_1) = V^{AB'}(\hat{u}_1|\hat{u}_1)$ and $V^{B''}(\hat{u}_1) < V^{AB''}(\hat{u}_1|\hat{u}_1)$. Then, there exists $\hat{u}_2 > \hat{u}_1$ such that $V^{AB'}(\hat{u}_2|\hat{u}_1) = V^{BAB'}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ and $V^{AB''}(\hat{u}_2|\hat{u}_1) < V^{BAB''}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ and such $\hat{u}_2$ is unique. Moreover, $V^{AB'}(u|\hat{u}_1) > V^{BAB'}(u|\hat{u}_1, u)$ for all $u \in (\hat{u}_1, \hat{u}_2)$.*

Proof of Lemma OA.4. Using (OA.1.2), (B.7) and $V^{BAB}(u|\hat{u}_1, u) = V^{AB}(u|\hat{u}_1)$, $\phi V^{BAB\prime}(u|\hat{u}_1, u) - \phi V^{AB\prime}(u|\hat{u}_1)$ can be rewritten as follows:

$$\lambda_B \Pi - \lambda_A \left\{ V_M(u + \phi/\lambda_A) + u + \phi/\lambda_A \right\} + (\lambda_A - \lambda_B)(V^{AB}(u|\hat{u}_1) + u).$$

By performing a similar derivation as in (A.12) and using $\eta = \frac{\lambda_A}{\lambda_B} - 1$ and $y \equiv e^{-\frac{\lambda_A}{\phi}(u - \hat{u}_1)}$, the above expression can be further rewritten as follows:

$$\begin{aligned} \tilde{H}_2(y) \equiv\; & \frac{1-\eta}{1+\eta} c + (\lambda_A \Pi - c) e^{-1 - \frac{\lambda_A}{\phi}\hat{u}_1} \left[1 + \frac{\eta}{1+\eta} \log y \right] y \\ & + \eta \left[\frac{1-\eta}{1+\eta} c - (\lambda_B \Pi - c) e^{-\frac{\lambda_B}{\phi}\hat{u}_1} \right] y. \end{aligned} \tag{OA.1.6}$$

Note that $\tilde{H}_2(1) = \phi V^{BAB\prime}(\hat{u}_1|\hat{u}_1, \hat{u}_1) - \phi V^{AB\prime}(\hat{u}_1|\hat{u}_1) = \phi V^{B\prime}(\hat{u}_1) - \phi V^{AB\prime}(\hat{u}_1|\hat{u}_1) \leq 0$ by assumption. By differentiating $\tilde{H}_2$ twice, we have

$$\tilde{H}_2''(y) = \frac{\eta}{1+\eta} (\lambda_A \Pi - c) e^{-1 - \frac{\lambda_A}{\phi}\hat{u}_1} \frac{1}{y} > 0.$$

Since $\Pi > \frac{c}{\lambda_B} > \frac{c}{\lambda_A}$, $\tilde{H}_2$ is strictly convex in y. Also note that

$$\lim_{y \to 0} \tilde{H}_2(y) = \frac{1-\eta}{1+\eta} c > 0.$$

By the convexity of $\tilde{H}_2$, there exists $y_2 \in (0,1)$ such that (i) $\tilde{H}_2(y) < 0$ for all $y \in (y_2, 1)$, (ii) $\tilde{H}_2(y_2) = 0$, and (iii) $\tilde{H}_2'(y_2) < 0$. Let $\hat{u}_2 = \hat{u}_1 - \frac{\phi}{\lambda_A} \log y_2$. Then, from (i) and (ii), we have $V^{AB\prime}(u|\hat{u}_1) > V^{BAB\prime}(u|\hat{u}_1, u)$ for all $u \in (\hat{u}_1, \hat{u}_2)$ and $V^{AB\prime}(\hat{u}_2|\hat{u}_1) = V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$. Additionally, since y is decreasing in u, $\tilde{H}_2'(y_2) < 0$ implies that $V^{AB\prime\prime}(\hat{u}_2|\hat{u}_1) < V^{BAB\prime\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$. $\square$

Next, when there is a switching point $\hat{u}_2$, the following lemma shows that the execution path is employed for all $u > \hat{u}_2$.

