

Simple Contracts for Due Diligence and Execution

Felipe Balmaceda*

René Caldentey†

Maia Cufre‡

September 23, 2025

Abstract

We study a continuous-time principal-agent model in which an expert privately acquires information about a project's profitability before recommending whether to execute it. We consider two contract classes: one in which the principal controls the time the agent spends collecting information, and another in which the agent chooses when to stop collecting information. This allows the principal to control the quality of the information acquired. The principal prefers the former to the latter when the agent's outside option is large, and also when the agent's speed of learning is small. When the principal chooses the contract that leaves the stopping decision to the agent, she also delegates the execution to the agent with a positive probability, which increases the set of implementable stopping times and reduces the rents given to the agent to induce him to choose the principal's desired stopping time.

Keywords: Moral Hazard, Private Information, Bayesian Learning, Optimal Stopping, Contracts, Execution, Brownian Motion, Task Bundling.

JEL-Classification: G32, J24, L26, M13

*Instituto de Política Económica, Andrés Bello University. Email: felipe.balmaceda@unab.cl

†Booth School of Business, The University of Chicago. Email: rene.caldentey@chicagobooth.edu

‡Northwestern University. Email: mcufre@u.northwestern.edu

1 Introduction

Firms frequently face complex, high-stakes decisions, such as strategic investments, mergers and acquisitions, and the development of new products. Successfully navigating these decisions relies heavily on acquiring and processing information, often through the engagement of external experts, such as consultants, financial advisors, or market analysts. These experts conduct thorough due diligence, gathering data and providing recommendations that significantly impact crucial business outcomes. This paper studies the optimality of simple information acquisition (due diligence) contracts.

In many industries, experts not only collect and communicate the information but also participate in the project’s execution—despite sometimes introducing cost inefficiencies—since this can be a powerful mechanism for aligning incentives and facilitating the implementation of specific information acquisition levels. Several industries exhibit versions of this structure in practice:

- **Pharmaceuticals:** In early-stage clinical trials, pharmaceutical firms often outsource feasibility studies to contract research organizations (CROs). To ensure diligence, CROs are sometimes offered continued involvement in later stages (e.g., regulatory submissions), even though firms may have more efficient internal teams for execution. The bundling encourages thorough information collection.
- **Software consulting:** Companies engaging consultants to assess and upgrade legacy IT systems often award implementation rights to the same consultant. Even when in-house teams could execute at lower cost, bundling ensures that the consultant does not underinvest in diagnosis or issue premature recommendations.
- **Infrastructure procurement:** Public-private partnerships (PPPs) frequently

1 delegate both feasibility assessment and construction/operation of infrastruc-
2 ture to a single private firm. This bundling incentivizes careful early-stage
3 evaluation, as the contractor bears downstream consequences.

4 • **Creative industries:** In the film sector, “first-look” deals with producers or
5 directors often bundling creative development (scriptwriting or treatment) with
6 the right to direct or produce, ensuring commitment to high-quality idea explo-
7 ration.

8 • **Startup incubation:** Accelerators like Y Combinator offer mentorship and
9 early-stage funding in bundled form. The continued support acts as both a
10 reward and an incentive for founders to signal project viability truthfully after
11 exploratory development.

12 This paper studies a continuous-time principal–agent model of optimal information
13 acquisition and provides conditions under which a contract that bundles information
14 acquisition and execution is optimal. A risk-neutral principal considers executing a
15 project whose profitability depends on an unknown binary state of the world. The
16 principal may hire a risk-neutral expert (the agent) to gather information over time
17 by observing a stream of noisy signals, gradually refining the belief about the project’s
18 value. While learning, the agent incurs a flow opportunity cost and may eventually
19 choose to recommend either pursuing or abandoning the project. Execution is costly
20 and may be assigned either to the principal or to the agent, who faces higher execution
21 costs.

22 The challenge arises from the agent’s private learning. Neither party can influence the
23 signal-generating process, which depends solely on the agent’s expertise, which is com-
24 mon knowledge. However, the principal cannot monitor intermediate signals—only
25 due diligence duration, the decision to implement or abandon the project, and the
26 project’s realized outcome are contractible. Hence, optimal contracting must align

1 the agent's incentives to acquire information and make the right recommendation
 2 with the principal's goals, while respecting the agent's limited liability constraints.

3 We analyze two classes of simple contracts that reflect real-world institutional prac-
 4 tices. First, we consider *fixed-term* contracts. In them, the principal fully commits
 5 to a due diligence (experimentation) horizon in advance and pays the agent based
 6 only on the duration of the due diligence process. This reflects common arrange-
 7 ments such as internships, fixed-fee due diligence contracts, or clinical trial protocols.
 8 Second, we study *free-term* contracts. In these cases, the principal fully commits to
 9 letting the agent determine when to cease due diligence, and the compensation is
 10 contingent upon whether the project is abandoned or completed. When due diligence
 11 ends and the project is abandoned, the agent receives a non-negative payment. When
 12 the project is executed, the agent receives a payment contingent upon the project's
 13 outcome. This resembles freelance advisory contracts or entrepreneurial partnerships,
 14 where timing and discretion over recommendations rest with the agent.¹

15 Each contract type has strengths and weaknesses. Fixed-Term contracts allow the
 16 principal to control the length of due diligence, and do not require the payment of
 17 informational rents or delegating the project execution to the agent. In contrast,
 18 Free-Term contracts offer greater flexibility, adapting the length of due diligence and
 19 the recommendations to the quality of accumulated evidence. However, they often re-
 20 quire the payment of incentive-compatible and limited-liability rents and inefficiently
 21 delegate execution to the agent. This happens because: the set of implementable op-
 22 timal stopping times when the limited liability constraint binds and delegation does
 23 not occur is a smaller set than that when delegation of execution takes place; and
 24 the rent required to induce the agent to conduct due diligence until the principal's
 25 desired hitting time is reached is smaller under delegation.

¹There is a third natural family of contracts that we do not consider in which the principal conditions the payments when the project is implemented, not only on the project's outcome but also on the time it takes to make a recommendation. This requires solving partial differential equations whose solutions are unknown or do not have analytical solutions.

1 Our analysis characterizes the optimal timing and reporting thresholds under each
2 contract, and shows (numerically) that from the point of view of the principal, neither
3 class uniformly dominates. Fixed-Term contracts strictly dominate when the agent’s
4 outside option is high, and Free-Term contracts do so when the agent’s outside option
5 is low, despite the due diligence and limited liability rents, and the inefficient agent’s
6 execution. In addition, sensitivity analysis shows that the larger the agent’s cost
7 relative to the principal’s, the smaller the outside option threshold below which the
8 principal prefers the Free-Term contract. Concerning the learning speed, we find that
9 Free-Term contracts tend to perform well when the agent’s learning speed is high,
10 and Fixed-Term contracts do so when learning speed is low. This happens because,
11 in a fixed horizon, higher-quality information is gathered without the need to pay the
12 agent rent and inefficiently delegate the project’s execution to him.

13 In summary, our analysis provides insights into a fundamental tension between the
14 commitment to the duration of due diligence and the flexibility required in decision-making
15 based on experts’ advice. Two novel and key insights emerge: fixed due diligence con-
16 tracts collect lower-quality information, but they are cheaper to implement; in con-
17 trast, flexible due diligence contracts improve information acquisition but are more
18 expensive to implement. When flexible contracts are optimal, bundling informa-
19 tion acquisition with execution—despite introducing cost inefficiencies—is a powerful
20 mechanism for aligning incentives and facilitating implementation of specific informa-
21 tion quality levels. By examining the trade-offs between controlling the duration of
22 information acquisition versus allowing the agent to adapt to evolving information,
23 we shed light on how principals can structure incentives to elicit better information,
24 faster decisions, and greater project value.

25 The remainder of the paper formalizes our model and results. Section 3 introduces
26 the model, describing the signal process, belief dynamics, and contract space. The
27 following section discusses Fixed-Term contracts. Section 5 derives the structure of
28 the optimal Free-Term contract, highlighting the roles of belief thresholds, deadlines,

1 and outcome-contingent rewards. We characterize the optimal Free-Term contract;
2 however, we cannot analytically prove whether the optimal Free-Term contract dom-
3 inates the Fixed-Term contract. In the next Section, Section 6, we provide numerical
4 exercises comparing profits and welfare under each type of contract and some numer-
5 ical comparative statics concerning the main primitives. Section 7 concludes with a
6 summary of key findings and implications for real-world contract design.

7 **2 Related Literature**

8 This paper builds on three related streams of literature—delegated expertise, statisti-
9 cal information acquisition, and multitasking with moral hazard and limited liability
10 —which we review in tandem. Across these literatures, increasing attention has been
11 devoted to how agents balance the cost, timing, and informativeness of learning, and
12 implementation, all of which are central to our setup and results. Both areas comprise
13 extensive and well-established bodies of work, and our review is not intended to be
14 exhaustive. Rather, we highlight the contributions most closely related to our setting
15 and refer the reader to these references for additional background.

16 **2.1 Delegated Expertise**

17 A substantial body of work examines how a principal can contract with an informed
18 agent or expert to gather, process, and truthfully report information. Key contribu-
19 tions in this area include Crémer and Khalil (1992), Lewis and Sappington (1997), and
20 Compte and Jehiel (2008), which explore issues such as pre-contractual information
21 acquisition, optimal screening mechanisms, and the design of incentive-compatible
22 contracts. Our paper extends this literature by focusing on the *dynamic* aspects of
23 information acquisition and by explicitly modeling the principal’s *control* over the
24 agent’s search process. Unlike Crémer and Khalil (1992) and Lewis and Sappington

1 (1997), we do not consider pre-contract information acquisition.

2 Our work is also related to the stochastic continuous-time principal-agent literature,
3 which typically focuses on the agent’s experimentation to learn about a risky bandit
4 arm (Szalay (2005), Keller et al. (2005), Gerardi and Maestri (2012), Green and
5 Taylor (2016), Klein (2016), Henry and Ottaviani (2019), Georgiadis and Szentes
6 (2020), McClellan (2022), Madsen (2022), and Feng et al. (2024)). This literature
7 examines different types of contracts that allow for time-dependent actions, transfers,
8 and incentives. Gerardi and Maestri (2012) study a dynamic principal-agent model
9 where the agent sequentially acquires costly information about an unknown binary
10 state. However, they rely on signal structures that are not absolutely continuous
11 across states, whereas our model features these structures instead.

12 Unlike these models, our paper focuses explicitly on how the principal can strate-
13 gically trade off commitment to a given due diligence horizon against flexibility in
14 the information acquisition process, allowing better control of information quality,
15 and considers delegation of execution as instrument to increase feasibility and reduce
16 rents.

17 2.2 Statistical Information Acquisition

18 A separate but related literature examines how agents acquire information over time
19 to improve decision-making under uncertainty. Foundational work in this area in-
20 cludes Wald (1947), Wald and Wolfowitz (1948), Arrow et al. (1949), which laid the
21 groundwork for sequential hypothesis testing. This literature has been extended in
22 various directions, including the use of continuous-time Bayesian formulations Ara-
23 man and Caldentey (2022) and drift-diffusion models Roberts and Weitzman (1981),
24 Bolton and Harris (1999) and Moscarini and Smith (2001). Our paper contributes to
25 this literature by providing a flexible and tractable foundation for analyzing learning
26 dynamics under incentive constraints.

More recently, scholars have begun to explore the endogenous design of information structures (Lang (2019), Zhong (2022), and Wong (2025)). Zhong (2022) generalizes this approach by allowing the agent to fully control the evolution of beliefs among a broad class of martingale processes and subject to convex flow costs. Different from such work, our contribution lies on showing how the interaction between information frictions, contractual instruments, and execution costs shapes optimal delegation and timing.

In contrast to this literature, our paper does not allow for exogenous design of information structure but focuses on combining learning and execution, even when the agent is less efficient at performing the execution task. By examining these trade-offs, our model sheds light on how principals can structure incentives to elicit better information, faster decisions, and greater project value.

2.3 Multitask Principal-Agent:

Since Holmström and Milgrom (1991), the multitasking principal-agent problem with moral hazard has been extensively studied. The one related to this paper considers multiple risk-neutral agents subject to limited liability that must perform numerous tasks (e.g., Laux (2001), Dewatripont et al. (2000), Bond and Gomes (2009), Balmaceda (2016), Winter (2006), Winter (2009), Winter (2010), Basov and Danilkina (2010)).

Winter (2010) studies a model with multiple agents performing tasks in a fixed sequence. Agents are risk-neutral and have limited liability. He shows that the optimal contract often involves agents performing later tasks in the sequence receiving higher rewards than those performing earlier tasks. Winter (2010) demonstrates that transparency facilitates the implementation of the optimal effort profile. Balmaceda (2016) shows that when an agent is responsible for more than one task that are complementary, there are effort profiles that cannot be implemented, despite being optimal.

1 A specialized job design solves the implementation problem. [Basov and Danilkina](#)
 2 [\(2010\)](#) argue that when the number of effort dimensions exceeds the number of per-
 3 formance measures observed by the principal, hidden action leads to an additional
 4 welfare loss due to the impossibility of implementing certain effort profiles.
 5 We draw from the literature the fact that, under moral hazard and limited liability,
 6 certain actions cannot be implemented through standard incentive contracts, and an
 7 alternative instrument is required. In contrast to [Winter \(2010\)](#), [Balmaceda \(2016\)](#),
 8 and [Basov and Danilkina \(2010\)](#), we show that delegating both tasks to the agent may
 9 alleviate implementation and incentive problems. [Winter \(2010\)](#) demonstrates that
 10 transparency facilitates the implementation of the optimal effort profile, and [Winter](#)
 11 [\(2006\)](#) shows that discriminating agents implement the optimal effort profile.
 12 Our paper combines these three streams of literature, providing a framework for ana-
 13 lyzing dynamic contracting with private information. Our key contributions are three-
 14 fold: (1) we characterize the optimal contract structure (fixed-term vs. free-term) as
 15 a function of the agent’s learning speed and outside option; (2) we demonstrate that
 16 strategic delegation of execution can be used to reduce information rents and increase
 17 the set of implementable stopping times, and (3) we provide insights into the interplay
 18 between commitment and flexibility in expert-based decision-making. By focusing on
 19 the principal’s *control* over the information acquisition process, we provide a more
 20 nuanced understanding of how to design effective incentive contracts in environments
 21 where information is costly to acquire and difficult to verify.

22 3 Model Setup

23 We consider the problem faced by a risk-neutral firm (the principal) that is presented
 24 with a business opportunity to execute a project whose return is unknown. Initially,
 25 the principal has three possible courses of action: to abandon the project entirely, to
 26 proceed immediately with its execution, or to hire a risk-neutral external firm (the

1 agent) with appropriate expertise to conduct due diligence and gather additional
 2 information about the project’s return. Additionally, if requested, the agent can also
 3 execute the project on behalf of the principal. The principal and the agent have a
 4 common discount $r > 0$.

5 **Project Characteristics:** The project can be, *a priori*, either “good” or “bad”.
 6 We will use θ to represent the unknown type of the project, with $\theta = 0$ indicating
 7 that the project is good and $\theta = 1$ that it is bad. The prior belief –common to the
 8 principal and the agent– that the project is bad is denoted by $\delta = \mathbb{P}(\theta = 1)$, and we
 9 assume that $\delta \in (0, 1)$.

10 The project’s type influences its return. Let $\mathcal{R}$ denote the project’s random return,
 11 and let F_θ be the cumulative distribution function of $\mathcal{R}$ given the project type θ ,
 12 which we assume has a density function f_θ . We denote by $\mathcal{R}_\theta := \mathbb{E}_\theta[\mathcal{R}]$ the expected
 13 value of $\mathcal{R}$ under F_θ .

14 The execution cost depends on whether the project is carried out by the principal or
 15 the agent. We denote the cost by $\mathcal{C}_p$ when executed by the principal and by $\mathcal{C}_A$ when
 16 executed by the agent. For simplicity of exposition, we assume that the execution
 17 cost is deterministic and independent of the project type θ , and that execution time
 18 is negligible. We assume that the distribution function F_θ and the execution costs $\mathcal{C}_p$
 19 and $\mathcal{C}_A$ are common knowledge.

20 **Due Diligence and Belief Process:** When the agent conducts due diligence ac-
 21 tivities, he privately gathers additional information about the project and use it to
 22 update his belief about its type. Let δ_t denote the agent’s belief after conducting
 23 due diligence for t time units, conditional on the information the agent has gathered
 24 during this period.

To model the stochastic evolution of the belief process δ_t , we adopt the *statistical experiment* framework from [Peskir and Shiryaev \(2006, Chapter VI, §21\)](#). Specifically, we consider a probability space $(\Omega, \mathcal{F}, \mathbb{P}_\delta, \delta \in [0, 1])$, where the random variable θ

satisfies $\mathbb{P}_\delta(\theta = 0) = 1 - \delta$ and $\mathbb{P}_\delta(\theta = 1) = \delta$. As the agent conducts due diligence, he privately observes the evolution of a signal process X given by

$$X_t = \theta t + \frac{W_t}{\sigma},$$

1 where W_t is a standard Brownian motion under $\mathbb{P}_\delta$, and $\sigma > 0$ determines the
 2 signal-to-noise ratio of the information generated through due diligence. A larger
 3 value of σ corresponds to a faster rate of learning, and we refer to σ as the “speed of
 4 learning” parameter.

In this setting, under Bayes’ rule, the belief process $\delta_t = \mathbb{P}_\delta(\theta = 1 \mid \mathcal{F}_t)$ is given by

$$\delta_t = \frac{\delta}{\delta + (1 - \delta) \mathcal{L}_t},$$

where $\mathcal{L}_t$ is the likelihood process, defined as the Radon–Nikodym derivative of $\mathbb{P}_0$ with respect to $\mathbb{P}_1$, and satisfies

$$\mathcal{L}_t = \frac{d[\mathbb{P}_0|\mathcal{F}_t]}{d[\mathbb{P}_1|\mathcal{F}_t]} = \exp\left(\sigma^2\left(\frac{t}{2} - X_t\right)\right),$$

5 where $\mathcal{F}_t$ is filtration generated by X_t .

6 Applying Itô’s lemma, we conclude that the agent’s information acquisition process
 7 causes his belief δ_t to evolve continuously over time according to the following stochas-
 8 tic differential equation:

$$d\delta_t = \delta_t(1 - \delta_t)\sigma dB_t \quad \text{with initial condition } \delta_0 = \delta, \quad (1)$$

9 where $B_t = \sigma(X_t - \int_0^t \delta_s ds)$ is a standard Brownian motion with respect to $\mathcal{F}_t$. The
 10 term $\delta_t(1 - \delta_t)$ reflects the idea that new information has a smaller effect on posterior
 11 beliefs when the agent is more certain about the project’s type, that is, when δ_t is
 12 close to 0 or 1. Under (1), we interpret $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ as the filtration generated by B_t ,

1 and we define $\mathbb{T}$ to denote the set of $\mathbb{F}$ -stopping times.