Lemma OA.5. *Suppose that* $\pi > 1$, $V^{AB\prime}(\hat{u}_2|\hat{u}_1) = V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ *and* $V^{AB\prime\prime}(\hat{u}_2|\hat{u}_1) < V^{BAB\prime\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$. *Then,* $\lambda_A \left(V_M(u + \phi/\lambda_A) - V^{BAB}(u|\hat{u}_1, \hat{u}_2)\right) - \phi V^{BAB\prime}(u|\hat{u}_1, \hat{u}_2) - c < 0$ *for all* $u > \hat{u}_2$.

Proof of Lemma OA.5. By differentiating (B.6) and (B.7), we have

$$\begin{aligned}\phi V^{AB\prime\prime}(u|\hat{u}_1) &= \lambda_A \left(V_M'\left(u + \phi/\lambda_A\right) + 1\right) - \lambda_A \left(V^{AB\prime}(u|\hat{u}_1) + 1\right), \\ \phi V^{BAB\prime\prime}(u|\hat{u}_1, \hat{u}_2) &= -\lambda_B \left(V^{BAB\prime}(u|\hat{u}_1, \hat{u}_2) + 1\right).\end{aligned}$$

Then, $V^{AB\prime}(\hat{u}_2|\hat{u}_1) = V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ and $V^{AB\prime\prime}(\hat{u}_2|\hat{u}_1) < V^{BAB\prime\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ imply that

$$\begin{aligned}&(\lambda_A - \lambda_B)(1 + V^{AB\prime}(\hat{u}_2|\hat{u}_1)) > \lambda_A(V_M'(\hat{u}_2 + \phi/\lambda_A) + 1) \\ \iff \quad &\eta(1 + V^{AB\prime}(\hat{u}_2|\hat{u}_1)) > (\eta + 1)\left(\frac{\lambda_A \Pi - c}{\phi}\right) e^{-\frac{\lambda_A}{\phi}\hat{u}_2 - 1}. \qquad \text{(OA.1.7)}\end{aligned}$$

Define a function $H_3 : [\hat{u}_2, \infty) \to \mathbb{R}$ as

$$H_3(u) \equiv \lambda_A \left[V_M(u + \phi/\lambda_A) - V^{BAB}(u|\hat{u}_1, \hat{u}_2)\right] - \phi V^{BAB\prime}(u|\hat{u}_1, \hat{u}_2) - c.$$

With some algebra, we can derive that

$$H_3(u) = (\eta - 1)\, c - (\lambda_A \Pi - c) e^{-\frac{\lambda_A}{\phi}\hat{u}_2 - 1} \cdot e^{\frac{\lambda_A}{\phi}(\hat{u}_2 - u)} + \eta\phi \left(V^{AB\prime}(\hat{u}_2|\hat{u}_1) + 1\right) e^{\frac{\lambda_B}{\phi}(\hat{u}_2 - u)}.$$

Also note that

$$\begin{aligned}H_3(\hat{u}_2) &= \lambda_A \left[V_M(\hat{u}_2 + \phi/\lambda_A) - V^{AB}(\hat{u}_2|\hat{u}_1)\right] - c - \phi V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2) \\ &= \phi V^{AB\prime}(\hat{u}_2|\hat{u}_1) - \phi V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2) = 0.\end{aligned}$$

Define $x \equiv e^{\frac{\lambda_B}{\phi}(\hat{u}_2-u)}$. Then, $H_3(u)$ can be rewritten as follows:

$$\tilde{H}_3(x) = (\eta-1)c - (\lambda_A \Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_2-1}x^{\eta+1} + \eta\phi\left(V^{AB'}(\hat{u}_2|\hat{u}_1)+1\right)x$$

and $\tilde{H}_3(1) = H_3(\hat{u}_2) = 0$.

Note that

$$\tilde{H}_3'(x) = -(\eta+1)(\lambda_A \Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_2-1}x^{\eta} + \eta\phi\left(V^{AB'}(\hat{u}_2|\hat{u}_1)+1\right).$$

By (OA.1.7), we can derive that

$$\tilde{H}_3'(1) = -(\eta+1)(\lambda_A \Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_2-1} + \eta\phi\left(V^{AB'}(\hat{u}_2|\hat{u}_1)+1\right) > 0.$$

Also note that

$$\tilde{H}_3''(x) = -(\eta+1)\eta(\lambda_A \Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_2-1}x^{\eta-1} < 0.$$

Therefore, $\tilde{H}_3'(x) > 0$ for all $0 < x < 1$. Since $\tilde{H}_3(1) = 0$, $\tilde{H}_3(x) < 0$ for all $x \in (0,1)$. Thus, $\lambda_A\left(V_M(u+\phi/\lambda_A) - V^{BAB}(u|\hat{u}_1,\hat{u}_2)\right) - \phi V^{BAB'}(u|\hat{u}_1,\hat{u}_2) - c < 0$ for all $u \geq \hat{u}_2$. $\square$

Lastly, we show that the resulting value function is concave, and that L^B and L^A—the functions specifying deviations, defined in (B.13) and (B.14)—are nonpositive.

Lemma OA.6. *Suppose that $\pi > 1$ and $\eta < 1$.*

(a) $\mathcal{V}$ is concave;

(b) for any $u \geq 0$, $L^B(u,R) \leq 0$ for all $R \geq u+\phi/\lambda_B$, and $L^A(u,w) \leq 0$ for all $w \geq u+\phi/\lambda_A$.

Proof of Lemma OA.6. (a) When $u < \hat{u}_1$, $\mathcal{V}''(u) = V^{B''}(u) < 0$ from Lemma 8 (c).

When $\hat{u}_1 < u < \hat{u}_2$, the inequality (B.12) is still applicable, and from $\lambda_A < 2\lambda_B$, we have $\mathcal{V}''(u) = V^{AB''}(u|\hat{u}_1) \leq 0$.

When $u > \hat{u}_2$, by differentiating (OA.1.1) twice, we have

$$V^{BAB''}(u|\hat{u}_1, \hat{u}_2) = -\left(\frac{\lambda_B}{\phi}\right)^2 \cdot \left(\Pi - \frac{c}{\lambda_B} - (V^{AB}(\hat{u}_2|\hat{u}_1, \hat{u}_2) + \hat{u}_2)\right) e^{-\frac{\lambda_B}{\phi}(u-\hat{u}_2)}.$$

Note that $V^{AB}(\hat{u}_2|\hat{u}_1, \hat{u}_2) + \hat{u}_2$ cannot exceed the first-best expected surplus $\Pi - c/\lambda_B$, thus, the above expression is negative.

Since these component functions are smoothly pasted at $\hat{u}_1$ and $\hat{u}_2$, the entire value function is concave.

(b) As in the no efficiency loss case, $\mathcal{V}(u) + u$ is increasing in u, thus, $\mathcal{V}'(u) \geq -1$ and it gives $\frac{\partial L^B}{\partial R} \leq 0$. Thus, it is sufficient to show that $L^B(u, u + \phi/\lambda_B) \leq 0$ for all $u \geq 0$. Observe that from (B.6), (B.7) and (OA.1.2), we have

$$L^B(u, u + \phi/\lambda_B) = \begin{cases} 0, & \text{if } u \leq \hat{u}_1 \text{ or } u \geq \hat{u}_2, \\ \phi V^{BAB'}(u|\hat{u}_1, u) - \phi V^{AB'}(u|\hat{u}_1), & \text{if } u \in (\hat{u}_1, \hat{u}_2). \end{cases}$$

Since $\hat{u}_1$ and $\hat{u}_2$ are chosen to satisfy $V^{BAB'}(u|\hat{u}_1, u) < V^{AB'}(u|\hat{u}_1)$ for all $u \in (\hat{u}_1, \hat{u}_2)$, L^B is always nonpositive.