Remark 1. The continuous-time stochastic evolution of δ_t in (1) can be understood as the limit (in the sense of weak convergence) of a discrete-time belief process. In this discrete-time setting, every $\Delta > 0$ time units, the agent generates a new piece of information that is used to update his belief about the project's complexity. Let δ_n^Δ denote the agent's belief after collecting the n^{th} piece of information. Then, δ_n^Δ evolves according to Bayes's rule

$$\delta_n^\Delta = \delta_{n-1}^\Delta + (1 - \delta_{n-1}^\Delta) \delta_{n-1}^\Delta \left(\frac{1 - \mathcal{L}_n^\Delta}{\delta_{n-1}^\Delta + (1 - \delta_{n-1}^\Delta) \mathcal{L}_n^\Delta} \right),$$

2 where $\mathcal{L}_n^\Delta$ is the (random) likelihood ratio associated with the n^{th} piece of information.
3 By letting $\Delta \downarrow 0$ and allowing $\mathcal{L}_n^\Delta$ converge to 1 (a.s.) at a rate of $O(\sqrt{\Delta})$, one can
4 show that δ_n^Δ converges weakly to the continuous-time process δ_t in (1) (see Araman and
5 Caldenty, 2022 for details). With this interpretation, we consider the continuous-time
6 model as a mathematically convenient approximation of a discrete-time model, where the
7 agent collects new information at a high frequency, though each new piece carries limited
8 informative value. The advantage of a continuous-time formulation is that it will enable
9 us to apply the tools of stochastic calculus to the martingale process δ_t . $\square$

10 **Assumption 1.** *We impose the following conditions on the project's execution cost,*
11 *payoffs and information.*

- 12 (i) *The agent privately observes the evolution of δ_t while conducting due diligence*
13 *activities. Aside from the opportunity cost of allocating time to due diligence*
14 *activities, the agent's actual cost of conducting due diligence is negligible.*
- 15 (ii) *Upon execution, the project's payoff $\mathcal{R}$ is observable to both parties.*
- 16 (iii) *The principal's execution cost is less than or equal to the agent's, i.e., $\mathcal{C}_P \leq \mathcal{C}_A$.*
- 17 (iv) *Good projects have higher expected returns than bad ones, i.e., $\mathcal{R}_0 \geq \mathcal{R}_1$. Fur-*
18 *thermore, only good projects are worth executing, meaning that $\mathcal{R}_0 - \mathcal{C}_A \geq 0 \geq \mathcal{R}_1 - \mathcal{C}_P$.*

(v) The likelihood ratio $L(x) = f_1(x)/f_0(x)$ of the project's payoff is bounded above and monotonically decreasing. For future reference, let us define $\mathcal{L} := \lim_{x \rightarrow \infty} L(x)$ and note that $\mathcal{L} \in [0, 1]$. To exclude the trivial case in which the project's type has no impact on its return, we further assume that $\mathcal{L} < 1$.

Conditions (i) and (ii) capture the type of asymmetric information structure between the firms that we would expect to see in practice. Condition (iii) is not essential for the analysis that follows but helps emphasize that the principal prefers to engage the agent solely for due diligence activities. However, since the agent privately observes the evolution of δ_t , the principal may need to delegate the project's execution to the agent to ensure that an "optimal amount" of due diligence is performed. The first part of condition (iv) is intuitive, while the second part is necessary to prevent trivial solutions. If $\mathcal{R}_1 - \mathcal{C}_p \geq 0$, a risk-neutral principal would execute the project immediately, eliminating the need for any due diligence. Conversely, if $\mathcal{R}_0 - \mathcal{C}_A \leq 0$, the principal would never delegate the project's execution to the agent. Finally, condition (v) is imposed for mathematical tractability purposes as it simplifies the characterization of the set of implementable contracts (see Lemma 4).

Admissible Contracts: Given the nature of the relationship and the asymmetric information between the principal and the agent, there are two aspects of it that the principal seeks to control: first, the amount of time the agent devotes to due diligence activities; and second, the "quality" of the information the agent produces during these activities. While *a priori* granting the agent more time for due diligence may seem beneficial, it does not necessarily guarantee better information, given the stochastic, martingale nature of the belief process. Thus, the principal must design contracts that carefully balance the benefits of granting the agent sufficient time to conduct due diligence against the costs of unnecessarily delaying the execution or abandonment of the project.

The contracts considered specify three key aspects of their agreement: (i) the type

1 of service the agent is hired for, namely due diligence only or due diligence with the
 2 option to execute the project, (ii) the timeframe granted to the agent for conducting
 3 due diligence activities, and (iii) the compensation scheme the agent receives for the
 4 services performed.

5 The agent must accept the contract terms at time zero. In doing so, he must con-
 6 sider the opportunity cost associated with his outside option, denoted by $\widehat{\omega}$, which is
 7 available at any time and provides a fixed payoff that is collected immediately upon
 8 termination of the contractual relationship with the principal. This termination can
 9 occur either immediately after the completion of due diligence or after the project is
 10 executed.

11 We study two types of contracts. In the first, the principal specifies how much time
 12 the agent must devote to due diligence, treating the resulting information as a ran-
 13 dom outcome of that effort. The principal fully commits to the specified time and
 14 compensation. In the second scenario, the agent is given complete discretion over
 15 how long to conduct due diligence, while the principal uses the agent's compensation
 16 scheme to influence the agent's decision on when to stop. The principal fully commits
 17 to not intervening in the agent's stopping decision.²

18 We will refer to the first class of contracts as *Fixed-Term* contracts and to the second
 19 one as *Free-Term* contracts. Each contract is described by a pair $(\mathcal{T}, \mathcal{W})$, where
 20 $\mathcal{T}$ denotes the time allocated by the agent for conducting due diligence, and $\mathcal{W}$
 21 represents the compensation the agent receives for the services performed, which
 22 include both due diligence and, potentially, project execution. Under a Fixed-Term
 23 contract, the duration is determined in advance, i.e., $\mathcal{T} = T$ for some $T \geq 0$. In
 24 contrast, under an Free-Term contract, the agent is free to terminate the due diligence
 25 process at any time, so that the allocated time is a random variable $\mathcal{T} \in [0, \infty)$. In
 26 both cases, we impose limited liability on the agent's payoff by requiring that $\mathcal{W} \geq 0$

²Because of this, the principal will not benefit from receiving a report on the evolution of the prior belief; she only cares about the belief at the time the agent decides to stop conducting due diligence.

1 (a.s.).

2 In what follows, we investigate the solution to the principal’s problem by examining
3 the different types of contracts individually. We begin in Section 4 with contracts
4 that impose a restricted due diligence period, as they are simpler to analyze. Then,
5 in Section 5, we turn to the more complex contracts that place no restrictions on the
6 agent’s due diligence period.

7 4 Fixed-Term Contracts

8 When the due diligence duration is predetermined, it is never optimal for the principal
9 to grant the agent the option to execute the project. This follows from the fact that
10 the agent does not influence the dynamics of the belief-learning process in equation
11 (1), and as stated in point (iii) of Assumption 1, the principal incurs a lower exe-
12 cution cost than the agent. As a result, the defining characteristic of a Fixed-Term
13 contract $(\mathcal{T}, \mathcal{W})$ is that it specifies a fixed due diligence period $\mathcal{T} = T$, for some
14 fixed non-negative scalar T . Regarding the compensation $\mathcal{W}$, it is straightforward
15 to see that the principal’s optimal choice is to offer the agent a fixed payment of
16 $\mathcal{W} = \widehat{\omega}(e^{rT} - 1)$. This payment scheme compensates the agent exactly for his oppor-
17 tunity cost of conducting due diligence for T units of time while ensuring that the
18 agent truthfully reports the value of δ_T to the principal (both the compensation is
19 independent of the belief report and the agent’s preferences are independent of the
20 beliefs, so the agent does not have incentives to misreport). The principal can then
21 use this information to determine whether to proceed with the project.

22 An optimal Fixed-Term contract can then be found by maximizing the principal’s
23 expected payoff over the choice of T . For a given value of δ_T , the principal will
24 proceed with executing the project if and only if $\mathbb{E}_{\delta_T}[\mathcal{R} - \mathcal{C}_p] \geq 0$. The notation $\mathbb{E}_{\delta}[\cdot]$
25 denotes the conditional expectation operator given a belief δ ³.

³Specifically, $\mathbb{E}_{\delta}[\cdot] := (1 - \delta)\mathbb{E}_0[\cdot] + \delta\mathbb{E}_1[\cdot]$, where $\mathbb{E}_i[\cdot] := \mathbb{E}[\cdot | \theta = i]$ is the conditional expectation

1 Let us define $\Pi_\theta := \mathbb{E}_\theta[\mathcal{R} - \mathcal{C}_p]$ for $\theta = 0, 1$, which represents the principal's expected
 2 execution payoff net of execution costs when the project is of type θ . It follows
 3 that $\mathbb{E}_{\delta_T}[\mathcal{R} - \mathcal{C}_p] = (1 - \delta_T)\Pi_0 + \delta_T\Pi_1$. By Assumption 1, part (iv), we know that
 4 $\Pi_0 \geq 0 \geq \Pi_1$, implying that the principal will execute the project only if δ_T belongs
 5 to the execution region $\mathcal{E} := [0, \hat{\delta}]$, where $\hat{\delta} := \Pi_0/(\Pi_0 - \Pi_1)$. However, at time $t = 0$,
 6 the exact value of δ_T is unknown. Instead, the principal must determine the optimal
 7 value of T by solving the following optimization problem:

$$\begin{aligned}
 \Pi^F(\delta) &:= \max_{T \geq 0} \mathbb{E}_\delta \left[e^{-rT} \left(\pi(\delta_T) - \hat{\omega}(e^{rT} - 1) \right) \right] \\
 &= \max_{T \geq 0} \Pi_0 e^{-rT} \int_0^{\hat{\delta}} \frac{\mathbb{P}_\delta(\delta_T \leq x)}{\hat{\delta}} dx - \hat{\omega}(1 - e^{-rT}), \tag{2}
 \end{aligned}$$

8 where $\pi(\delta) := [(1 - \delta)\Pi_0 + \delta\Pi_1]^+$ is the project's expected payoff when the decision
 9 to execute or abandon is made at belief δ , and the second equality follows from an
 10 integration by parts argument.
 11 Solving (2) requires characterizing the probability distribution $\mathbb{P}_\delta(\delta_T \leq x)$ of δ_T given
 12 the initial belief δ and the stochastic dynamics described in (1).

Lemma 1. *Let Φ_θ denote the cumulative distribution function of a Normal random variable with mean $\mu_\theta = (\frac{1}{2} - \theta)\sigma^2 T$ and variance $\sigma_\theta^2 = \sigma^2 T$. Then, given an initial belief $\delta_0 = \delta$,*

$$\mathbb{P}_\delta(\delta_T \leq x) = (1 - \delta) \left[1 - \Phi_0 \left(\log \left(\frac{\delta(1 - x)}{(1 - \delta)x} \right) \right) \right] + \delta \left[1 - \Phi_1 \left(\log \left(\frac{\delta(1 - x)}{(1 - \delta)x} \right) \right) \right], \quad x \in (0, 1).$$

13 An explicit analytical solution to (2) appears to be unavailable, and in general, this
 14 problem must be solved numerically. Figure 1 depicts the optimal value of T^* (left
 15 panel) and the principal's payoff $\Pi^F(\delta)$ (right panel) as functions of the initial belief δ .
 16 The dashed piecewise linear function $\pi(\delta)$ in the right panel represents the principal's
 17 payoff in the absence of additional information, i.e., when the decision to execute or
 given $\theta = i$ for $i = 0, 1$.




Figure 1: Illustration of the optimal duration (T^*) of a Fixed-Term contract and the associated optimal principal's payoff (Π^D) as functions of the initial belief δ . The dashed piecewise linear function, $\pi(\delta)$, in the right panel represents the principal's payoff if a decision to execute to abandon the project is made at time $t = 0$.

- 1 abandon the project is made without any due diligence at time $t = 0$.
- 2 The optimal solution is characterized by a pair of thresholds, $\underline{\delta}^F$ and $\bar{\delta}^F$, that define the
- 3 region of initial beliefs for which it is optimal for the principal to hire the agent under
- 4 a Fixed-Term contract. For sufficiently small beliefs ($\delta \leq \underline{\delta}^F$) or sufficiently large
- 5 beliefs ($\delta \geq \bar{\delta}^F$), we have $T^* = 0$, and hiring the agent to conduct due diligence is not
- 6 beneficial. In the former case, it is optimal for the principal to execute the project
- 7 immediately without any due diligence. In the latter, it is optimal to abandon the
- 8 project immediately. For values of $\delta \in (\underline{\delta}^F, \bar{\delta}^F)$, hiring the agent to conduct due
- 9 diligence is optimal and $T^* > 0$.

10 The following property follows from the convexity of $\Pi^F(\delta)$ in δ .

11 **Lemma 2.** For all $\delta \in (0, 1)$, $\Pi^F(\hat{\delta}) - \pi(\hat{\delta}) \geq \Pi^F(\delta) - \pi(\delta)$.

12 In words, Lemma 2 establishes that the incremental value of hiring the agent to

13 conduct due diligence is maximized at the point of indifference, where $\delta = \hat{\delta}$. This

14 result also implies that $\underline{\delta}^F \leq \hat{\delta} \leq \bar{\delta}^F$. Furthermore, if $\Pi^F(\hat{\delta}) > \pi(\hat{\delta})$, then $\underline{\delta}^F < \hat{\delta}$,

15 indicating that for initial beliefs in the interval $\delta \in (\underline{\delta}^F, \hat{\delta})$, the principal is willing to

16 hire the agent even though $\delta \in \mathcal{E}$ and executing the project immediately without any

17 due diligence yields a positive expected payoff.

1 Additionally, there exists a minimum amount of due diligence, $T_{\min} > 0$, that is worth
 2 conducting. That is, T^* exhibits discontinuities at the boundaries $\delta = \underline{\delta}^F$ and $\delta = \bar{\delta}^F$,
 3 indicating that it is not optimal for the principal to engage the agent for only a small
 4 amount of time. The value of information generated during a short due diligence
 5 phase is insufficient to justify the cost required to hire the agent. In other words,
 6 $T = 0$ is a local maximum of the principal's payoff function.

7 The discontinuity at $\delta = \bar{\delta}^F$ follows from the fact that $\bar{\delta}^F \notin \mathcal{E}$ and the project would
 8 be abandoned without any due diligence. Thus, if the principal decides to hire the
 9 agent, she must do so for a sufficiently long duration to allow the agent to collect
 10 enough information so that the likelihood of the posterior belief $\delta_T \in \mathcal{E}$ moving into
 11 the execution region is large enough to justify the cost of hiring the agent. On the
 12 other hand, the discontinuity of T^* at $\delta = \underline{\delta}^F$ occurs within the execution region $\mathcal{E}$.
 13 In this case, the rationale for hiring the agent is to collect information that might
 14 lead to the abandonment of the project.⁴ Therefore, T must be large enough so that
 15 the probability of $\delta_T \notin \mathcal{E}$ is sufficiently high to justify the fixed cost of engaging the
 16 agent.

17 Practically speaking, a Fixed-Term contract must allocate enough time for belief up-
 18 dating to generate meaningful information that can guide execution or abandonment
 19 decisions. This reflects the realistic need for the agent to not only investigate but
 20 also internalize the project's details. A due diligence window that is too short fails to
 21 extract value from the agent's effort while still incurring hiring costs. This outcome
 22 suggests that the agent's learning process exhibits increasing returns in early stages,
 23 requiring a minimum duration for effective information production.

24 Extending the fixed due diligence period T creates a trade-off. Longer durations im-
 25 prove learning and lead to a more polarized posterior belief distribution, reducing

⁴In fact, if the principal were to execute the project at time T regardless of the value of δ_T , her expected payoff would be $\mathbb{E}_\delta[(1 - \delta_T)\Pi_0 + \delta_T\Pi_1] = (1 - \delta)\Pi_0 + \delta\Pi_1$, which is exactly the payoff for executing the project at time $t = 0$. Thus, there is no value in hiring the agent if the project is ultimately going to be executed.

1 the risk of acting on weak signals. However, this comes at the cost of longer engage-
2 ments, which require greater compensation. Thus, the optimal T balances the value
3 of sharper information with the cost of inducing effort.

4 Despite its simplicity, the fixed-term contract offers limited control: the principal
5 cannot influence the specific posterior belief realized at the end of due diligence, only
6 its distribution. This residual uncertainty limits the precision of project screening.

7 **5 Free-Term Contracts**

8 In this section, we investigate the principal's problem of designing an optimal con-
9 tract $(\mathcal{T}, \mathcal{W})$ within the class of Free-Term contracts with an unrestricted due dili-
10 gence period, i.e., $\mathcal{T} \in [0, \infty)$. The compensation $\mathcal{W}$ received by the agent depends
11 on the decision made and, potentially, on the realized value of $\mathcal{R}$ if the project is
12 executed. Formally, we consider a class of randomized contracts defined by a tuple
13 $\mathcal{W} = (\mathcal{W}_\emptyset, \mathcal{W}^A(\mathcal{R}), \mathcal{W}^P(\mathcal{R}), \alpha)$, where $\mathcal{W}_\emptyset$ is a fixed compensation received by the
14 agent if, following the due diligence period, the project is abandoned. The functions
15 $\mathcal{W}^A(\mathcal{R})$ and $\mathcal{W}^P(\mathcal{R})$ represent the agent's compensation, contingent on the realized
16 return $\mathcal{R}$ of the project, when it is executed by the agent or the principal, respectively.
17 Finally, the parameter $\alpha \in [0, 1]$ denotes the probability that the agent executes the
18 project if the principal proceeds with it.

While separating the agent's compensation based on who executes the project is
practically meaningful, for the purposes of the analysis that follows, it is convenient
to combine $\mathcal{W}^A(\mathcal{R})$ and $\mathcal{W}^P(\mathcal{R})$ into a single compensation scheme net of execution
costs

$$\mathcal{W}(\mathcal{R}) := \alpha (\mathcal{W}^A(\mathcal{R}) - \mathcal{C}_A) + (1 - \alpha) \mathcal{W}^P(\mathcal{R}),$$

19 which captures the agent's net compensation conditional on the project being exe-
20 cuted. Accordingly, we redefine the contract's compensation as $\mathcal{W} = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}), \alpha)$.

1 The dependence of $\mathcal{W}$ on α is reflected in the limited liability requirement. Specif-
 2 ically, since $\mathcal{W}(\mathcal{R})$ represents the agent's payoff net of execution costs—rather than
 3 the total amount paid by the principal—the limited liability constraints on $\mathcal{W}$ require
 4 that $\mathcal{W}_\emptyset \geq 0$ and $\mathcal{W}(\mathcal{R}) + \alpha \mathcal{C}_A \geq 0$ almost surely.⁵

In sum, we define the family of compensations schemes

$$\mathbb{W} := \left\{ \mathcal{W} = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}), \alpha) : \alpha \in [0, 1], \mathcal{W}_\emptyset \geq 0 \text{ and } \mathcal{W}(\mathcal{R}) + \alpha \mathcal{C}_A \geq 0 \text{ (a.s.)} \right\}.$$

5 Since the contracts in $\mathbb{W}$ have the distinctive property of including a randomization
 6 device that serves as a mechanism to probabilistically bundle the agent's services
 7 of conducting due diligence and executing the project, we refer to them as *partially*
 8 *bundled* contracts, and to α as the *bundling parameter* (or *bundling probability*) that
 9 characterizes the contract.