Likewise, from the concavity of $\mathcal{V}$, $\frac{\partial L^A}{\partial u} \geq 0$. Thus, it is sufficient to show that $L^A(u, u + \phi/\lambda_A) \leq 0$ for all $u \geq 0$. Then, we have

$$L^A(u, u + \phi/\lambda_A) = \begin{cases} \phi V^{AB'}(u|u) - \phi V^{B'}(u), & \text{if } u \leq \hat{u}_1, \\ 0, & \text{if } u \in (\hat{u}_1, \hat{u}_2), \\ \begin{aligned} &\lambda_A(V_M(u + \phi/\lambda_A) - V^{BAB}(u|\hat{u}_1, \hat{u}_2)) \\ &\qquad -c - \phi V^{BAB'}(u|\hat{u}_1, \hat{u}_2), \end{aligned} & \text{if } u \geq \hat{u}_2. \end{cases}$$

By Lemma OA.3 and Lemma OA.5, L^A is always nonpositive.

□

Now we prove Proposition 4.

Proof of Proposition 4. We start by showing that, for each condition in (b) and (c), the switching points are as stated and the conditions in (a) also hold.

(b-i & c-i) By Lemma OA.3 (a), $V^{B\prime}(u) > V^{AB\prime}(u|u)$ for all $u > 0$. Note that $V^B(u) = V^{BAB}(u|0,0)$ and

$$\phi V^{AB\prime}(u|u) = \lambda_A(V_M(u + \phi/\lambda_A) - V^{BAB}(u|0,0)) - c$$

by (B.7) and $V^{AB}(u|u) = V^B(u)$. Therefore, with $\hat{u}_1 = \hat{u}_2 = 0$, the conditions (a-i)–(a-iv) hold trivially, and the condition (a-v) holds as demonstrated above.

(b-ii & c-ii) By Lemma OA.3 (b), there exists $\hat{u}_1 > 0$ such that the conditions (a-i) and (a-ii) hold. Next, by Lemma OA.4, there exists $\hat{u}_2 > \hat{u}_1$ such that the conditions (a-iii) and (a-iv) hold. By Lemma OA.5, the condition (a-v) holds.

(c-iii) By Lemma OA.2, $V^{AB\prime}(0|0) > V^B(0)$. By setting $\hat{u}_1 = 0$, the conditions (a-i) and (a-ii) hold trivially. Next, by Lemma OA.4, there exists $\hat{u}_2 > 0$ such that the conditions (a-iii) and (a-iv) hold. By Lemma OA.5, the condition (a-v) holds.

(d) Following the same steps of the proof of Proposition 1, $\mathcal{V}$ solves (HJB_V) subject to (IC). □

OA.2 Proofs of Proposition 2 and Proposition 3

Lemma OA.7. *Suppose that $\pi > \tilde{\pi}_M(\eta)$ and $\eta \leq c/(c+\phi)$. Then, $\hat{u}_2$ is less than $\bar{u}$.*

Proof of Lemma OA.7. Since $\mathcal{V}$ is strictly concave, $\hat{u}_2 < \bar{u}$ is equivalent to $0 < \mathcal{V}'(\hat{u}_2) = V^{AB\prime}(\hat{u}_2|\hat{u}_1) = V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$. Then, $0 < V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ is equivalent to:

$$\lambda_B(\hat{u}_2 + V^{AB}(\hat{u}_2|\hat{u}_1)) < \lambda_B \Pi - c - \phi. \tag{OA.2.1}$$

Also note that $V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2) = V^{BAB\prime}(\hat{u}_2|\hat{u}_1)$ and $V^{BAB}(\hat{u}_2|\hat{u}_1, \hat{u}_2) = V^{AB}(\hat{u}_2|\hat{u}_1)$ imply that

$$\lambda_B(\Pi - \hat{u}_2 - V^{AB}(\hat{u}_2|\hat{u}_1)) = \lambda_A \left(V_M\left(\hat{u}_2 + \phi/\lambda_A\right) + \hat{u}_2 + \phi/\lambda_A\right) - \lambda_A \left(V^{AB}(\hat{u}_2|\hat{u}_1) + \hat{u}_2\right)$$

by (B.6) and (B.7). By plugging (B.1) into the above equation, we can derive that

$$\begin{aligned}
(\lambda_A - \lambda_B)(V^{AB}(\hat{u}_2|\hat{u}_1) + \hat{u}_2) =& \lambda_A \left(\Pi - \frac{c}{\lambda_A}\right)\left(1 - e^{-\frac{\lambda_A}{\phi}\hat{u}_2 - 1}\right) - \lambda_B \Pi \\
\iff \qquad \eta\lambda_B(V^{AB}(\hat{u}_2|\hat{u}_1) + \hat{u}_2) =& \eta\lambda_B\Pi - c - (\lambda_A\Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_2 - 1}.
\end{aligned}$$

Then, by plugging this into (OA.2.1), $0 < V^{BAB\prime}(\hat{u}_2|\hat{u}_1, \hat{u}_2)$ is equivalent to

$$\eta(c+\phi) - c < (\lambda_A\Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_2 - 1}.$$

Since $\Pi > c/\lambda_B > c/\lambda_A$, the right hand side of the above inequality is always greater than 0. Since it is assumed that $\eta \leq c/(c+\phi)$, the left hand side of the above inequality is always less than or equal to 0. Therefore, the above inequality always holds, i.e., $\hat{u}_2$ is less than $\bar{u}$. $\square$

Lemma OA.8. *Suppose that $\eta > \overline{\eta} = \max\{1/(e-1), \sqrt{c/(c+\phi)}\}$. There exists $\tilde{\pi}_B(\eta) \in (\tilde{\pi}_M(\eta), \tilde{\pi}_A(\eta))$ such that $\hat{u}_1 < \overline{u}$ if and only if $\pi > \tilde{\pi}_B(\eta)$.*

Proof of Lemma OA.8. From $\eta > 1/(e-1)$, $\tilde{\pi}_A(\eta)$ exists. Suppose that $\pi \geq \tilde{\pi}_A(\eta)$. By Lemma OA.2, $\hat{u}_1 = 0$. Note that

$$\tilde{\pi}_A(\eta) = \frac{e-1}{(e-1)\eta - 1} \geq \frac{e-1}{e-2} > 2 \geq \frac{c+\phi}{c} = \pi_F.$$

Then, the project is feasible and $\bar{u}$ is greater than 0, i.e., $\bar{u} > \hat{u}_1$.

Now suppose that $\pi_M(\eta) < \pi < \pi_A(\eta)$. Since V is strictly concave, $\hat{u}_1 < \bar{u}$ is equivalent to $0 < \mathcal{V}'(\hat{u}_1) = V^{B\prime}(\hat{u}_1)$. Note that $0 < V^{B\prime}(\hat{u}_1)$ is equivalent to:

$$\frac{\phi}{\lambda_B \Pi - c} < e^{-\frac{\lambda_B}{\phi}\hat{u}_1} = \hat{x}_1. \tag{OA.2.2}$$

Recall that $\hat{x}_1$ is a solution where $\tilde{H}_1(x)$, as defined in (OA.1.4), equals zero. Additionally, $\pi_M(\eta) < \pi < \pi_A(\eta)$ implies that $\tilde{H}_1(1) < 0 < \tilde{H}_1(x^*)$ where x^* is defined in (OA.1.5).[17]

There are two possible cases that satisfy (OA.2.2): (i) $x^* \geq \frac{\phi}{\lambda_B \Pi - c}$; (ii) $\frac{\phi}{\lambda_B \Pi - c} > x^*$ and $\tilde{H}_1(\frac{\phi}{\lambda_B \Pi - c}) < 0$.