10 **Remark 2.** (Non-Randomized Contracts) An alternative non-randomized interpreta-
 11 tion of the class of contracts in $\mathbb{W}$ is to view project execution as consisting of a large
 12 number of small tasks. Under this interpretation, the contract stipulates that a fraction
 13 α of the tasks is executed by the agent, while the remaining fraction $1 - \alpha$ is carried out
 14 by the principal.

15 In cases where project execution cannot be split, we can still recover simpler, non-randomized
 16 contracts by restricting the value of α to the discrete set $\{0, 1\}$.

17 1. Information-Only Contracts ($\alpha = 0$): Contracts in which the agent is hired ex-
 18 clusively to conduct due diligence, while the principal is responsible for execution.

⁵Note that for any combined compensation scheme, $\mathcal{W}(\mathcal{R})$, that satisfies the limited liability condi-
 tion $\mathcal{W}(\mathcal{R}) + \alpha \mathcal{C}_A \geq 0$, we can construct actual compensation functions $\mathcal{W}^A(\mathcal{R})$ and $\mathcal{W}^P(\mathcal{R})$ that
 satisfy the limited liability requirements $\mathcal{W}^A(\mathcal{R}) \geq 0$ and $\mathcal{W}^P(\mathcal{R}) \geq 0$ almost surely. In fact, any
 pair $\mathcal{W}^A(\mathcal{R})$ and $\mathcal{W}^P(\mathcal{R})$ that meets the conditions

$$\mathcal{W}^A(\mathcal{R}) = \frac{\mathcal{W}(\mathcal{R}) + \alpha \mathcal{C}_A - (1 - \alpha) \mathcal{W}^P(\mathcal{R})}{\alpha} \quad \text{and} \quad 0 \leq \mathcal{W}^P(\mathcal{R}) \leq \frac{\mathcal{W}(\mathcal{R}) + \alpha \mathcal{C}_A}{1 - \alpha}$$

satisfies $\mathcal{W}^A(\mathcal{R}) \geq 0$ and $\mathcal{W}^P(\mathcal{R}) \geq 0$.

1 In this case, the agent's compensation is given by $d \mathcal{W}^P(\mathcal{R}) + (1 - d) \mathcal{W}_\emptyset$, where
 2 $d = 1$ if the project is executed and $d = 0$ if it is abandoned.

3 2. Fully Bundled Contracts ($\alpha = 1$): Contracts in which the agent is hired to perform
 4 both due diligence and execution. Here, the agent's compensation is given by
 5 $d(\mathcal{W}^A(\mathcal{R}) - \mathcal{C}_A) + (1 - d) \mathcal{W}_\emptyset$. $\diamond$

6 5.1 Agent's Optimal Strategy

7 Presented with a contract $(\mathcal{T}, \mathcal{W})$ with $\mathcal{T} \in [0, \infty)$ and $\mathcal{W} \in \mathbb{W}$, the agent first decides
 8 whether to accept or reject it. If the agent rejects the contract, he immediately
 9 receives his outside option payoff $\widehat{\omega}$. If the agent accepts, he proceeds with due
 10 diligence activities until a (potentially random) stopping time $\tau \in \mathbb{T}$. At that point,
 11 based on the updated belief δ_τ , the agent makes a recommendation $d \in \{0, 1\}$, either
 12 to execute the project ($d = 1$) or to abandon it ($d = 0$), and receives an expected
 13 compensation (net of execution cost) equal to $\mathbb{E}_{\delta_\tau}[d \mathcal{W}(\mathcal{R}) + (1 - d) \mathcal{W}_\emptyset]$, in addition
 14 to his outside option $\widehat{\omega}$.

15 The agent determines his optimal strategy (τ^*, d^*) using backward programming. At
 16 time τ , conditional on the value δ_τ , the agent selects

$$d^* = \mathbb{1}(\mathcal{W}_\emptyset \leq \mathbb{E}_{\delta_\tau}[\mathcal{W}(\mathcal{R})]).$$

17 Let us define the agent's expected compensation net of execution cost as follows:

$$\mathcal{W}_\theta := \mathbb{E}_\theta[\mathcal{W}(\mathcal{R})] \quad \text{for } \theta \in \{0, 1\}. \quad (3)$$

It follows that $d^* = \mathbb{1}(\mathcal{W}_\emptyset \leq (1 - \delta_\tau) \mathcal{W}_0 + \delta_\tau \mathcal{W}_1)$. As a result, the agent's expected payoff under d^* equals:

$$V(\delta_\tau) := \widehat{\omega} + \max \left\{ \mathcal{W}_\emptyset, (1 - \delta_\tau) \mathcal{W}_0 + \delta_\tau \mathcal{W}_1 \right\}.$$

Thus, the tuple $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1, \alpha)$ summarizes all the payoff-relevant characteristics of the contract's payment $\mathcal{W} = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}), \alpha) \in \mathbb{W}$. As such, we will treat them as the decision variables that the principal must select.

Remark 3. Contracts with $\mathcal{W}_\emptyset > 0$ offer the agent an arbitrage opportunity, in the sense that the agent can choose $\tau^* = 0$ and $d^* = 0$, thereby immediately receiving a total compensation of $\mathcal{W}_\emptyset + \widehat{\omega} > \widehat{\omega}$ at no cost and without providing any valuable information to the principal. Therefore, any contract with a payment scheme $\mathcal{W} = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}), \alpha)$ where $\mathcal{W}_\emptyset > 0$ must ensure that $\mathcal{W}(\mathcal{R})$ provides the agent with the proper incentive to engage in a meaningful level of due diligence. $\diamond$

The agent determines τ^* by solving an optimal stopping problem:

$$\mathcal{V}(\delta) = \sup_{\tau \in \mathbb{T}} \mathbb{E}_\delta [e^{-r\tau} V(\delta_\tau)] \quad \text{subject to} \quad d\delta_t = \delta_t(1 - \delta_t)\sigma dB_t, \quad \delta_0 = \delta. \quad (4)$$

Intuitively, the solution to the agent's problem in (4) involves partitioning the belief domain $[0, 1]$ into a *continuation* region, where the agent actively conducts due diligence, and an *intervention* region, where the agent stops and selects an optimal strategy d^* , as discussed above. Consequently, the optimal stopping time τ^* is either zero if the initial belief δ belongs to the intervention region or equal to the first exit time of the belief process δ_t from the continuation region. From the continuity of the belief process, it follows that when δ belongs to the continuation region, the optimal solution to (4) is defined by a pair of thresholds $\underline{\delta}$ and $\bar{\delta}$, with $\underline{\delta} < \delta < \bar{\delta}$, such that $\tau^* = \inf\{t \geq 0: \delta_t \notin (\underline{\delta}, \bar{\delta})\}$.

In what follows, we formalize the previous intuition using a quasi-variational inequality (QVI) approach, similar to [Araman and Caldentey \(2022\)](#). To this end, let us define the set of continuously differentiable functions

$$\widehat{\mathcal{C}}^2 := \left\{ f \in \mathcal{C}^1[0, 1] : f''(\delta) \text{ exists } \forall \delta \in [0, 1] \setminus S(f) \text{ for some finite set } S(f) \subseteq [0, 1] \right\} \quad (5)$$

1 and the operator $\mathcal{H}$ on $\widehat{\mathcal{C}}^2$

$$(\mathcal{H}f)(\delta) := \frac{1}{2} \sigma^2 \delta^2 (1 - \delta)^2 f''(\delta) - r f(\delta), \quad \text{for all } \delta \in [0, 1] \setminus S(f). \quad (6)$$

2 **Definition 1.** *The function $f \in \widehat{\mathcal{C}}^2$ satisfies the quasi-variational inequalities for the*
 3 *agent's optimal stopping problem in (4), if for all $\delta \in [0, 1] \setminus S(f)$*

$$\begin{aligned} f(\delta) - V(\delta) &\geq 0 \\ (\mathcal{H}f)(\delta) &\leq 0 \\ (f(\delta) - V(\delta)) (\mathcal{H}f)(\delta) &= 0. \quad \square \end{aligned} \quad (\text{QVI})$$

For every solution $f \in \widehat{\mathcal{C}}^2$ of the (QVI) conditions, we associate a stopping time τ_f given by

$$\tau_f = \inf\{t > 0: f(\delta_t) = V(\delta_t)\}.$$

4 **Theorem 1.** (VERIFICATION) *Let $f \in \widehat{\mathcal{C}}^2$ be a solution of (QVI). Then, $f(\delta) \geq \mathcal{V}(\delta)$*
 5 *for every $\delta \in [0, 1]$. In addition, if there exists control τ_f associated with f such that*
 6 *$\mathbb{E}[\tau_f] < \infty$, then τ_f is optimal and $f(\delta) = \mathcal{V}(\delta)$.*

7 According to the previous result, at optimality, the QVI conditions partition the
 8 interval $[0, 1]$ into a continuation region where $\mathcal{V}(\delta) > V(\delta)$ and an intervention region
 9 where $\mathcal{V}(\delta) = V(\delta)$. In the continuation region, the third QVI condition implies that
 10 $\mathcal{V}(\delta)$ solves $(\mathcal{H}\mathcal{V})(\delta) = 0$, that is,

$$\frac{(\sigma \delta (1 - \delta))^2}{2} \mathcal{V}''(\delta) - r \mathcal{V}(\delta) = 0. \quad (7)$$

11 The two independent solutions to this ODE are given by $F(\delta)$ and $F(1 - \delta)$ with

$$F(\delta) := \frac{(1 - \delta)^\gamma}{\delta^{\gamma-1}} \quad \text{and} \quad \gamma := \frac{1 + \sqrt{1 + 8r/\sigma^2}}{2}. \quad (8)$$

1 The general solution to (7) is of the form $\mathcal{V}(\delta) = A_0 F(\delta) + A_1 F(1 - \delta)$, where
 2 A_0 and A_1 are constants of integration, whose values are determined by imposing
 3 value-matching and smooth-pasting conditions.

4 Proposition 1 below characterizes the optimal solution to the agent's problem, as
 5 a function of the triplet $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$. The proposition is formulated under the
 6 additional condition $\mathcal{W}_1 < \mathcal{W}_\emptyset < \mathcal{W}_0$, which, as we will show later in Lemma 5, must
 7 be satisfied by an optimal contract⁶. This condition is sufficient for the agent to make
 8 a recommendation consistent with the information acquired through due diligence,
 9 since he has no preferences for either state or over his information. In addition, in
 10 stating this result, we utilize an auxiliary function that will play a central role in the
 11 analysis that follows:

$$\widehat{\mathcal{V}}(\delta; \bar{\delta}) := \frac{(\gamma - \bar{\delta})}{(2\gamma - 1)} \frac{F(\delta)}{F(\bar{\delta})} + \frac{(\gamma + \bar{\delta} - 1)}{(2\gamma - 1)} \frac{F(1 - \delta)}{F(1 - \bar{\delta})} \quad \text{for } 0 < \delta \leq \bar{\delta}. \quad (9)$$

12 When viewed as a function of δ for fixed $\bar{\delta}$, $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ corresponds to a solution to (7)
 13 resulting from imposing value-matching and smooth-pasting conditions, $\widehat{\mathcal{V}}(\bar{\delta}; \bar{\delta}) = 1$
 14 and $\widehat{\mathcal{V}}'(\bar{\delta}; \bar{\delta}) = 0$ at $\delta = \bar{\delta}$. Here, $\widehat{\mathcal{V}}'(\delta; \bar{\delta})$ denotes the derivative of $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ with
 15 respect to its first argument δ . The function $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ is decreasing and convex in δ
 16 and increasing in $\bar{\delta}$, which are properties that we use to derive the following result.

17 **Proposition 1.** *Consider a contract specified by $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$ such that $\mathcal{W}_1 < \mathcal{W}_\emptyset < \mathcal{W}_0$*

⁶ In fact, it is straightforward to see that an optimal contract must satisfy the weaker condition

$$\min\{\mathcal{W}_0, \mathcal{W}_1\} < \mathcal{W}_\emptyset < \max\{\mathcal{W}_0, \mathcal{W}_1\}.$$

Otherwise, if $\mathcal{W}_\emptyset \geq \max\{\mathcal{W}_0, \mathcal{W}_1\}$, the agent would always choose $\tau = 0$ and recommend $d = 0$. Conversely, if $\mathcal{W}_\emptyset \leq \min\{\mathcal{W}_0, \mathcal{W}_1\}$, the agent would again choose $\tau = 0$ and, in this case, always recommend $d = 1$. In either case, the contract provides no information to the principal. Intuitively, the additional requirement $\mathcal{W}_\emptyset \geq \mathcal{W}_1$ follows from the principle that an optimal contract should align the incentives of the principal and the agent to ensure that only good projects are executed. Thus, the agent should be incentivized to report $d = 1$ only when this outcome is more likely (see Lemma 5 for details).

1 and let the agent's expected discounted payoff, $\mathcal{V}(\delta)$, be given as in (4). Then,

$$\mathcal{V}(\delta) = \begin{cases} (1 - \delta)(\hat{\omega} + \mathcal{W}_0) + \delta(\hat{\omega} + \mathcal{W}_1) & \text{if } 0 \leq \delta \leq \underline{\delta}^* \\ (\hat{\omega} + \mathcal{W}_0) \hat{\mathcal{V}}(\delta; \bar{\delta}^*) & \text{if } \underline{\delta}^* < \delta < \bar{\delta}^* \\ \hat{\omega} + \mathcal{W}_0 & \text{if } \bar{\delta}^* \leq \delta \leq 1, \end{cases} \quad (10)$$

2 where the thresholds $\underline{\delta}^*$ and $\bar{\delta}^*$ are determined imposing value-matching ($\mathcal{V}(\delta) = V(\delta)$)
3 and smooth-pasting ($\mathcal{V}'(\delta) = V'(\delta)$) conditions at $\delta = \underline{\delta}^*$ and $\delta = \bar{\delta}^*$, and satisfy
4 $\underline{\delta}^* < (\mathcal{W}_0 - \mathcal{W}_0)/(\mathcal{W}_0 - \mathcal{W}_1) < \bar{\delta}^*$. The agent's optimal strategy (τ^*, d^*) is given by
5 $\tau^* = \inf\{t > 0: \delta_t \notin (\underline{\delta}^*, \bar{\delta}^*)\}$ and $d^* = \mathbb{1}(\delta_{\tau^*} \leq \underline{\delta}^*)$.

6 The agent's optimal solution in Proposition 1 is illustrated in Figure 2.




Figure 2: Agent's expected discounted payoff $\mathcal{V}(\delta)$ as a function of the belief δ . The range of beliefs is partitioned into three regions: (i) for $\delta \in [0, \underline{\delta}^*]$ the agent accepts the contract, (ii) for $\delta \in (\underline{\delta}^*, \bar{\delta}^*)$ the agent conducts due diligence and (iii) for $\delta \in [\bar{\delta}^*, 1]$ the agent rejects the contract.

7 Let us now discuss the economic intuition behind the structure of the compensation
8 scheme illustrated in Figure 2. Since the agent is risk-neutral, only the expected
9 payments matter for incentive provision. The key statistics are the expected com-
10 pensations conditional on the state: $\mathcal{W}_0$ captures the expected net compensation
11 when the project is executed and it turns out to be of high quality, while $\mathcal{W}_1$ is the
12 corresponding expected compensation when the project is of low quality. To ensure
13 truthful revelation through the stopping rule, the contract offers a reward in the

1 good state ($\mathcal{W}_0 > \widehat{\omega}$) and a punishment in the bad state ($\mathcal{W}_1 < \widehat{\omega}$), relative to the
 2 agent’s outside option. Otherwise, the agent will always recommend the project to
 3 be executed.

4 This contract closely parallels the logic of the classic static moral hazard model with
 5 a risk-neutral agent and limited liability. There, compensation is made contingent on
 6 observable output to induce unobservable effort; here, compensation is contingent on
 7 the (endogenous) stopping decision and on the observable output in case of execution
 8 to induce experimentation and an honest recommendation. In our dynamic setting, the
 9 compensation scheme effectively pins down the agent’s stopping behavior, making the
 10 width of the due diligence region ($\bar{\delta} - \underline{\delta}$) endogenous to the reward–punishment spread
 11 between $\mathcal{W}_0$ and $\mathcal{W}_1$ (smooth pasting conditions) and to the expected compensation
 12 (value-matching condition).

13 5.2 Alternative Parametrization

14 The result in Proposition 1 characterizes the agent’s best response strategy (τ^*, d^*) in
 15 terms of the pair of thresholds $(\underline{\delta}^*, \bar{\delta}^*)$ for a given compensation scheme $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$.
 16 To solve the principal’s problem, we find it more convenient to treat the thresholds
 17 $(\underline{\delta}, \bar{\delta})$ as decision variables and express $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$ in terms of these associated
 18 thresholds. This alternative parametrization is justified by the fact that, as a corol-
 19 lary of Proposition 1, there exists a one-to-one correspondence between a pair of
 20 thresholds $(\underline{\delta}, \bar{\delta})$ and an optimal triplet $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$ associated with a contract (see
 21 Proposition 4). Here, optimality should be understood from the principal’s perspec-
 22 tive, in the sense of minimizing the compensation required to incentivize the agent to
 23 conduct due diligence as long as the belief δ_t remains within the interval $(\underline{\delta}, \bar{\delta})$.

24 To this end, Lemma 3 expresses $(\mathcal{W}_0, \mathcal{W}_1)$ in terms of $\mathcal{W}_\emptyset$ and the thresholds $(\underline{\delta}, \bar{\delta})$.
 25 In its statement—and throughout the analysis that follows—we will make extensive

1 use of the following shorthand notation:

$$\mathcal{A}(\underline{\delta}, \bar{\delta}) := \widehat{\mathcal{V}}(\underline{\delta}; \bar{\delta}) - \underline{\delta} \widehat{\mathcal{V}}'(\underline{\delta}; \bar{\delta}) \quad \text{and} \quad \mathcal{B}(\underline{\delta}, \bar{\delta}) := \widehat{\mathcal{V}}(\underline{\delta}; \bar{\delta}) + (1 - \underline{\delta}) \widehat{\mathcal{V}}'(\underline{\delta}; \bar{\delta}). \quad (11)$$

2 Recall that $\widehat{\mathcal{V}}'(\delta; \bar{\delta})$ denotes the derivative of the function $\widehat{\mathcal{V}}(\delta; \bar{\delta})$, defined in (9), with
 3 respect to δ . Additionally, we note that the optimality condition $\mathcal{W}_1 < \mathcal{W}_0 < \mathcal{W}_0$
 4 implies $\underline{\delta} < \bar{\delta}$, which in turn leads to $\mathcal{B}(\underline{\delta}, \bar{\delta}) < 1 < \mathcal{A}(\underline{\delta}, \bar{\delta})$.

5 **Remark 4.** To simplify the notation, we will regularly omit the explicit dependence of
 6 quantities such as $\mathcal{A}(\underline{\delta}, \bar{\delta})$, $\mathcal{B}(\underline{\delta}, \bar{\delta})$, and $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ on the pair of thresholds $(\underline{\delta}, \bar{\delta})$, and instead
 7 write $\mathcal{A}$, $\mathcal{B}$, and $\widehat{\mathcal{V}}$. $\diamond$

8 **Lemma 3.** *For a given value of $\mathcal{W}_0 \geq 0$ and a pair of thresholds $\underline{\delta}$ and $\bar{\delta}$ with
 9 $0 < \underline{\delta} < \bar{\delta} < 1$ there exists a unique pair $\mathcal{W}_0$ and $\mathcal{W}_1$ with $\mathcal{W}_1 < \mathcal{W}_0 < \mathcal{W}_0$ such that
 10 the agent's optimal strategy is to conduct due diligence in the interval $\delta \in (\underline{\delta}, \bar{\delta})$ when
 11 offered a contract $(\mathcal{W}_0, \mathcal{W}_0, \mathcal{W}_1)$. In particular,*

$$\mathcal{W}_0 + \widehat{\omega} = \mathcal{A}(\mathcal{W}_0 + \widehat{\omega}) \quad \text{and} \quad \mathcal{W}_1 + \widehat{\omega} = \mathcal{B}(\mathcal{W}_0 + \widehat{\omega}). \quad (12)$$

12 Furthermore, the principal's payment to the agent is given by
 13 $\mathcal{W}_0 + d^* ((\widehat{\mathcal{V}} - 1)(\widehat{\omega} + \mathcal{W}_0) + \alpha \mathcal{C}_A)$, while the agent's realized payoff net of execu-
 14 tion costs is equal to $(1 + d^* (\widehat{\mathcal{V}} - 1))(\mathcal{W}_0 + \widehat{\omega})$.

15 We find it instructive to visualize the previous result in terms of the function $\widehat{\mathcal{V}}$, as
 16 depicted in Figure 3. For a given pair of thresholds $(\underline{\delta}, \bar{\delta})$ with $\underline{\delta} < \bar{\delta}$ and $\mathcal{W}_0 \geq 0$, the
 17 corresponding values of $\mathcal{W}_0$ and $\mathcal{W}_1$ are determined based on the intercepts at $\delta = 0$
 18 and $\delta = 1$ of the tangent line to the function $\widehat{\mathcal{V}}(\delta, \bar{\delta})$ at $\delta = \underline{\delta}$.

19 The benefit of treating the thresholds $(\underline{\delta}, \bar{\delta})$ as decision variables is that it allows for a
 20 precise probabilistic characterization of the agent's optimal strategy (τ^*, d^*) when due
 21 diligence is conducted, which proves useful in solving the principal's problem. This
 22 result is derived using the dynamics of the belief process δ_t , as detailed in Equation (1),




Figure 3: Illustration of Lemma 3.

1 the first-exit time representation of τ^* in Proposition 1, and Dynkin's formula (see
2 Øksendal, 2013).

3 **Proposition 2.** Suppose $\delta \in (\underline{\delta}, \bar{\delta})$ and let $\tau = \inf\{t > 0: \delta_t \notin (\underline{\delta}, \bar{\delta})\}$ be the first-exit
4 time of δ_t from the interval $(\underline{\delta}, \bar{\delta})$. Then, τ has the moment generating function

$$\mathbb{E}_\delta[e^{-r\tau}] = \frac{(F(1 - \bar{\delta}) - F(1 - \underline{\delta})) F(\delta) + (F(\underline{\delta}) - F(\bar{\delta})) F(1 - \delta)}{F(\underline{\delta}) F(1 - \bar{\delta}) - F(\bar{\delta}) F(1 - \underline{\delta})}$$

and satisfies $\mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \underline{\delta})] = \frac{F(\delta) F(1 - \bar{\delta}) - F(\bar{\delta}) F(1 - \delta)}{F(\underline{\delta}) F(1 - \bar{\delta}) - F(\bar{\delta}) F(1 - \underline{\delta})}$, where $F(\delta)$ is given in (8).

The expected amount of time the agent spends conducting due diligence is equal to

$$\mathbb{E}_\delta[\tau] = \left(\frac{\bar{\delta} - \delta}{\bar{\delta} - \underline{\delta}} \right) g(\underline{\delta}) + \left(\frac{\delta - \underline{\delta}}{\bar{\delta} - \underline{\delta}} \right) g(\bar{\delta}) - g(\delta), \quad \text{where} \quad g(\delta) = \frac{2(1 - 2\delta)}{\sigma^2} \ln \left(\frac{1 - \delta}{\delta} \right).$$

Finally, the probabilities that the agent recommends either the execution or abandonment of the project are given by

$$\mathbb{P}_\delta(d = 1) = \frac{\bar{\delta} - \delta}{\bar{\delta} - \underline{\delta}} \quad \text{and} \quad \mathbb{P}_\delta(d = 0) = \frac{\delta - \underline{\delta}}{\bar{\delta} - \underline{\delta}}, \quad \text{respectively.}$$

5.3 Implementation

Given the agent's best response strategy –characterized by the thresholds $(\underline{\delta}, \bar{\delta})$ and the payment $\mathcal{W}_\emptyset$, as derived in Proposition 2– we now turn to identifying the class of admissible thresholds $(\underline{\delta}, \bar{\delta})$ for which the associated compensation $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$, determined in Lemma 3, can be implemented under limited liability. Recall that the limited liability constraints require $\mathcal{W}_\emptyset \geq 0$ and $\mathcal{W}_\theta + \alpha \mathcal{C}_A \geq 0$ for $\theta = 0, 1$.

As we will see, in general, for an implementable pair of thresholds $(\underline{\delta}, \bar{\delta})$, there exists a continuum of compensation schemes $\mathcal{W} = (\mathcal{W}_0, \mathcal{W}_1, \mathcal{W}_\emptyset)$ that can implement it. Thus, we also address the problem of characterizing an optimal scheme $\mathcal{W}^* = (\mathcal{W}_0^*, \mathcal{W}_1^*, \mathcal{W}_\emptyset^*)$, in the sense that it minimizes the principal's expected compensation to the agent.

Without considering the problem of implementation, the principal would ideally minimize the value of $\mathcal{W}_\emptyset$. This follows from Lemma 3, as reducing $\mathcal{W}_\emptyset$ lowers the agent's payment $\mathcal{W}_\emptyset + d^* ((\hat{\mathcal{V}} - 1)(\hat{\omega} + \mathcal{W}_\emptyset) + \alpha \mathcal{C}_A)$, without affecting the agent's strategy $(\underline{\delta}, \bar{\delta})$, that is, the time spent on due diligence or the final information provided. However, setting $\mathcal{W}_\emptyset$ to its minimum feasible value, $\mathcal{W}_\emptyset = 0$, may not always be viable, as limited liability also requires $\mathcal{W}_\theta + \alpha \mathcal{C}_A \geq 0$. For instance, the particular pair of thresholds $(\underline{\delta}, \bar{\delta})$ depicted in Figure 3 is not implementable under an information-only contract ($\alpha = 0$) since in this case $\mathcal{W}_1 \geq 0$, which necessarily implies $\mathcal{B} \geq 0$.

To address the implementation problem, we introduce the auxiliary function

$$\mathbf{W}_1(\omega_0) := \inf_{\mathcal{W}(\cdot)} \mathbb{E}_1[\mathcal{W}(\mathcal{R})] \quad \text{subject to} \quad \mathbb{E}_0[\mathcal{W}(\mathcal{R})] = \omega_0 \quad \text{and} \quad \mathcal{W}(\mathcal{R}) + \alpha \mathcal{C}_A \geq 0 \quad (\text{a.s.}), \quad (13)$$

which provides a lower bound on the agent's expected compensation for executing a bad project, $\mathcal{W}_1$, given a fixed expected compensation ω_0 for executing a good project, i.e., given that $\mathcal{W}_0 = \omega_0$. From condition (v) in Assumption 1, we get the following characterization of $\mathbf{W}_1(\omega_0)$. Recall that $\mathcal{L} = \lim_{x \rightarrow \infty} L(x)$.

Lemma 4. *Under Assumption 1, point (v), $\mathbf{W}_1(\omega_0) = \mathcal{L} \omega_0 - (1 - \mathcal{L}) \alpha \mathcal{C}_A$ for all*

$$\omega_0 \geq -\alpha \mathcal{C}_A.$$

It follows that for any implementable contract $(\mathcal{W}_0, \mathcal{W}_0, \mathcal{W}_1)$, we must have $\mathcal{W}_1 \geq \mathcal{L} \mathcal{W}_0 - (1 - \mathcal{L}) \alpha \mathcal{C}_A$. From the lemma, it is explicit that the bundling parameter directly relaxes the limited liability constraint. Furthermore, according to Lemma 3, if a contract implements a pair of thresholds $(\underline{\delta}, \bar{\delta})$, then $\mathcal{W}_0$ and $\mathcal{W}_1$ satisfy $\mathcal{B} \mathcal{W}_0 - \mathcal{A} \mathcal{W}_1 = (\mathcal{A} - \mathcal{B}) \hat{\omega}$. In addition, since we also require $\mathcal{W}_0 \geq 0$, (12) implies that $\mathcal{W}_0 \geq (\mathcal{A} - 1) \hat{\omega}$. Define

$$\mathcal{I} := \left\{ (\mathcal{W}_0, \mathcal{W}_1) : \mathcal{B} \mathcal{W}_0 - \mathcal{A} \mathcal{W}_1 = (\mathcal{A} - \mathcal{B}) \hat{\omega}, \quad \mathcal{W}_1 \geq \mathcal{L} \mathcal{W}_0 - (1 - \mathcal{L}) \alpha \mathcal{C}_A, \quad \mathcal{W}_0 \geq (\mathcal{A} - 1) \hat{\omega} \right\}. \quad (14)$$

We conclude that a pair of thresholds $(\underline{\delta}, \bar{\delta})$ can be implemented by a feasible compensation scheme $(\mathcal{W}_0, \mathcal{W}_1, \mathcal{W}_0)$ if the set $\mathcal{I}$ is non-empty. Combining the first two conditions, we see that $\mathcal{I}$ is non-empty if there exists a $\mathcal{W}_0$ such that

$$(\mathcal{B} - \mathcal{L} \mathcal{A}) \mathcal{W}_0 \geq (\mathcal{A} - \mathcal{B}) \hat{\omega} - \mathcal{A} (1 - \mathcal{L}) \alpha \mathcal{C}_A \quad \text{and} \quad \mathcal{W}_0 \geq (\mathcal{A} - 1) \hat{\omega}.$$

This leads to the next result.

Proposition 3. *For $\alpha \in [0, 1]$, the set $\mathcal{X}(\alpha)$ of threshold pairs $(\underline{\delta}, \bar{\delta})$ that can be implemented through a partially bundled contract with bundling parameter α is given by:*

$$\mathcal{X}(\alpha) = \left\{ \mathcal{B} > \mathcal{L} \mathcal{A} \right\} \cup \left\{ \mathcal{B} \geq \mathcal{L} \mathcal{A} + (1 - \mathcal{L}) \left(\frac{\hat{\omega} - \alpha \mathcal{C}_A}{\hat{\omega}} \right) \right\}.$$

The fact that $\mathcal{L} < 1$ yields the following immediate corollary.

Corollary 1. *The set $\mathcal{X}(\alpha)$ is non-decreasing in α . In particular, the set $\mathcal{X}(0)$ corresponding to an information-only contract is a subset of $\mathcal{X}(1)$, the set of implementable thresholds under a fully bundled contract. Moreover, $\mathcal{X}(0) = \mathcal{X}(1)$ if and only if $\mathcal{C}_A < \hat{\omega}$.*

1 It follows that a pair of thresholds $(\underline{\delta}, \bar{\delta})$ is implementable by a partially bundled
 2 contract if and only if it is implementable by a fully bundled contract, that is,
 3 $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(1)$. Moreover, the range of feasible bundling parameters α that can be
 4 used to implement $(\underline{\delta}, \bar{\delta})$ is given by

$$\mathbb{1}(\mathcal{B} \leq \mathcal{L} \mathcal{A}) \bar{\alpha} \leq \alpha \leq 1, \quad \text{where} \quad \bar{\alpha} = \bar{\alpha}(\underline{\delta}, \bar{\delta}) := \frac{\hat{\omega}}{\mathcal{C}_A} \left(1 - \frac{\mathcal{B} - \mathcal{L} \mathcal{A}}{1 - \mathcal{L}} \right). \quad (15)$$

5 Figure 4 illustrates the region $\mathcal{X}(\alpha)$ of implementable pairs $(\underline{\delta}, \bar{\delta})$ for different values
 6 of the parameters σ and $\mathcal{C}_A$, with $r = 0.05$, $\mathcal{L} = 0.3$ and $\hat{\omega} = 1$ fixed.⁷ The top row
 7 has a fixed value of $\mathcal{C}_A = 5$, while the bottom row has $\sigma = 0.5$. The inner region
 8 (lighter shade) represents $\mathcal{X}(0)$, which, as stated in Corollary 1, is a subset of $\mathcal{X}(1)$
 9 (darker shade).




Figure 4: Region $\mathcal{X}(\alpha)$ of implementable thresholds $(\underline{\delta}; \bar{\delta})$ for different values of γ and $\mathcal{C}_A$. The inner region (lighter shade) represents $\mathcal{X}(0)$, which, as stated in Corollary 1, is a subset of $\mathcal{X}(1)$ (the union of darker and lighter shade regions).

10 The top plots show that the region of implementable thresholds is maximal when

⁷Note that from the definition of $\hat{\mathcal{V}}(\delta, \bar{\delta})$ in (9), the values of $\mathcal{A}$ and $\mathcal{B}$ in (11) depends exclusively on the parameter $\gamma := \frac{1 + \sqrt{1 + 8r/\sigma^2}}{2}$, as defined in (8).

$\sigma = \infty$. The region shrinks as σ decreases, converging to the diagonal $\underline{\delta} = \bar{\delta}$, as $\sigma \rightarrow 0$. A higher σ implies a faster effective speed of learning relative to discounting, allowing the agent to update his beliefs more quickly. As a result, for any given contract, the upper threshold will rise and the lower threshold will fall. This enlarges the due diligence region and directly expands the feasible set. The economic force behind this result relies on the reduction of the opportunity cost of experimentation. Perhaps counterintuitively, the bottom plots show that the region $\mathcal{X}(1)$ of implementable thresholds $(\underline{\delta}, \bar{\delta})$ under a fully bundled contract expands as the agent's execution cost $\mathcal{C}_A$ increases. From Lemma 4, it is easy to see that the agent execution cost expands the feasibility set in an analogous way to the bundling parameter. This happens because a larger cost makes executing the project less attractive as the agent's expected utility falls with $\mathcal{C}_A$.

For a pair of implementable thresholds $(\underline{\delta}, \bar{\delta})$ in the sense of Proposition 3, there are infinitely many compensation schemes $(\mathcal{W}_\emptyset, \mathcal{W}_0, \mathcal{W}_1)$ that implement $(\underline{\delta}, \bar{\delta})$. These are precisely the schemes for which $(\mathcal{W}_0, \mathcal{W}_1) \in \mathcal{I}$, as defined in (14). Among these, the principal selects the one that minimizes the expected compensation to the agent.

Proposition 4. *For a given $\alpha \in [0, 1]$, let $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(\alpha)$ be a pair of implementable thresholds in the sense of Proposition 3. Let $\mathcal{W}^*(\alpha) = (\mathcal{W}_\emptyset^*(\alpha), \mathcal{W}_0^*(\alpha), \mathcal{W}_1^*(\alpha))$ denote an optimal compensation scheme net of execution cost that implements $(\underline{\delta}, \bar{\delta})$. Then,*

$$\mathcal{W}_0^*(\alpha) = \mathcal{A} \mathcal{W}_\emptyset^*(\alpha) + (\mathcal{A} - 1) \hat{\omega}, \quad \mathcal{W}_1^*(\alpha) = \mathcal{B} \mathcal{W}_\emptyset^*(\alpha) + (\mathcal{B} - 1) \hat{\omega},$$

and the value of $\mathcal{W}_\emptyset^*(\alpha)$ is given by

$$\mathcal{W}_\emptyset^*(\alpha) = \mathbb{I}(\mathcal{B} > \mathcal{L} \mathcal{A}) \left[\left(\frac{1 - \mathcal{L}}{\mathcal{B} - \mathcal{L} \mathcal{A}} \right) (\hat{\omega} - \alpha \mathcal{C}_A) - \hat{\omega} \right]^+.$$

It is noteworthy that implementing a pair of thresholds $(\underline{\delta}, \bar{\delta})$ for which $\mathcal{L} \mathcal{A} < \mathcal{B} < \mathcal{L} \mathcal{A} + \frac{(1 - \mathcal{L})(\hat{\omega} - \mathcal{C}_A)}{\hat{\omega}}$ requires using a compensation scheme with $\mathcal{W}_\emptyset^*(\alpha) > 0$.

As noted in Remark 3, these types of compensation offer the agent a form of arbitrage that the principal must counterbalance by making the compensation when execution occurs more attractive. This poses a challenge for contracts with small values of the bundling parameter. In fact, implementing $(\underline{\delta}, \bar{\delta})$ with an information-only contract ($\alpha = 0$) is only possible if the condition $\mathcal{L}\mathcal{A} < \mathcal{B} < \mathcal{L}\mathcal{A} + \frac{(1-\mathcal{L})(\hat{\omega}-\mathcal{C}_A)}{\hat{\omega}}$ is satisfied, as

$$\mathcal{L}\mathcal{A} + (1-\mathcal{L}) \left(\frac{\hat{\omega} - \alpha \mathcal{C}_A}{\hat{\omega}} \right) = \mathcal{L}\mathcal{A} + (1-\mathcal{L}) > 1 > \mathcal{B}.$$

1 As we mentioned before, the exit payoff $\mathcal{W}_0$ serves as a tool for the principal to
 2 deal with the limited liability constraint. Specifically, certain combinations of belief
 3 thresholds $(\underline{\delta}, \bar{\delta})$ are only feasible if the agent receives strictly positive compensation
 4 upon abandonment. These are the pairs included in the feasibility set via the first
 5 set in $\mathcal{X}(\alpha)$ (see Proposition 3, when $\mathcal{B} > \mathcal{L}\mathcal{A}$).