The first case is equivalent to $\tilde{H}_1'(\frac{\phi}{\lambda_B \Pi - c}) < 0$. By algebra, we can show that it is equivalent to

$$\frac{(\eta+1)\pi - 1}{(\pi-1)^{\eta+1}} < \frac{\eta e}{\eta+1}\phi^{-\eta}c^{\eta}. \tag{OA.2.3}$$

The second case is equivalent to $\tilde{H}_1'(\frac{\phi}{\lambda_B \Pi - c}) \geq 0$ and $\tilde{H}_1(\frac{\phi}{\lambda_B \Pi - c}) < 0$. By algebra, we can show that it is equivalent to

$$\frac{\eta e}{\eta+1}\phi^{-\eta}c^{\eta} \leq \frac{(\eta+1)\pi - 1}{(\pi-1)^{\eta+1}} < (\eta(c+\phi) - c)\, e\phi^{-\eta-1}c^{\eta}. \tag{OA.2.4}$$

[17] For simplicity, Π is omitted from the definition of $\tilde{H}_1$ and x^*.

Last, by the proof of Lemma OA.3, we can show that $\pi > \pi_M(\eta)$ is equivalent to

$$\frac{(\eta+1)\pi - 1}{(\pi-1)^{\eta+1}} < \left(\frac{\eta^2}{1-\eta^2}\right)^{\eta} \frac{\eta e}{1+\eta}. \tag{OA.2.5}$$

Now we compare the above three conditions. Using $\eta > \sqrt{c/(c+\phi)}$, by simple algebra, we can show that

$$\frac{\eta e}{\eta+1}\phi^{-\eta}c^{\eta} < (\eta(c+\phi)-c)e\phi^{-\eta-1}c^{\eta} < \left(\frac{\eta^2}{1-\eta^2}\right)^{\eta} \frac{\eta e}{1+\eta}.$$

Therefore, the inequality

$$\frac{(\eta+1)\pi - 1}{(\pi-1)^{\eta+1}} < (\eta(c+\phi)-c)\, e\phi^{-\eta-1}c^{\eta}$$

imply that (OA.2.3), (OA.2.4) and (OA.2.5). Define $\pi_B(\eta)$ be the value of π that makes both sides of the above inequality equal. Then, $\pi_B(\eta) > \pi_M(\eta)$ since $\pi < \pi_M(\eta)$ implies $\pi < \pi_B(\eta)$. Therefore, there exists $\pi_B(\eta) > \pi_M(\eta)$ such that $\hat{u}_1 < \bar{u}$ if and only if $\pi > \pi_B(\eta)$. $\square$

Lemma OA.9. *Suppose that $\eta > \overline{\eta} = \sqrt{c/(c+\phi)}$. Then, $\hat{u}_2 \geq \overline{u}$.*

Proof of Lemma OA.9. By following the proof of Lemma OA.7, $\hat{u}_2 \geq \bar{u}$ is equivalent to

$$\overline{y} \equiv \frac{(\eta-1)c + \eta\phi}{(\lambda_A \Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_1 - 1}} \geq e^{\frac{\lambda_A}{\phi}(\hat{u}_1 - \hat{u}_2)} = \hat{y}_2 \tag{OA.2.6}$$

By the proof of Lemma OA.4, $\hat{y}_2$ is the solution, which is not equal to 1, of $\tilde{H}_2(y) = 0$.[18] Since $\hat{u}_2 \geq \hat{u}_1$, if $\overline{y} \geq 1$, the above inequality holds, thus, we restrict attention to the case of $\overline{y} < 1$. Observe that the inequality $\tilde{H}_2(\overline{y}) \leq 0$ implies (OA.2.6) because $\tilde{H}_2$ is strictly convex in y and $\tilde{H}_2(1) \leq 0$.

[18] The function $\tilde{H}_2$ is defined in (OA.1.6).

Using the definition of $\tilde{H}_1$ in (OA.1.4) and $\hat{x}_1 \equiv e^{-\frac{\lambda_B}{\phi}\hat{u}_1}$, $\tilde{H}_2(y)$ can be rewritten as follows:

$$\tilde{H}_2(y) = \frac{1-\eta}{1+\eta}c - \tilde{H}_1(\hat{x}_1)y + \left[-\frac{1-\eta}{1+\eta}c + \frac{\eta}{1+\eta}(\lambda_A \Pi - c)e^{-\frac{\lambda_A}{\phi}\hat{u}_1 - 1}\log y\right] y.$$

Also note that $\hat{x}_1$ is chosen to satisfy $\tilde{H}_1(\hat{x}_1)$ being greater than equal to zero.