6 Let's analyze the economic intuition behind the previous result and highlight the
 7 forces that enable new combinations of belief thresholds to be feasible. By marginally
 8 increasing the due diligence fee $\mathcal{W}_0$, the principal affects two key margins in the
 9 agent's problem. First, since the agent receives a higher payoff upon abandonment,
 10 the incentive to continue due diligence diminishes. For any fixed pair $(\mathcal{W}_0, \mathcal{W}_1)$,
 11 the agent now exits at a lower belief threshold, i.e., $\bar{\delta}$ decreases. Second, there is a
 12 direct effect on the lower threshold: as the abandonment payoff increases, the agent's
 13 continuation value rises (see Proposition 1). This makes the agent more willing to
 14 stay in the experimentation phase longer, thereby reducing $\underline{\delta}$. There is, however,
 15 a counteracting indirect effect: the reduction in $\bar{\delta}$ lowers the expected continuation
 16 payoff, which could push the agent to recommend execution sooner. But this force is
 17 second-order relative to the direct gain from $\mathcal{W}_0$, and thus the net effect on $\underline{\delta}$ remains
 18 negative. In conclusion, increasing $\mathcal{W}_0$ shifts the entire due diligence region leftward,
 19 concentrating experimentation in a more optimistic belief range and introduces new
 20 feasible pairs $(\underline{\delta}, \bar{\delta})$.

1 Proposition 4 has several important implications for designing an optimal contract.
 2 Perhaps the most striking insight arises when examining how the cost to the principal
 3 of inducing the agent to conduct due diligence depends on the probability of bundling,
 4 α . Specifically, for some threshold pairs $(\underline{\delta}, \bar{\delta})$, it may be more costly for the principal
 5 to offer a pure due diligence contract with $\alpha = 0$ than a fully bundled contract with
 6 $\alpha = 1$, even though the latter requires the agent to bear the execution cost fully.
 7 In such instances, the principal minimizes cost by incentivizing the agent to both
 8 conduct due diligence and, potentially, execute the project, rather than limiting the
 9 agent's role to due diligence alone.

To illustrate this point, suppose $\mathcal{C}_A < \hat{\omega}$ and consider a pair $(\underline{\delta}, \bar{\delta})$ such that
 $0 < \mathcal{B} - \mathcal{L}\mathcal{A} < \frac{(1-\mathcal{L})(\hat{\omega}-\mathcal{C}_A)}{\hat{\omega}}$. According to Proposition 4, the optimal compensation
 scheme satisfies

$$\mathcal{W}_\emptyset^*(\alpha) = \left(\frac{1-\mathcal{L}}{\mathcal{B}-\mathcal{L}\mathcal{A}} \right) (\hat{\omega} - \alpha \mathcal{C}_A) - \hat{\omega},$$

for all $\alpha \in [0, 1]$. Moreover, from Lemma 3, the principal's payment to the agent is
 given by

$$\mathcal{P}(\alpha) := \mathcal{W}_\emptyset + d^* ((\hat{\mathcal{V}} - 1) (\hat{\omega} + \mathcal{W}_\emptyset) + \alpha \mathcal{C}_A).$$

Differentiating with respect to α yields:

$$\frac{\partial \mathcal{P}(\alpha)}{\partial \alpha} = -(1 + d^* (\hat{\mathcal{V}} - 1)) \mathcal{C}_A \left(\frac{1-\mathcal{L}}{\mathcal{B}-\mathcal{L}\mathcal{A}} \right) + d^* \mathcal{C}_A < 0,$$

10 where the inequality holds for both $d^* = 0, 1$ and follows from the fact that $\hat{\mathcal{V}} > 1$ for
 11 $\underline{\delta} < \bar{\delta}$. Therefore, in this case, the principal minimizes the payment to the agent by
 12 choosing $\alpha^* = 1$.

However, setting $\alpha = 1$ is not necessarily optimal in general. In fact, in some cases,
 the principal may prefer to set α as small as possible. For instance, suppose $\hat{\omega} < \mathcal{C}_A$
 and consider a pair $(\underline{\delta}, \bar{\delta})$ such that $\mathcal{L}\mathcal{A} + \frac{(1-\mathcal{L})(\hat{\omega}-\mathcal{C}_A)}{\hat{\omega}} \leq \mathcal{B} \leq \mathcal{L}\mathcal{A}$. In this case, by
 (15), the pair $(\underline{\delta}, \bar{\delta})$ is implementable for all $\alpha \in [\bar{\alpha}, 1]$, with $\bar{\alpha} < 1$. Furthermore, the

optimal compensation scheme in Proposition 4 satisfies $\mathcal{W}_\emptyset^*(\alpha) = 0$ in this case. It follows that

$$\frac{\partial \mathcal{P}(\alpha)}{\partial \alpha} = d^* \mathcal{C}_A \geq 0,$$

1 and the principal is better off selecting the smallest feasible value of α that implements
2 $(\underline{\delta}, \bar{\delta})$, namely, $\alpha^* = \bar{\alpha} < 1$ in this case.

3 Below, in Proposition 5, we show that the two specific cases discussed above are
4 representative of the general structure of the optimal bundling parameter α^* .

5 5.4 Optimal Free-Term Contract

6 Equipped with the optimal compensation scheme $\mathcal{W}^*(\alpha) = (\mathcal{W}_\emptyset^*(\alpha), \mathcal{W}_0^*(\alpha), \mathcal{W}_1^*(\alpha))$
7 in Proposition 4, we now turn to the principal's problem of selecting an optimal
8 contract from the class $\mathbb{W}$ of partially bundled contracts. Under the parametrization
9 introduced in Section 5.2, this problem is formulated using the thresholds $(\underline{\delta}, \bar{\delta})$ and
10 the bundling parameter α as the decision variables.

11 To this end, the following result is useful, as it provides a set of necessary optimality
12 conditions that further restrict the possible values of the optimal thresholds $\underline{\delta}^*$ and $\bar{\delta}^*$,
13 as well as their associated compensation components $\mathcal{W}_0^*$, $\mathcal{W}_1^*$, and $\mathcal{W}_\emptyset^*$ in an optimal
14 contract.

15 **Lemma 5.** *If it is optimal for the principal to offer a contract $\mathcal{W} = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}), \alpha) \in \mathbb{W}$
16 that induces the agent to initially conduct due diligence, then it must hold that
17 $\mathcal{W}_1 < \mathcal{W}_\emptyset < \mathcal{W}_0$. Furthermore, the principal's expected payoff from executing a bad
18 project is negative under an optimal contract.*

19 According to Lemma 5, when inducing due diligence is optimal, the principal offers
20 a contract that results in a non-positive expected payoff if the project is bad. In-
21 tuitively, this optimality condition reflects the fact that any contract inducing the

1 agent to conduct due diligence while guaranteeing the principal a non-negative pay-
 2 off regardless of project complexity is dominated by the alternative of executing the
 3 project immediately at time $t = 0$ without any due diligence. Delaying the execution
 4 decision is costly because the principal must provide the agent with a learning rent,
 5 and the belief-updating process follows a martingale, implying that the principal's
 6 expected discounted profits will be lower.

7 The principal's realized payoff, net of compensation costs and the cost of executing
 8 the project, is given by

$$\begin{aligned}\Pi(\underline{\delta}, \bar{\delta}, \alpha) &:= d^* (\mathbb{E}_{\underline{\delta}}[\mathcal{R}] - (1 - \alpha)\mathcal{C}_P - ((1 - \underline{\delta})\mathcal{W}_0 + \underline{\delta}\mathcal{W}_1 + \alpha\mathcal{C}_A)) - (1 - d^*)\mathcal{W}_\emptyset^*(\alpha) \\ &:= d^* (\mathbb{E}_{\underline{\delta}}[\mathcal{R}] - \mathcal{C}_P - \alpha(\mathcal{C}_A - \mathcal{C}_P)) - \left(\mathcal{W}_\emptyset^*(\alpha) + d^* (\widehat{\mathcal{V}}(\underline{\delta}; \bar{\delta}) - 1) (\mathcal{W}_\emptyset^*(\alpha) + \widehat{\omega}) \right),\end{aligned}\tag{16}$$

9 where d^* is defined in Proposition 1 and $\widehat{\mathcal{V}}(\underline{\delta}; \bar{\delta})$ in (9).

10 We approach the optimization of $\Pi(\underline{\delta}, \bar{\delta}, \alpha)$ in two steps. First, for fixed values of
 11 $\underline{\delta}$ and $\bar{\delta}$, we determine the optimal bundling probability α^* . Then, we solve for the
 12 optimal thresholds $\underline{\delta}^*$ and $\bar{\delta}^*$.

13 From Proposition 4, the function $\mathcal{W}_\emptyset^*(\alpha)$ is piecewise linear in α , and therefore so
 14 is $\Pi(\underline{\delta}, \bar{\delta}, \alpha)$. This property induces a bang-bang behavior in the optimal bundling
 15 probability α^* , for fixed thresholds $\underline{\delta}$ and $\bar{\delta}$.

16 **Proposition 5.** *Let $\delta \in (0, 1)$ and $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(1)$ be a pair of implementable thresholds
 17 such that $\underline{\delta} < \delta < \bar{\delta}$. Recall the definition of $\bar{\alpha}$ in (15). Then:*

18 (i) *If $\mathcal{B} > \mathcal{L}\mathcal{A}$, then $\alpha^* = 1 \wedge \bar{\alpha}$.*

19 (ii) *If $\mathcal{B} \leq \mathcal{L}\mathcal{A}$, then $\alpha^* = 1 \wedge \bar{\alpha}$.*

20 In light of (15) and Proposition 5, the set of implementable threshold pairs $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(1)$
 21 can be partitioned into two regions: one in which the principal prefers to maximize

the value of α^* (case (i)), and another in which the principal chooses to minimize it (case (ii)).

Since, $\bar{\alpha} > 0$, an information-only contract, in which the agent is hired exclusively to conduct due diligence, is never optimal. Moreover, when $\hat{\omega} < \mathcal{C}_A$, we have $\bar{\alpha} < 1$, which implies $\alpha^* < 1$, and thus a fully bundled contract is never optimal. Therefore, in such cases, we have $0 < \alpha^* < 1$, and the principal will always offer the agent a contract that partially bundles due diligence and execution.

Figure 5 illustrates the result in Proposition 5. In each panel, the shaded region corresponds to the set of implementable pairs $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(1)$ satisfying $\underline{\delta} < \delta < \bar{\delta}$. The intensity of the shading represents the value of α^* , with light gray indicating values of α^* close to zero and dark gray indicating values close to one. The circular marker denotes the location of the optimal pair $(\underline{\delta}^*, \bar{\delta}^*)$ that maximizes the principal's expected payoff in (16).




Figure 5: Optimal value of randomization parameter $\alpha^*(\delta, \underline{\delta}, \bar{\delta})$ as defined in Proposition 5. The intensity of the shading represents the value of α^* , with light gray indicating values of α^* close to zero and dark gray indicating values close to one. The circular marker denotes the location of the optimal pair $(\underline{\delta}^*, \bar{\delta}^*)$ that maximizes the principal's expected payoff $\Pi^*(\delta)$ in (17). DATA: $\delta = 0.6$, $\mathcal{L} = 0.3$, $\gamma = 1.1$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, $\mathcal{C}_A = 4$, $\hat{\omega} = 1$ (left panel), and $\hat{\omega} = 5$ (right panel).

The probability that the execution is delegated to the agent is the result of a fundamental trade-off when the limited liability constraint binds. On the one hand,

1 increasing delegation relaxes the limited liability constraint and reduces the incentive
2 compatibility rent. This expands the set of implementable thresholds, which is maxi-
3 mum when the delegation is complete. On the other hand, delegation is costly, as the
4 principal is more efficient at execution than the agent. Thus, delegating execution
5 to the agent with positive probability, holding the expected compensation constant,
6 reduces the expected net return. Therefore, α provides the principal with the flexi-
7 bility to balance the incentives and feasibility for implementing the desired stopping
8 behavior against execution efficiency.

9 With a slight abuse of notation, let us define $\Pi(\underline{\delta}, \bar{\delta}) = \Pi(\underline{\delta}, \bar{\delta}, \alpha^*)$, where $\Pi(\underline{\delta}, \bar{\delta}, \alpha)$
10 is defined in (16) and α^* is the optimal bundling parameter in Proposition 5. The
11 principal's optimization problem reduces to

$$\Pi^*(\delta) = \sup_{\underline{\delta} < \delta < \bar{\delta}} \mathbb{E}_{\delta} [e^{-r\tau^*} \Pi(\underline{\delta}, \bar{\delta})] \quad \text{subject to} \quad (\underline{\delta}, \bar{\delta}) \in \mathcal{X}(1) \text{ and } \tau^* = \inf\{t > 0: \delta_t \notin (\underline{\delta}, \bar{\delta})\}. \quad (17)$$

12 In general, (17) must be solved numerically to determine the optimal values of $\underline{\delta}^*$ and
13 $\bar{\delta}^*$ as functions of the initial belief δ . A computational investigation of this problem
14 is presented in Section 6, where we study various properties of the optimal solution
15 and compare the optimal Free-Term with the optimal Fixed-Term contract.

16 6 Numerical Analysis

17 In this section, we first present an analysis of when to delegate under different learning
18 regimes. Second, we derive the optimal contract numerically for different instances of
19 the agent's outside payoff. Third, we compare welfare across contracts and between
20 them and two different benchmarks.

6.1 Due Diligence and Learning Speed

In this subsection, we examine how changes in the agent's learning speed affect the set of prior beliefs under which the principal hires the agent to conduct due diligence under both the Free-Term and Fixed-Term contracts.

Figure 6 depicts the region of prior beliefs for which the principal prefers each contract over immediate action (execution or abandonment), as a function of the signal precision parameter. The red region corresponds to the Free-Term contract, and the




Figure 6: Contract Design as Function of Signal Precision.
DATA: $\mathcal{L} = 0.3$, $\mathcal{R}_0 = 5$, $\mathcal{R}_1 = 0$, $\mathcal{C}_p = 2.5$, and $\mathcal{C}_A = 2.9$.

blue region represents the Fixed-Term contract.

Figure 6 also includes the first-best region (dashed line), which represents the set of priors under which it would be socially efficient to conduct due diligence. By construction, the set of implementable due diligence sets for the Fixed-Term and the Free-Term contract are subsets of the efficient one: whenever the principal finds it optimal to hire the agent under either the Fixed-Term or the Free-Term contract, it must also be optimal in the first-best allocation.

The figure highlights the key strengths and weaknesses of each contract, showing that neither dominates. Instead, the optimal choice depends on the interplay between belief dynamics, cost, and return properties. We identify three main frictions driving the results:

When learning is slow, hiring the agent is profitable only for the Fixed-Term contract.

1 In this regime, beliefs evolve slowly, and inducing the agent to continue learning under
2 the Free-Term contract is too expensive due to the long expected duration, relative to
3 the information gains. To induce the agent to participate, the principal would need to
4 offer a large compensation, making the Free-Term contract unprofitable. In contrast,
5 the Fixed-Term contract allows the principal to cap the cost of learning by choosing a
6 short fixed duration, since this contract doesn't require giving the agent a rent. The
7 principal benefits from the variance of the belief distribution at the end of the period,
8 even if the expected change in beliefs is small, the principal may benefit from the
9 tail outcomes that provide a clearer recommendation. In this sense, the Fixed-Term
10 contract exploits the second moment of belief evolution, while the Free-Term contract
11 depends critically on the first moment (i.e., the speed of convergence), which is too
12 slow in this setting.

13 When the agent starts with a very favorable belief, only the Fixed-Term contract
14 is profitable. Under a Free-Term contract, the principal would execute the project
15 immediately. In contrast, a short Fixed-Term contract introduces a delay, allowing
16 the belief to evolve before committing to execution, thereby enabling the principal to
17 obtain a recommendation for a sufficiently pessimistic posterior with positive proba-
18 bility and low costs, and subsequently abort the project. This gives the Fixed-Term
19 contract screening power even in optimistic environments, at a low cost. This behav-
20 ior stems from the discontinuity at $T = 0$ in the Fixed-Term contract discussed in
21 Section 4.

22 When learning is fast and the prior is pessimistic, only the Free-Term contract is
23 profitable for collecting information: when signals are highly informative, the agent's
24 belief evolves rapidly. Under the Free-Term contract, the agent can respond quickly
25 to the arrival of new information, abandoning the project if pessimism intensifies
26 or continuing to experiment if the signals become more favorable. The contract
27 effectively filters out bad trajectories while preserving upside potential, and does so
28 without a long expected due diligence period. In contrast, under the Fixed-Term

contract, the principal commits to paying the agent for a fixed time, regardless of whether the belief improves. When the prior is already pessimistic, the belief is likely to remain low throughout the fixed period, leading to execution delays with little updates. This makes the Fixed-Term contract inefficient: the principal pays for learning but doesn't get meaningful screening.

6.2 The Optimal Contract

In this subsection, we compute the principal's equilibrium profit π under both the Fixed-Term and Free-Term contract, the induced thresholds $\underline{\delta}$ and $\bar{\delta}$ as functions of the agent's outside payoff $\hat{\omega}$.

First, we consider the optimality as a function of the agent's outside payoff $\hat{\omega}$ (see Figure 7a and 7b). The solid lines correspond to the Free-Term contract, while the dot-dashed lines represent the Fixed-Term contract.

A priori, it is not evident whether the principal achieves higher profits under the Free-Term or Fixed-Term contract. The Free-Term contract allows the principal to control the quality of information better, but doing so comes at a cost of providing the agent with an information-acquisition and limited liability rent. The Fixed-Term contract controls the time better and does not require giving the agent a rent, but the principal cannot control the quality of information.

As $\hat{\omega}$ increases, the profit gap between the two contracts narrows and eventually changes sign. The Free-Term contract dominates when the outside option is small, since informational rents are limited and the contract can closely approximate the first-best. As $\hat{\omega}$ rises, however, informational rents grow and make the Free-Term contract increasingly costly, while the Fixed-Term contract—whose compensation only reflects the opportunity cost—becomes more attractive. Consequently, the Fixed-Term contract dominates for high values of $\hat{\omega}$.

Under either contract, a higher learning speed allows the principal to obtain the




Figure 7: Optimal Due Diligence Region in Free-Term Contract.

DATA: Panel (a) $\delta = 0.5$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, and $\mathcal{C}_A = 4$. Panel (b) $\delta = 0.65$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 2.5$, and $\mathcal{C}_A = 2.8$. Panel (c) $\delta = 0.5$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, and $\hat{\omega} = 0.5$. Panel (d) $\hat{\omega} = 1.25$, $\mathcal{L} = 0.3$, $r = 0.05$, $\sigma = 0.7$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, and $\mathcal{C}_A = 3.25$

1 same level of informational content at a lower compensation cost, since less time is
2 required. This can be interpreted as a reduction in the marginal cost of learning.
3 As a result, the principal is incentivized to expand screening (i.e., encourage more
4 learning). Overall, the improvement in learning speed leads to higher expected profits
5 under both contracts. Figure 7a shows that as the learning speed rises, the threshold
6 for the agent's outside option below which the Free-Term contract dominates the
7 Fixed-Term contract rises. This shows that higher learning speed favors Free-Term
8 contracts. However, Figure 7b shows a non-monotonic relationship -the threshold rises
9 first and then falls. Thus, the relationship between learning speed and the profitability
10 difference between contracts is non-monotonic. Finally, as shown in Figure 6, there is
11 a region of small σ in which only Fixed-Term profitably implements screening. Low

1 learning speed diminishes the option value of adaptive stopping under Free-Term,
 2 while Fixed-Term caps costs by fixing the investigation horizon.

3 Figure 7c shows the relationship between the agent's execution cost and profits.
 4 Changes in $\mathcal{C}_A$ have a straightforward effect on profitability under the Fixed-Term
 5 contract because the principal's payoff is independent of $\mathcal{C}_A$. In contrast, under the
 6 Free-Term contract, her payoff decreases with the agent's cost. Thus, as $\mathcal{C}_A$ rises, the
 7 threshold $\hat{\omega}$ above which the profits of the Free-Term contract exceed those of the
 8 Fixed-Term contract falls. Whenever $\alpha^* < 1$, an increase in $\mathcal{C}_p$ favors the Free-Term
 9 contract since its weight on profits under the Free-Term contract is $1 - \alpha^* < 1$ and
 10 under the Fixed-Term contract is 1.

11 Figure 7d shows the relationship between profits and the prior δ . Higher optimism
 12 (lower δ) makes Fixed-Term relatively more attractive: the principal can set a short
 13 due diligence window T that screens for adverse evidence at low cost, compensating
 14 only the agent's opportunity cost (no informational rents). By contrast, under more
 15 pessimistic priors, Free-Term is preferred because its compensation scheme supports
 16 state-contingent learning: continuation occurs only along paths with favorable signals
 17 and is stopped immediately otherwise. For intermediate priors, the numerical results
 18 suggest that the profit difference is monotone in δ .

19 Based on the numerical analysis conducted previously, we draw the following obser-
 20 vations.

21 **Observation 1.** *There exists a threshold for $\hat{\omega}$, $\mathcal{C}_A$, and δ such that the Free-Term is*
 22 *optimal when $\hat{\omega}$ is higher than, $\mathcal{C}_A$ is lower than, and δ is higher than the corresponding*
 23 *threshold.*

24 Figure 8 shows that as $\hat{\omega}$ increases, the principal chooses a narrower due diligence re-
 25 gion - i.e., a smaller $\bar{\delta} - \underline{\delta}$ - under the Free-Term contract. This occurs because a higher
 26 $\hat{\omega}$ raises the informational rent associated with any given $\bar{\delta}$ (specifically, $\mathcal{V}(\delta_0; \bar{\delta}) - \hat{\omega}$,
 27 as shown in (10)). To limit these rents, the principal reduces $\bar{\delta}$. However, lowering the

1 upper threshold decreases the probability of execution, which the principal partially
 2 offsets by increasing $\underline{\delta}$. In addition, a higher σ corresponds to faster effective learning,
 3 making due diligence more informative. Consequently, the value of experimentation
 4 rises since the principal's cost of eliciting the same information under either contract
 5 is now smaller. Thus, the principal provides the agent with more powerful incentives
 6 to acquire information under the Free-Term contract, and thereby the optimal $\bar{\delta} - \underline{\delta}$
 7 rises with the learning speed.




Figure 8: Optimal Due Diligence Region in Free-Term Contract.
 DATA: $\delta = 0.5$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, and $\mathcal{C}_A = 4$.

8 In Figure 9, we plot the optimal expected duration of due diligence and the probability
 9 that the project will be implemented under the Fixed-Term and Free-Term contracts
 10 against the agent's outside option.




Figure 9: Implementation Probability and Expected Due-Diligence Duration.
 DATA: $\delta = 0.5$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, and $\mathcal{C}_A = 4$.

11 The optimal expected due diligence duration decreases with the agent's outside op-

tion under both contracts as the hiring costs increase with $\hat{w}$ and with the learning speed as a consequence of the improvement in the agent's efficiency. Interestingly, the Fixed-Term contract leads to longer expected due diligence times. This outcome reflects the fact that, under the fixed contract, the principal can afford to specify more extended experimentation periods without incurring larger incentive costs. Since there are no informational rents to pay and the agent's effort is predetermined, extending the duration is relatively cheap. In contrast, the Free-Term contract requires compensating the agent for endogenous stopping, which becomes more costly the longer the continuation region. As a result, the principal has more powerful incentives to limit the expected duration under the adaptive regime.

The probability that the project will be implemented decreases with the agent's outside option for the Free-Term contract. This is explained by the changes in the optimal thresholds discussed above, where the reduction in $\bar{\delta}$ dominates over the increment in $\underline{\delta}$.

The effect of learning speed works in the opposite direction: an increase in σ reduces the cost of learning by shortening the time required. As a result, the optimal level of learning rises, reflected in a higher $\bar{\delta}$ and a lower $\underline{\delta}$. The effect on $\underline{\delta}$ dominates, leading to a lower probability of execution; however, the expected benefit from execution increases, resulting in a higher overall expected profit.

In contrast, the probability that the idea is implemented increases for the Fixed-Term contract. Because the rise in hiring costs implies a reduction in the optimal due diligence duration, and given that the initial parameters are such that the principal prefers to execute immediately rather than abandon, the likelihood of being inside the execution region decreases as the due diligence time increases. Furthermore, it falls with σ at any given outside option level because the larger the learning speed, the shorter the expected duration of due diligence.

6.3 Welfare

In this sub-section, we examine how the welfare provided by the optimal Fixed-Term and Free-Term contracts varies with the agent's execution costs, his outside option, and learning speed. To gain a better understanding of the welfare losses under these two different contracts, we present two benchmarks against which to compare our results before discussing the numerical exercises.

The first benchmark corresponds to the case in which a benevolent central planner chooses the optimal stopping time and execution under perfect and complete information. Because there are no informational or contractual frictions, the agent is compensated according to his discounted outside option, $\mathcal{W} = \widehat{\omega}(e^{r\tau} - 1)$, and the planner chooses the optimal stopping and execution policy to maximize total surplus.

Let's define the total surplus $S(\delta, \alpha) =: \max\{\mathbb{E}_{\delta_\tau}[\mathcal{R}] - \mathcal{C}_P - \alpha(\mathcal{C}_A - \mathcal{C}_P) + \widehat{\omega}, \widehat{\omega}\}$. Then, the social planner solves the following optimal stopping problem,

$$\mathcal{S}(\delta) = \sup_{\tau \in \mathbb{T}, \alpha \in [0,1]} \mathbb{E}_\delta [e^{-r\tau} S(\delta_\tau, \alpha)] \quad \text{subject to} \quad d\delta_t = \delta_t(1 - \delta_t)\sigma dB_t \quad \text{and} \quad \delta_0 = \delta. \quad (18)$$

It readily follows that the principal chooses $\alpha^{FB} = 0$ and its expected discounted total surplus is given by

$$\mathcal{S}(\delta) = \begin{cases} (1 - \delta)(\mathcal{R}_0 + \widehat{\omega} - \mathcal{C}_P) + \delta(\mathcal{R}_1 + \widehat{\omega} - \mathcal{C}_P) & \text{if } 0 \leq \delta \leq \underline{\delta}^{FB} \\ \widehat{\omega} \widehat{\mathcal{V}}(\delta; \bar{\delta}^{FB}) & \text{if } \underline{\delta}^{FB} < \delta < \bar{\delta}^{FB} \\ \widehat{\omega} & \text{if } \bar{\delta}^{FB} \leq \delta \leq 1, \end{cases} \quad (19)$$

where the thresholds $\underline{\delta}^{FB}$ and $\bar{\delta}^{FB}$ are determined imposing value-matching ($\mathcal{S}(\delta) = S(\delta, 0)$)

and smooth-pasting ($\mathcal{S}(\delta) = S_\delta(\delta, 0)$) conditions at $\delta = \underline{\delta}^{FB}$ and $\delta = \bar{\delta}^{FB}$, and satisfy

$\underline{\delta}^{FB} < (\mathcal{C}_P - \mathcal{R}_1)/(\mathcal{R}_0 - \mathcal{R}_1) < \bar{\delta}^{FB}$. The benevolent planer's optimal strategy

(τ^{FB}, d^{FB}) is given by

1 $\tau^{FB} = \inf\{t > 0: \delta_t \notin (\underline{\delta}^{FB}, \bar{\delta}^{FB})\}$ and $d^{FB} = \mathbb{1}(\delta_{\tau^{FB}} \leq \underline{\delta}^{FB})$.

2 Thus, the social planner assigns the project execution to the agent with probability
3 zero and chooses a high threshold $\bar{\delta}^{FB} > \bar{\delta}^*$ and a low $\underline{\delta}^{FB} < \underline{\delta}^*$. This happens
4 because there is no need to provide the agent with a due diligence rent or a limited
5 liability rent. Thus, the social planner fully appropriates the whole surplus.

6 A second benchmark worth studying is one in which the principal observes the agent's
7 belief updating process, i.e., there is no private information. The principal can choose
8 $(\underline{\delta}, \bar{\delta})$ directly since, in the absence of private information concerning the evolution of
9 the prior, she can induce the agent to stop whenever she wants. To do so, the contract
10 must be such that the agent does not exercise his outside option while conducting
11 due diligence until the principal instructs the agent to stop. The way to prevent the
12 agent from stopping due diligence any time before or after the principal wishes the
13 agent to stop is to pay him a fixed wage equal to $\mathcal{W} = \hat{\omega} (e^{r\tau} - 1)$. Because this is
14 positive for all τ , it satisfies limited liability.

15 The principal's optimal stopping problem is identical to the central planner's, up to a
16 constant $-\hat{\omega}$, the principal never delegates the execution to the agent and implements
17 the socially optimal thresholds, i.e., $\bar{\delta}^{SB} = \bar{\delta}^{FB}$ and $\underline{\delta}^{SB} = \underline{\delta}^{FB}$. The principal chooses
18 these thresholds since she can extract the whole surplus minus the agent's expected
19 present value of the opportunity cost $\hat{\omega}$ of conducting due diligence.

20 It is worthwhile noting that the first-best contract shares the exact fundamental
21 nature as the Free-Term contract. Namely, in both cases, the optimal stopping rule
22 takes the form of a stopping time defined by two belief thresholds since they rely
23 on the same type of adaptive experimentation rule. Nevertheless, because the agent
24 does not internalize the principal's marginal revenues and costs, his privately optimal
25 decision under Free-Term contracts is inefficient. Under certain parameterizations,
26 this provides the principal's incentives to offer a Fixed-Term contract, which is of a
27 different nature from the Free-Term contract. This is done to control the duration

1 of experimentation so as to minimize the compensation cost at the cost of worsening
 2 the quality of information.

3 The previous discussion shows that the inefficiencies of the Free-Term contract come
 4 from the agent's private observation of the prior belief evolution. When the agent's
 5 learning is private, he must be given a due diligence rent to prevent him from taking
 6 his outside payoff plus the fixed abandonment payment immediately, and to con-
 7 duct due diligence until either of the principal's optimal thresholds is hit. The
 8 principal delegates the execution with positive probability to the agent to lower the
 9 information-acquisition rent and to enlarge the set of implementable thresholds. As
 10 a result, the equilibrium entails an inefficiently low level of due diligence, resulting in
 11 lower welfare compared to the first-best and second-best outcomes.




Figure 10: Welfare.

DATA: Panel (a) $\delta = 0.5$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 3$, and $\mathcal{C}_A = 4$.

Panel (b) $\delta = 0.65$, $\mathcal{L} = 0.3$, $r = 0.05$, $\mathcal{R}_0 = 7$, $\mathcal{R}_1 = 0$, $\mathcal{C}_P = 2.5$, and $\mathcal{C}_A = 2.8$

12 In Figure 10, we compute total welfare for different levels of the agent's outside option
 13 and learning speeds for both the Fixed-Term and Free-Term contracts. Under either
 14 contract, welfare tends to rise with the agent's outside option. In both contracts, the
 15 frictions have the same source: on one hand, the learning cost increases, which leads
 16 to a decline in the due diligence duration, diminishing the expected revenues from the
 17 project. On the other hand, the higher outside option increases the surplus generated
 18 once the agent finishes learning, which positively affects the total surplus.

19 From Figure 7 and 10, we can see that the contract that maximizes profit does

1 not maximize the total surplus. For low values of the outside option, the Free-Term
 2 contract generates higher welfare. This result reflects the Free-Term contract's ability
 3 to tailor experimentation to the agent's posterior belief in real time, enabling more
 4 efficient project screening. Because the stopping time is endogenous, the agent ends
 5 due diligence precisely when the belief crosses a critical threshold, thereby avoiding
 6 both wasteful continuation and premature exit. In contrast, the Fixed-Term contract
 7 imposes a rigid schedule that may lead to misaligned continuation decisions. In Figure
 8 10a, when the agent's outside option is high, the Fixed-Term maximizes welfare. In
 9 that case, the welfare loss from inefficient stopping is outweighed by the savings from
 10 avoiding the inefficient delegation. In contrast, in Figure 10b, the Free-Term contract
 11 dominates the Fixed-Term one for all outside option values due to that the difference
 12 in execution costs is small.

13 **7 Concluding Remarks**

14 This paper develops a dynamic principal–agent model of private information acqui-
 15 sition in which an expert gathers unobservable signals about a binary-state project and
 16 ultimately recommends whether to pursue or abandon it. We analyze two contract
 17 types: Fixed-Term contracts, which specify a predetermined investigation period, and
 18 Free-Term contracts, which tie compensation to the decision to abandon the project
 19 and the project's outcome upon execution. The principal can also allocate execution
 20 rights to the agent as part of the contract.

21 We show that neither contract class uniformly dominates. Fixed-Term contracts are
 22 simple to implement and do not require informational rents or inefficient task alloca-
 23 tion, making them optimal when the agent's outside option is high and learning speed
 24 is low. In contrast, Free-Term contracts allow greater flexibility and are better suited
 25 when the agent's outside option is small and learning speed is high. However, they
 26 often require paying incentive-compatible rents and suffer from implementation prob-

1 lems that necessitate delegating execution to the agent, despite his higher execution
2 cost.

3 A key and novel result is that the optimal Free-Term contract demands bundling
4 information acquisition with project execution, which helps mitigate the incentive
5 problems and implement a larger set of information acquisition qualities, despite
6 the inefficiency of delegating execution. This mechanism highlights a novel trade-off
7 between productive efficiency and incentive alignment.

8 Several avenues merit further exploration. First, future work could examine settings
9 where agents can influence the quality or rate of signal arrival through their own
10 efforts, introducing a second dimension of moral hazard. Second, consider contracts
11 that can condition payments on both the project's outcome and the time the project
12 is implemented. This will entail dealing with time-dependent thresholds and solving a
13 stochastic partial differential equation. Finding analytical (closed-form) solutions to
14 them is often challenging, and many do not have explicit solutions. Third, extending
15 the model to multi-agent settings, where multiple experts simultaneously acquire
16 information or compete for contracts, could discipline agents and yield insights into
17 screening, collusion, or information aggregation.

1 A Proofs

PROOF OF LEMMA 1: First, note that

$$\mathbb{P}_\delta(\delta_T \leq x) = (1 - \delta) \mathbb{P}_0(\delta_T \leq x) + \delta \mathbb{P}_1(\delta_T \leq x).$$

From Section 3, the belief process satisfies

$$\delta_T = \frac{\delta}{\delta + (1 - \delta) \exp\left(\sigma^2 \left(\frac{T}{2} - X_T\right)\right)}, \quad \text{where } X_T = \theta T + \frac{W_T}{\sigma}.$$

2 Conditional on the value of θ , the term $\sigma^2 \left(\frac{T}{2} - X_T\right)$ is normally distributed with
 3 mean $\mu_\theta = \left(\frac{1}{2} - \theta\right) \sigma^2 T$ and variance $\sigma_\theta^2 = \sigma^2 T$. From this, the quantities $\mathbb{P}_0(\delta_T \leq x)$
 4 and $\mathbb{P}_1(\delta_T \leq x)$ can be computed directly in terms of the standard normal cumulative
 5 distribution functions Φ_θ to recover the expression for $\mathbb{P}_\delta(\delta_T \leq x)$ in the statement of
 6 the lemma. $\square$

7 PROOF OF LEMMA 2: From Lemma 1 it follows that $\mathbb{P}_\delta(\delta_T \leq x)$ is decreasing in δ
 8 for all $T \geq 0$ and $x \in (0, 1)$, which implies that $\Pi^F(\delta)$ is also decreasing in δ . Thus,
 9 for $\delta \geq \hat{\delta}$, we have $\Pi^F(\hat{\delta}) \geq \Pi^F(\delta)$. We also have $\pi(\hat{\delta}) = \pi(\delta) = 0$. Thus, for $\delta \geq \hat{\delta}$, we
 10 conclude that $\Pi^F(\hat{\delta}) - \pi(\hat{\delta}) \geq \Pi^F(\delta) - \pi(\delta)$ as required.

11 Consider the case $\delta < \hat{\delta}$. In this case, we have that $\pi(\hat{\delta}) - \pi(\delta) = (\hat{\delta} - \delta) (\Pi_1 - \Pi_0)$
 12 and the inequality $\Pi^F(\hat{\delta}) - \pi(\hat{\delta}) \geq \Pi^F(\delta) - \pi(\delta)$ is equivalent to $\frac{\Pi^F(\hat{\delta}) - \Pi^F(\delta)}{\hat{\delta} - \delta} \geq \Pi_1 - \Pi_0$.
 13 To prove this inequality we show that (i) $\Pi^F(\delta)$ is convex in δ and the right derivative
 14 of $\Pi^F(\delta)$ at $\delta = 0$ is greater than or equal to $\Pi_1 - \Pi_0$.

15 To prove the convexity, let $H(T) = e^{-rT} \left(\pi(\delta_T) - \hat{\omega}(e^{rT} - 1) \right)$ and $\delta = \alpha \delta_1 + (1 - \alpha) \delta_2$

1 for $\alpha \in [0, 1]$ and $\delta_1, \delta_2 \in (0, 1)$. We have

$$\begin{aligned}
\mathbb{E}_\delta[H(T)] &= (1 - \delta) \mathbb{E}_0[H(T)] + \delta \mathbb{E}_1[H(T)] \\
&= (1 - \alpha \delta_1 - (1 - \alpha) \delta_2) \mathbb{E}_0[H(T)] + (\alpha \delta_1 + (1 - \alpha) \delta_2) \mathbb{E}_1[H(T)] \\
&= \alpha ((1 - \delta_1) E_0[H(T)] + \delta_1 E_1[H(T)]) + (1 - \alpha) ((1 - \delta_2) E_0[H(T)] + \delta_2 E_1[H(T)]) \\
&= \alpha \mathbb{E}_{\delta_1}[H(T)] + (1 - \alpha) \mathbb{E}_{\delta_2}[H(T)].
\end{aligned}$$

2 It follows that

$$\begin{aligned}
\Pi^F(\delta) &= \max_{T \geq 0} \mathbb{E}_\delta[H(T)] = \max_{T \geq 0} \left\{ \alpha \mathbb{E}_{\delta_1}[H(T)] + (1 - \alpha) \mathbb{E}_{\delta_2}[H(T)] \right\} \\
&\leq \alpha \max_{T \geq 0} \mathbb{E}_{\delta_1}[H(T)] + (1 - \alpha) \max_{T \geq 0} \mathbb{E}_{\delta_2}[H(T)] \\
&= \alpha \Pi^F(\delta_1) + (1 - \alpha) \Pi^F(\delta_2),
\end{aligned}$$

3 which shows the convexity of $\Pi^F(\delta)$.