By plugging the definition of $\overline{y}$ into the above equation, we can derive that

$$\tilde{H}_2(\overline{y}) = \frac{1-\eta}{1+\eta}c(1-\overline{y}) - \tilde{H}_1(\hat{x}_1)\overline{y} + \frac{\eta}{1+\eta}\left((\eta-1)c + \eta\phi\right)\log \overline{y}.$$

Now define a new function G as follows:

$$G(y) \equiv \frac{1-\eta}{1+\eta}c(1-y) - \tilde{H}_1(\hat{x}_1)y + \frac{\eta}{1+\eta}\left((\eta-1)c + \eta\phi\right)\log y,$$

and it is enough to show that $G(y) \leq 0$ for all $y < 1$.

Note that

$$G''(y) = -\frac{\eta}{1+\eta}\left(\frac{(\eta-1)c + \eta\phi}{y^2}\right) < 0$$

from $\eta \geq \sqrt{c/(c+\phi)} > c/(c+\phi)$. Also note that

$$G'(1) = -\tilde{H}_1(\hat{x}_1) + \frac{1}{1+\eta}\left((\eta^2-1)c + \eta^2\phi\right) < 0.$$

from $\eta \geq \sqrt{c/(c+\phi)}$ and $\tilde{H}_1(\hat{x}_1) \geq 0$. Lastly, note that $G(1) = -\tilde{H}_1(\hat{x}_1) \leq 0$. Therefore, for all $y < 1$, $G(y) \leq G(1) + G'(1)(1-y) \leq 0$. Therefore, $\tilde{H}_2(\overline{y}) \leq 0$ and $u_2 \geq \bar{u}$. $\square$

Now we prove Proposition 2 and Proposition 3.

Proof of Proposition 2. Note that $\hat{u}_2$ is always greater than $\bar{u}$ by Lemma OA.9 since $\eta >$

$\sqrt{c/(c+\phi)}$. Additionally, using Lemma OA.2, Lemma OA.3 and Lemma OA.8, we have

$$\mathcal{V}(\overline{u}) = \begin{cases} V^B(\overline{u}), & \text{if } \pi_F < \pi < \tilde{\pi}_B(\eta), \\ V^{AB}(\overline{u}|\hat{u}_1), & \text{if } \tilde{\pi}_B(\eta) < \pi < \tilde{\pi}_A(\eta), \\ V^{AB}(\overline{u}|0), & \text{if } \tilde{\pi}_A(\eta) < \pi. \end{cases}$$

As in Theorem 2, the value functions above can be implemented by execution-only, one-switch and training-only contracts, respectively. □

Proof of Proposition 3. By Proposition 4 (b-i), when $\pi \leq \tilde{\pi}_M(\eta)$, $\hat{u}_1 = \hat{u}_2 = 0$. By Proposition 4 (b-ii) and Lemma OA.7, when $\pi > \tilde{\pi}_M(\eta)$, $\overline{u} > \hat{u}_2 > \hat{u}_1 > 0$. Therefore,

$$\mathcal{V}(\overline{u}) = \begin{cases} V^{BAB}(\overline{u}|0,0) = V^B(\overline{u}), & \text{if } \pi_F < \pi \leq \tilde{\pi}_M(\eta), \\ V^{BAB}(\overline{u}|\hat{u}_1, \hat{u}_2), & \text{if } \tilde{\pi}_M(\eta) < \pi. \end{cases}$$

By Lemma OA.1, $(V^{BAB}(\overline{u}|\hat{u}_1, \hat{u}_2), \overline{u})$ can be implemented by a two-switch contract. Therefore, when $\pi \in (\pi_F, \tilde{\pi}_M(\eta)]$, the execution-only contract is optimal, and when $\pi > \tilde{\pi}_M(\eta)$, there exists a two-switch contract that is optimal. □