4 Finally, note that for δ small enough $\Pi^F(\delta) \geq \pi(\delta) = (1 - \delta) \Pi_0 + \delta \Pi_1$. Since,
5 $\Pi^F(0) = \Pi_0$, the right-hand derivative of $\pi(\delta)$ at $\delta = 0$ is a lower bound for the
6 right-hand derivative of $\Pi^F(\delta)$ at $\delta = 0$. Thus, we conclude that it is greater than or
7 equal to $\Pi_1 - \Pi_0$, which completes the proof. $\square$ to show that the right derivative of
8 $\Pi^F(\delta)$ at $\delta = 0$ is greater than or equal to $\Pi_1 - \Pi_0$, it follows that $\frac{\Pi^F(\delta) - \Pi^F(0)}{\delta} \geq \Pi_1 - \Pi_0$,
9 which completes the proof. $\square$

10 **PROOF OF THEOREM 1:** For an $f \in \widehat{\mathcal{C}}^2$ that solves (QVI) we have

$$\begin{aligned}
e^{-r\tau} f(\delta_\tau) &= f(\delta) + \int_0^\tau e^{-rt} \mathcal{H}f(\delta_t) dt + \int_0^\tau e^{-rt} \sigma \delta_t (1 - \delta_t) f'(\delta_t) dB_t \\
&\leq f(\delta) + \int_0^\tau e^{-rt} \sigma \delta_t (1 - \delta_t) f'(\delta_t) dB_t,
\end{aligned}$$

11 where the equality follows from Itô's lemma and the inequality follows from the fact
12 that $\mathcal{H}f(\delta) \leq 0$ (second QVI condition). Taking expectation and canceling the

1 stochastic integral, we get $\mathbb{E}[e^{-r\tau} f(\delta_\tau)] \leq f(\delta)$. From the first QVI condition it
 2 follows that $\mathbb{E}[e^{-r\tau} V(\delta_\tau)] \leq \mathbb{E}[e^{-r\tau} f(\delta_\tau)] \leq f(\delta)$. Taking the supreme over all stop-
 3 ping times $\tau \geq 0$, we conclude that $f(\delta) \geq \mathcal{V}(\delta)$. Finally, all the inequalities above
 4 become equalities for the QVI-control associated to f . This follows from Dynkin's
 5 formula (see [Øksendal, 2013](#)) and the fact that the QVI-control is the first exit time
 6 from the continuation region $\mathcal{C}$. $\square$

PROOF OF PROPOSITION 1: Let $\mathcal{V}(\delta)$ be the function defined in (10). We will show
 that $\mathcal{V}(\delta)$ satisfies the (QVI) optimality conditions and so by Theorem 1 it is equal to
 the agent's optimal expected payoff. To this end, note that $\mathcal{V}(\delta) \in \widehat{\mathcal{C}}^2$, which follows
 from the smooth-pasting and value-matching conditions. Also, as we show below, the
 function $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ satisfies the ODE in (7) so it follows that

$$(\mathcal{H}\mathcal{V})(\delta) = \begin{cases} -r V(\delta) & \text{if } 0 \leq \delta < \underline{\delta} \\ 0 & \text{if } \underline{\delta} < \delta < \bar{\delta} \\ -r V(\delta) & \text{if } \bar{\delta} < \delta \leq 1. \end{cases}$$

7 From this, and the definition of $\mathcal{V}(\delta)$, it follows that $(\mathcal{H}\mathcal{V})(\delta) \leq 0$ and $(\mathcal{V}(\delta) - V(\delta)) (\mathcal{H}\mathcal{V})(\delta) = 0$
 8 for all $\delta \in [0, 1] \setminus \{\underline{\delta}, \bar{\delta}\}$. Thus, $\mathcal{V}(\delta)$ satisfies the second and third (QVI) condition.

Next, we show the existence and uniqueness of a function $\mathcal{V}(\delta)$ satisfying the condition
 in the proposition. Recall that $F(\delta) = (1 - \delta)^\gamma \delta^{1-\gamma}$ in (8) solves the ODE in (7).
 Let us also recall the definition an auxiliary function $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ in (9) and for notational
 convenience, let us extend the domain of its first argument to $\delta \in (0, 1)$ so that

$$\widehat{\mathcal{V}}(\delta; \bar{\delta}) = \begin{cases} A_0(\bar{\delta}) F(\delta) + A_1(\bar{\delta}) F(1 - \delta) & \text{if } 0 < \delta < \bar{\delta} \\ 1 & \text{if } \bar{\delta} \leq \delta \leq 1, \end{cases}$$

where the constants

$$A_0(\bar{\delta}) = \frac{(\gamma - \bar{\delta})}{(2\gamma - 1)F(\bar{\delta})} \quad \text{and} \quad A_1(\bar{\delta}) = \frac{(\gamma + \bar{\delta} - 1)}{(2\gamma - 1)F(1 - \bar{\delta})}$$

- 1 are chosen so that $\hat{\mathcal{V}}(\delta)$ is continuously differentiable at $\delta = \bar{\delta}$. Since $\gamma > 1$ it follows
- 2 that $A_0(\bar{\delta})$ and $A_1(\bar{\delta})$ are both positive for $\bar{\delta} \in (0, 1)$. Furthermore, $A_0(\bar{\delta}) \uparrow \infty$ as
- 3 $\bar{\delta} \uparrow 1$ and $A_1(\bar{\delta}) \uparrow \infty$ as $\bar{\delta} \downarrow 0$.

Besides being continuously differentiable in $(0, 1)$ by construction, the $\hat{\mathcal{V}}(\delta; \bar{\delta})$ is also decreasing and strictly convex in the region $\delta \in (0, \bar{\delta})$. To see this, note that in this region $\hat{\mathcal{V}}(\delta; \bar{\delta})$ satisfies the differential equation (7) and so

$$\frac{(\sigma\delta(1-\delta))^2}{2} \hat{\mathcal{V}}''(\delta; \bar{\delta}) - r \hat{\mathcal{V}}(\delta; \bar{\delta}) = 0 \quad \implies \quad \hat{\mathcal{V}}''(\delta; \bar{\delta}) = \frac{2r \hat{\mathcal{V}}(\delta; \bar{\delta})}{(\sigma\delta(1-\delta))^2} \geq \frac{2r}{(\sigma\delta(1-\delta))^2} > 0.$$

- 4 To complete the proof, we will show that there exists a value of $\bar{\delta} > \hat{\delta} := (\mathcal{W}_0 - \mathcal{W}_\emptyset)/(\mathcal{W}_0 - \mathcal{W}_1)$
- 5 such that the associated value function $(\hat{\omega} + \mathcal{W}_\emptyset) \hat{\mathcal{V}}(\delta; \bar{\delta})$ satisfies value-matching and
- 6 smooth-pasting conditions with the function $(1 - \delta)(\hat{\omega} + \mathcal{W}_0) + \delta(\hat{\omega} + \mathcal{W}_1)$ at some $\underline{\delta} < \hat{\delta}$. Figure 11 illustrates the argument. On panel (a), the value of $\bar{\delta} = 0.7$ is too low




Figure 11: Function $\hat{\mathcal{V}}(\delta, \bar{\delta})$ for three values of $\bar{\delta}$. DATA: $r = 0.1$, $\sigma = 1$, $\hat{\omega} = 0$, $\mathcal{W}_\emptyset = 100$, $\mathcal{W}_0 = 200$ and $\mathcal{W}_1 = 20$.

- 7
- 8 and the function $(\hat{\omega} + \mathcal{W}_\emptyset) \hat{\mathcal{V}}(\delta; \bar{\delta})$ intersects with the function $V(\delta)$ in a non-smooth
- 9 way in the region $(0, \bar{\delta}]$. On the flip side, on panel (c), the value of $\bar{\delta} = 0.9$ is too high

1 and the function $(\hat{\omega} + \mathcal{W}_0) \hat{\mathcal{V}}(\delta; \bar{\delta})$ does not intersect at all the function $V(\delta)$ in the re-
 2 gion $(0, \bar{\delta}]$. Finally, on panel (b), the value of $\bar{\delta} = 0.76406$ is such that $(\hat{\omega} + \mathcal{W}_0) \hat{\mathcal{V}}(\delta; \bar{\delta})$
 3 intersects smoothly the function $V(\delta)$ at $\underline{\delta} = 0.34103$.

4 Mathematically, the argument combines the following facts:

5 (i) The function $\hat{\mathcal{V}}(\delta; \bar{\delta})$ is monotonically decreasing and strict convex in $(0, \bar{\delta}]$ as
 6 argued above.

7 (ii) The function $V(\delta)$ is piece-wise linear in $(0, 1)$.

8 (iii) $\hat{\mathcal{V}}(\delta; \bar{\delta})$ is monotonic in $\bar{\delta}$, that is, $\hat{\mathcal{V}}(\delta; \bar{\delta}_1) \leq \hat{\mathcal{V}}(\delta; \bar{\delta}_2)$ for $\bar{\delta}_1 \leq \bar{\delta}_2$.

9 (iv) For all $\bar{\delta} \in (0, 1)$ we have that $\hat{\mathcal{V}}(\delta; \bar{\delta}) \uparrow \infty$ as $\delta \downarrow 0$.

10 (v) For $\bar{\delta}$ sufficiently large $(\hat{\omega} + \mathcal{W}_0) \hat{\mathcal{V}}(\delta; \bar{\delta}) > V(\delta)$ for all $\delta \in (0, \bar{\delta})$.

11 Point (iii) is obtained directly from the derivative:

$$\frac{\partial \hat{\mathcal{V}}(\delta; \bar{\delta})}{\partial \bar{\delta}} = \frac{(\gamma - \bar{\delta})(\gamma + \bar{\delta} - 1) - \bar{\delta}(1 - \bar{\delta})}{(2\gamma - 1)\bar{\delta}(1 - \bar{\delta})} \left(\frac{\delta}{\bar{\delta}}\right)^{1-\gamma} \left(\frac{1 - \delta}{1 - \bar{\delta}}\right)^{\gamma} \left[1 - \left(\frac{\delta}{\bar{\delta}}\right)^{2\gamma-1} \left(\frac{1 - \delta}{1 - \bar{\delta}}\right)^{1-2\gamma}\right] > 0$$

12 where the inequality holds as $\gamma > 1$ so $\gamma - \bar{\delta} > 1 - \bar{\delta}$ and $\gamma + \bar{\delta} - 1 > \bar{\delta}$ and the first
 13 factor is positive. In addition, as $\delta \geq \bar{\delta}$, the last term between brackets is not greater
 14 than 1.

15 Point (iv) follows from noticing that $F(0) \uparrow \infty$ as $\delta \downarrow 0$. Finally, (v) follows the fact
 16 that $A_0(\bar{\delta})$, which is non-negative and strictly increasing, grows unboundedly as $\bar{\delta} \uparrow 1$.

17 Combining points (i) and (iv) it follows that if $\bar{\delta} \leq \hat{\delta}$, the function $(\hat{\omega} + \mathcal{W}_0) \hat{\mathcal{V}}(\delta; \bar{\delta})$ will
 18 intersect $V(\delta)$ in a non-smooth way in the region $(0, \bar{\delta})$ as in panel (a) in Figure 11.

19 Thus, smooth-pasting can only be achieve if $\bar{\delta} > \hat{\delta}$. On the other hand, by point (v)
 20 for δ sufficiently large the function $(\hat{\omega} + \mathcal{W}_0) \hat{\mathcal{V}}(\delta; \bar{\delta})$ never intersects $V(\delta)$ in $(0, \bar{\delta})$ as
 21 in panel (c) in Figure 11 and so again there trivially no smooth-pasting in this region.

22 Thus, by the continuity $\hat{\mathcal{V}}(\delta; \bar{\delta})$ on δ and points (i) and (ii) there exists a $\bar{\delta}$ such that

- 1 $(\widehat{\omega} + \mathcal{W}_0) \widehat{\mathcal{V}}(\delta; \bar{\delta})$ intersects smoothly $V(\delta)$ in the region $(0, \bar{\delta})$. Finally, by point (iii)
 2 there is a unique $\bar{\delta} \in (\hat{\delta}, 1)$ for which $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ satisfies the smooth-pasting condition as
 3 in panel (b) in Figure 11. $\square$

PROOF OF LEMMA 3: According to Proposition 1, the connection between the thresholds $(\underline{\delta}, \bar{\delta})$ and the pair $(\mathcal{W}_0, \mathcal{W}_1)$ is given by the value-matching and smooth-pasting conditions imposed on the agent's value function $\mathcal{V}(\delta)$ at the values $\delta = \underline{\delta}$ and $\delta = \bar{\delta}$. Furthermore, in the region of due diligence $\delta \in (\underline{\delta}, \bar{\delta})$, the function $\mathcal{V}(\delta)$ is equal $\widehat{\mathcal{V}}(\delta; \bar{\delta})$ see (9) and (10). Thus, value-matching and smooth-pasting at $\underline{\delta}$ imply that

$$(\widehat{\omega} + \mathcal{W}_0) \widehat{\mathcal{V}}(\underline{\delta}; \bar{\delta}) = (1 - \underline{\delta})(\widehat{\omega} + \mathcal{W}_0) + \underline{\delta}(\widehat{\omega} + \mathcal{W}_1) \quad \text{and} \quad (\widehat{\omega} + \mathcal{W}_0) \widehat{\mathcal{V}}'(\underline{\delta}; \bar{\delta}) = \mathcal{W}_1 - \mathcal{W}_0.$$

Solving for $\mathcal{W}_0$ and $\mathcal{W}_1$ we get

$$\mathcal{W}_0 + \widehat{\omega} = \mathcal{A}(\underline{\delta}, \bar{\delta})(\mathcal{W}_0 + \widehat{\omega}) \quad \text{and} \quad \mathcal{W}_1 + \widehat{\omega} = \mathcal{B}(\underline{\delta}, \bar{\delta})(\mathcal{W}_0 + \widehat{\omega}),$$

- 4 where the values of $\mathcal{A}(\underline{\delta}, \bar{\delta})$ and $\mathcal{B}(\underline{\delta}, \bar{\delta})$ are defined in (11).
 5 Let us complete the proof by deriving the principal's payment to the agent and the
 6 agent's realized payoff net of execution costs. First, the principal payment is equal to

$$\begin{aligned} \mathcal{P} &= d(\mathbb{E}_{\underline{\delta}}[\mathcal{W}(\mathcal{R})] + \alpha \mathcal{C}_A) + (1 - d)\mathcal{W}_0 = d((1 - \underline{\delta})\mathcal{W}_0 + \underline{\delta}\mathcal{W}_1 + \alpha \mathcal{C}_A) + (1 - d)\mathcal{W}_0 \\ &= d((1 - \underline{\delta})(\widehat{\omega} + \mathcal{W}_0) + \underline{\delta}(\widehat{\omega} + \mathcal{W}_1) - (\widehat{\omega} + \mathcal{W}_0) + \alpha \mathcal{C}_A) + \mathcal{W}_0 \\ &= d((1 - \underline{\delta})\mathcal{A}(\underline{\delta}, \bar{\delta}) + \underline{\delta}\mathcal{B}(\underline{\delta}, \bar{\delta}) - 1)(\widehat{\omega} + \mathcal{W}_0) + d\alpha \mathcal{C}_A + \mathcal{W}_0 \\ &= d(\widehat{\mathcal{V}}(\underline{\delta}; \bar{\delta}) - 1)(\widehat{\omega} + \mathcal{W}_0) + d\alpha \mathcal{C}_A + \mathcal{W}_0 \\ &= \mathcal{W}_0 + d((\widehat{\mathcal{V}}(\underline{\delta}, \bar{\delta}) - 1)(\widehat{\omega} + \mathcal{W}_0) + \alpha \mathcal{C}_A), \end{aligned}$$

- 7 where the second-to-last equality uses the identity $(1 - \underline{\delta})\mathcal{A}(\underline{\delta}, \bar{\delta}) + \underline{\delta}\mathcal{B}(\underline{\delta}, \bar{\delta}) = \widehat{\mathcal{V}}(\underline{\delta}, \bar{\delta})$.
 8 On the hand, the agent's realized payoff net of execution costs is $d\mathbb{E}_{\underline{\delta}}[\widehat{\omega} + \mathcal{W}(\mathcal{R})] + (1 - d)(\widehat{\omega} + \mathcal{W}_0)$.
 9 Using a similar derivation as for $\mathcal{P}$, we get that this is equal to $(1 + d(\widehat{\mathcal{V}} - 1))(\mathcal{W}_0 + \widehat{\omega})$.

1 $\square$

PROOF OF PROPOSITION 2: To derive the moment generating function $\mathbb{E}_\delta[e^{s\tau}]$ of τ , let us consider a function $f(\delta)$ such that $f(\underline{\delta}) = f(\bar{\delta}) = 1$ and

$$\frac{1}{2} \sigma^2 \delta^2 (1 - \delta)^2 f''(\delta) + s f(\delta) = 0 \quad \text{for all } \delta \in [\underline{\delta}, \bar{\delta}].$$

For $s < \sigma^2/8$, the solution to this ODE is given by $f(\delta) = K_0 F_s(\delta) + K_1 F_s(1 - \delta)$ for two constants of integration K_0 and K_1 , where

$$F_s(\delta) = \frac{(1 - \delta)^{\eta(s)}}{\delta^{\eta(s)-1}} \quad \text{with } \eta(s) = \frac{1 + \sqrt{1 - 8s/\sigma^2}}{2}.$$

We find the values of K_0 and K_1 imposing the boundary conditions $f(\underline{\delta}) = f(\bar{\delta}) = 1$.

It follows that

$$f(\delta) = \frac{(F_s(1 - \bar{\delta}) - F_s(1 - \underline{\delta})) F_s(\delta) + (F_s(\underline{\delta}) - F_s(\bar{\delta})) F_s(1 - \delta)}{F_s(\underline{\delta}) F_s(1 - \bar{\delta}) - F_s(\bar{\delta}) F_s(1 - \underline{\delta})}.$$

From Dynkin's formula (see [Øksendal, 2013](#)) we get

$$\mathbb{E}_\delta[e^{s\tau} f(\delta_\tau)] = f(\delta) + \mathbb{E}_\delta \left[\int_0^\tau \left(\frac{1}{2} \sigma^2 \delta^2 (1 - \delta)^2 f''(\delta) + s f(\delta) \right) e^{st} dt \right] = f(\delta).$$

But since $f(\underline{\delta}) = f(\bar{\delta}) = 1$ we have that $\mathbb{E}_\delta[e^{s\tau} f(\delta_\tau)] = \mathbb{E}_\delta[e^{s\tau}]$. We conclude that

$$\mathbb{E}_\delta[e^{s\tau}] = \frac{(F_s(1 - \bar{\delta}) - F_s(1 - \underline{\delta})) F_s(\delta) + (F_s(\underline{\delta}) - F_s(\bar{\delta})) F_s(1 - \delta)}{F_s(\underline{\delta}) F_s(1 - \bar{\delta}) - F_s(\bar{\delta}) F_s(1 - \underline{\delta})}.$$

To compute the expected duration of due diligence, $\mathbb{E}_\delta[\tau]$, we can either evaluate the derivative of $\mathbb{E}_\delta[e^{s\tau}]$ with respect to s at $s = 0$. Alternatively, consider a function $g(\delta)$ such that

$$\frac{1}{2} \sigma^2 \delta^2 (1 - \delta)^2 g''(\delta) = 1 \quad \text{for all } \delta \in [\underline{\delta}, \bar{\delta}].$$

One particular solution is given by

$$g(\delta) = \frac{2(1-2\delta)}{\sigma^2} \ln \left(\frac{1-\delta}{\delta} \right).$$

Then, it follows that

$$\mathbb{E}_\delta[g(\delta_\tau)] = g(\delta) + \mathbb{E}_\delta \left[\int_0^\tau \frac{1}{2} \sigma^2 \delta^2 (1-\delta)^2 g''(\delta) dt \right] = g(\delta) + \mathbb{E}_\delta[\tau].$$

But since $\mathbb{E}_\delta[g(\delta_\tau)] = g(\underline{\delta}) \mathbb{P}_\delta(\delta_\tau = \underline{\delta}) + g(\bar{\delta}) \mathbb{P}_\delta(\delta_\tau = \bar{\delta})$, we have that

$$\mathbb{E}_\delta[\tau] = \mathbb{P}_\delta(\delta_\tau = \underline{\delta}) g(\underline{\delta}) + \mathbb{P}_\delta(\delta_\tau = \bar{\delta}) g(\bar{\delta}) - g(\delta).$$

Finally, we use a similar derivation to compute $\mathbb{P}_\delta(\delta_\tau = \underline{\delta})$ and $\mathbb{P}_\delta(\delta_\tau = \bar{\delta}) = 1 - \mathbb{P}_\delta(\delta_\tau = \underline{\delta})$.

Let us define the function $h(\delta)$ such that $h(\underline{\delta}) = 1$, $h(\bar{\delta}) = 0$ and

$$\frac{1}{2} \sigma^2 \delta^2 (1-\delta)^2 h''(\delta) = 0 \quad \text{for all } \delta \in [\underline{\delta}, \bar{\delta}].$$

It follows that $h(\delta) = (\bar{\delta} - \delta)/(\bar{\delta} - \underline{\delta})$. Then

$$\mathbb{P}_\delta(\delta_\tau = \underline{\delta}) = \mathbb{E}_\delta[\mathbb{1}(\delta_\tau = \underline{\delta})] = \mathbb{E}_\delta[h(\delta_\tau)] = h(\delta) + \mathbb{E}_\delta \left[\int_0^\tau \frac{1}{2} \sigma^2 \delta^2 (1-\delta)^2 h''(\delta) dt \right] = h(\delta) = \frac{\bar{\delta} - \delta}{\bar{\delta} - \underline{\delta}}.$$

And we conclude that

$$\mathbb{E}_\delta[\tau] = \left(\frac{\bar{\delta} - \delta}{\bar{\delta} - \underline{\delta}} \right) g(\underline{\delta}) + \left(\frac{\delta - \underline{\delta}}{\bar{\delta} - \underline{\delta}} \right) g(\bar{\delta}) - g(\delta). \quad \square$$

1

PROOF OF LEMMA 4: Consider the following constrained version of (13)

$$\mathbf{W}_1(\omega_0, \bar{\omega}) := \inf_{\mathcal{W}(\cdot)} \mathbb{E}_1[\mathcal{W}(\mathcal{R})] \quad \text{subject to} \quad \mathbb{E}_0[\mathcal{W}(\mathcal{R})] = \omega_0 \quad \text{and} \quad -\alpha \mathcal{C}_A \leq \mathcal{W}(\mathcal{R}) \leq \bar{\omega},$$

for some $\bar{\omega} \geq -\alpha \mathcal{C}_A$. We next solve for $\mathbf{W}_1(\omega_0, \bar{\omega})$ and then recover the value of $\mathbf{W}_1(\omega_0)$ in (13) by taking limit as $\bar{\omega} \rightarrow \infty$. To this end, let us “Lagrangianize” the first constraint to get the relaxation

$$\inf_{\mathcal{W}(\cdot)} \mathbb{E}_1[\mathcal{W}(\mathcal{R})] + \beta (\omega_0 - \mathbb{E}_0[\mathcal{W}(\mathcal{R})]) \quad \text{subject to} \quad -\alpha \mathcal{C}_A \leq \mathcal{W}(\mathcal{R}) \leq \bar{\omega}.$$

Recall that $L(x) = f_1(x)/f_0(x)$ denotes the likelihood ratio between the distributions of the project’ payoff under the two hypotheses and $\mathcal{L} = \lim_{x \rightarrow \infty} L(x)$. We can rewrite the relaxed objective as

$$\inf_{\mathcal{W}(\cdot)} \mathbb{E}_0[\mathcal{W}(\mathcal{R})(L(\mathcal{R}) - \beta)] + \beta \omega_0 \quad \text{subject to} \quad -\alpha \mathcal{C}_A \leq \mathcal{W}(\mathcal{R}) \leq \bar{\omega}.$$

It is not hard to see that –by optimizing the objective pointwise– the optimal solution is given by

$$\mathcal{W}^*(\mathcal{R}) = \bar{\omega} \mathbb{1}(L(\mathcal{R}) \leq \beta) - \alpha \mathcal{C}_A \mathbb{1}(L(\mathcal{R}) > \beta) = -\alpha \mathcal{C}_A + (\bar{\omega} + \alpha \mathcal{C}_A) \mathbb{1}(L(\mathcal{R}) \leq \beta).$$

Recall that under Assumption 1, condition (v), the likelihood ratio $L(x)$ is decreasing and we get that $\mathbb{1}(L(\mathcal{R}) \leq \beta) = \mathbb{1}(\mathcal{R} \geq r(\beta))$, where $r(\beta) = \inf\{r \geq 0 : L(\mathcal{R}) \leq \beta\}$. The value of β is obtained imposing the constraint $\mathbb{E}_0[\mathcal{W}^*(\mathcal{R})] = \omega_0$, which we can rewrite as

$$-\alpha \mathcal{C}_A + (\bar{\omega} + \alpha \mathcal{C}_A) \int_{r(\beta)}^{\infty} f_0(x) dx = \omega_0.$$

The resulting objective value is equal by

$$\mathbf{W}_1(\omega_0, \text{ow}) = -\alpha \mathcal{C}_A + (\bar{\omega} + \alpha \mathcal{C}_A) \int_{r(\beta)}^{\infty} f_1(x) dx = -\alpha \mathcal{C}_A + (\omega_0 + \alpha \mathcal{C}_A) \frac{\int_{r(\beta)}^{\infty} f_1(x) dx}{\int_{r(\beta)}^{\infty} f_0(x) dx}.$$

1 Finally, taking limit as $\bar{\omega} \rightarrow \infty$, is equivalent to take limit as $r(\beta) \rightarrow \infty$. We conclude

1 that

$$\begin{aligned}
\mathbf{W}_1(\omega_0) &= -\alpha \mathcal{C}_A + (\omega_0 + \alpha \mathcal{C}_A) \lim_{r \rightarrow \infty} \frac{\int_r^\infty f_1(x) \, dx}{\int_r^\infty f_0(x) \, dx} \\
&= -\alpha \mathcal{C}_A + (\omega_0 + \alpha \mathcal{C}_A) \lim_{r \rightarrow \infty} \frac{f_1(r)}{f_0(r)} \\
&= -\alpha \mathcal{C}_A + (\omega_0 + \alpha \mathcal{C}_A) \mathcal{L},
\end{aligned}$$

2 where the second equality follows by L'Hôpital's rule. $\square$

PROOF OF PROPOSITION 4: From condition (12) in Lemma 3, we have that the values of $\mathcal{W}_0$ and $\mathcal{W}_1$ can be expressed in terms of $\mathcal{W}_\emptyset$ and are equal to

$$\mathcal{W}_0 = \mathcal{A} \mathcal{W}_\emptyset + (\mathcal{A} - 1) \widehat{\omega}, \quad \mathcal{W}_1 = \mathcal{B} \mathcal{W}_\emptyset + (\mathcal{B} - 1) \widehat{\omega}.$$

3 Furthermore, from the same lemma, we have that the principal payment to the agent
4 is equal to $\mathcal{W}_\emptyset + d^* ((\widehat{\mathcal{V}} - 1) (\widehat{\omega} + \mathcal{W}_\emptyset) + \alpha \mathcal{C}_A)$. So, an optimal contract is found by
5 minimizing the value of $\mathcal{W}_\emptyset$ without violating the incentive compatibility constraints.
6 That is, by finding the minimum value of $\mathcal{W}_\emptyset$ so that the pair $(\mathcal{W}_0, \mathcal{W}_1)$ remains
7 feasible in the sense of $(\mathcal{W}_0, \mathcal{W}_1) \in \mathcal{I}$, as defined in (14).

This feasibility requirement reduces to the conditions

$$(\mathcal{B} - \mathcal{L} \mathcal{A}) (\widehat{\omega} + \mathcal{W}_\emptyset) \geq (1 - \mathcal{L}) (\widehat{\omega} - \alpha \mathcal{C}_A) \quad \text{and} \quad \mathcal{W}_\emptyset \geq 0.$$

8 We distinguish the following cases:

1. $\mathcal{B} \geq \mathcal{L} \mathcal{A}$. In this case, it is easy to see that

$$\mathcal{W}_\emptyset^* = \left[\left(\frac{1 - \mathcal{L}}{\mathcal{B} - \mathcal{L} \mathcal{A}} \right) (\widehat{\omega} - \alpha \mathcal{C}_A) - \widehat{\omega} \right]^+.$$

9 2. $\mathcal{B} = \mathcal{L} \mathcal{A}$. By the condition $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(\alpha)$ we must have that $\widehat{\omega} \leq \alpha \mathcal{C}_A$. Thus, it

1 follows trivially that in this case $\mathcal{W}_\emptyset^* = 0$.

3. $\mathcal{B} < \mathcal{L}\mathcal{A}$. In this case, the feasibility condition reduce to

$$0 \leq \mathcal{W}_\emptyset \leq \left(\frac{1 - \mathcal{L}}{\mathcal{B} - \mathcal{L}\mathcal{A}} \right) (\widehat{\omega} - \alpha \mathcal{C}_A) - \widehat{\omega}.$$

2 Again, from the condition $(\underline{\delta}, \bar{\delta}) \in \mathcal{X}(\alpha)$, it follows that there exists a feasible $\mathcal{W}_\emptyset$
 3 that implements $(\underline{\delta}, \bar{\delta})$, which implies that the right-hand side is non-negative.

4 Thus, setting $\mathcal{W}_\emptyset^* = 0$ is optimal.

We can combine the three cases above in a single condition

$$\mathcal{W}_\emptyset^* = \mathbb{1}(\mathcal{B} > \mathcal{L}\mathcal{A}) \left[\left(\frac{1 - \mathcal{L}}{\mathcal{B} - \mathcal{L}\mathcal{A}} \right) (\widehat{\omega} - \alpha \mathcal{C}_A) - \widehat{\omega} \right]^+. \quad \square$$

5

6 **PROOF OF LEMMA 5:** Suppose that it is optimal for the principal to offer a contract
 7 that induces the agent to conduct due diligence. Then, by the argument in Footnote 6,
 8 we have that $\min\{\mathcal{W}_0, \mathcal{W}_1\} < \mathcal{W}_\emptyset < \max\{\mathcal{W}_0, \mathcal{W}_1\}$. Let us show that we must have
 9 $\mathcal{W}_1 < \mathcal{W}_\emptyset < \mathcal{W}_0$ for an optimal contract.

10 Consider a contract $W = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}))$ and let us suppose (by contradiction) that
 11 $\mathcal{W}_0 < \mathcal{W}_\emptyset < \mathcal{W}_1$. By Proposition 1, we know that there are two cut-off beliefs $\{\underline{\delta}, \bar{\delta}\}$
 12 with $\underline{\delta} < \bar{\delta}$ such that the agent conducts due diligence as long as her belief $\delta_t \in (\underline{\delta}, \bar{\delta})$.
 13 In addition, if $\mathcal{W}_0 < \mathcal{W}_\emptyset < \mathcal{W}_1$ then in the boundary $\delta_t = \underline{\delta}$ the agent rejects the
 14 contract while in the boundary $\delta_t = \bar{\delta}$ the agent accepts the contract. Thus, we must
 15 have $\mathbb{E}_{\bar{\delta}}[\mathcal{W}(\mathcal{R})] \geq \mathcal{W}_\emptyset \geq \mathbb{E}_{\underline{\delta}}[\mathcal{W}(\mathcal{R})]$, that is,

$$(1 - \bar{\delta}) \mathcal{W}_0 + \bar{\delta} \mathcal{W}_1 \geq \mathcal{W}_\emptyset \geq (1 - \underline{\delta}) \mathcal{W}_0 + \underline{\delta} \mathcal{W}_1, \quad \text{it follows that} \quad (\underline{\delta} - \bar{\delta})(\mathcal{W}_0 - \mathcal{W}_1) \geq 0.$$

16 In addition, if a contract $(\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}))$ with these characteristics is optimal, the prin-

1 cipal must be weakly better off with this contract than not offering a contract at all,
 2 i.e., their expected payoff, if the agent accepts the contract, must be not negative,
 3 that is

$$\mathcal{U}(\delta, W) = \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \bar{\delta})] \mathbb{E}_{\bar{\delta}}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] - \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \underline{\delta})] \mathcal{W}_\emptyset \geq 0,$$

4 where $U(\mathcal{R}) = \mathcal{R} - \alpha \mathcal{C}_A - (1 - \alpha) \mathcal{C}_P$. We will use this expression for $\mathcal{U}(\delta, W)$ to show
 5 that the contract $W = (\mathcal{W}_\emptyset, \mathcal{W}(\mathcal{R}))$ cannot be optimal. To see this let us construct
 6 another contract $\widetilde{W} = (\widetilde{\mathcal{W}}_\emptyset, \widetilde{\mathcal{W}}(\mathcal{R}))$ that strictly dominates W . To this end, note that
 7 if the agent accepts the contract W when $\delta_\tau = \bar{\delta}$ then their expected payoff at this
 8 time τ satisfies $\tilde{w} := \mathbb{E}_{\bar{\delta}}[\mathcal{W}(\mathcal{R})] \geq \mathcal{W}_\emptyset$.

9 Let $\widetilde{W}$ be any contract such that $\widetilde{\mathcal{W}}_0 = \mathbb{E}_0[\widetilde{\mathcal{W}}(\mathcal{R})] = \tilde{w}$, $\widetilde{\mathcal{W}}_1 = \mathbb{E}_1[\widetilde{\mathcal{W}}(\mathcal{R})] = \tilde{w}$ and
 10 $\widetilde{\mathcal{W}}_\emptyset = \mathcal{W}_\emptyset$. Then, since $\tilde{w} \geq \mathcal{W}_\emptyset$, we have that the agent would accept the contract $\widetilde{W}$
 11 immediately if offered and recommend execution. Thus, under this contract $\widetilde{W}$, the
 12 principal's expected payoff is

$$\mathcal{U}(\delta, \widetilde{W}) = \mathbb{E}_\delta[U(\mathcal{R}) - \widetilde{\mathcal{W}}(\mathcal{R})] = \mathbb{E}_\delta[U(\mathcal{R})] - \tilde{w}.$$

13 But since $\delta < \bar{\delta}$ and $\mathbb{E}_0[U(\mathcal{R})] \geq \mathbb{E}_1[U(\mathcal{R})]$ by point (iv) in Assumption 1, it readily
 14 follows that this payoff is strictly higher than $\mathbb{E}_{\bar{\delta}}[U(\mathcal{R})] - \tilde{w}$. We conclude then that

$$\begin{aligned} \mathcal{U}(\delta, \widetilde{W}) &= \mathbb{E}_\delta[U(\mathcal{R})] - \tilde{w} > \mathbb{E}_{\bar{\delta}}[U(\mathcal{R})] - \tilde{w} = \mathbb{E}_{\bar{\delta}}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] \\ &\geq \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau \geq \bar{\delta})] \mathbb{E}_{\bar{\delta}}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] \\ &\geq \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \bar{\delta})] \mathbb{E}_{\bar{\delta}}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] - \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \underline{\delta})] \mathcal{W}_\emptyset \\ &= \mathcal{U}(\delta, W). \end{aligned}$$

15 This shows that contract $\widetilde{W}$ strictly dominates contract W . This contradiction implies
 16 that at optimality, we must have $\mathcal{W}_1 < \mathcal{W}_\emptyset < \mathcal{W}_0$ if the optimal contract induces the
 17 agent to conduct due diligence. This completes the proof of the first part.

1 Let us now show that at optimality $\mathbb{E}_1[U(\mathcal{R})] \leq \mathcal{W}_1$ and $\mathcal{W}_0 \leq \mathbb{E}_0[U(\mathcal{R})]$. For
 2 notational convenience, let us define $U_0 := \mathbb{E}_0[U(\mathcal{R})]$ and $U_1 := \mathbb{E}_1[U(\mathcal{R})]$

As above, let $\underline{\delta}$ and $\bar{\delta}$ denote the optimal thresholds that define the region of beliefs where the agent conducts due diligence. We have shown that $\mathcal{W}_1 < \mathcal{W}_\emptyset < \mathcal{W}_0$. Also, given the initial belief $\delta \in (\underline{\delta}, \bar{\delta})$, the principal's expected payoff is given by

$$\mathcal{U}(\delta, W) = \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \underline{\delta})] \mathbb{E}_{\bar{\delta}}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] - \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \bar{\delta})] \mathcal{W}_\emptyset$$

3 where $\tau = \inf\{t > 0: \delta_t \notin (\underline{\delta}, \bar{\delta})\}$. Let us assume, by contradiction, that at least one
 4 of the inequalities $U_1 \leq \mathcal{W}_1$ and $\mathcal{W}_0 \leq U_0$ is not satisfied.

5 Then, under Assumption 1 point (iv), one of the following three cases must hold:

(i) $\mathcal{W}_1 < U_1 < \mathcal{W}_\emptyset < \mathcal{W}_0 \leq U_0$. In this case $\mathbb{E}_{\delta_\tau}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] > 0$ (a.s.) and so

$$\mathcal{U}(\delta, W) < \mathbb{E}_\delta[\mathbb{E}_{\delta_\tau}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})]] = \mathbb{E}_\delta[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})],$$

6 where the equality follows from the fact that δ_t is a martingale. However,
 7 the assumption that the contract induces the agent to conduct due diligence
 8 implies that $\mathcal{W}_\emptyset \leq \mathbb{E}_\delta[\mathcal{W}(\mathcal{R})]$, which in turn yields $\mathcal{U}(\delta, W) < \mathbb{E}_\delta[U(\mathcal{R}) - \mathcal{W}_\emptyset]$.
 9 This inequality contradicts the optimality of the contract, since the right-hand
 10 side represents the expected payoff the principal would receive by offering any
 11 contract $\widetilde{W}$ with $\widetilde{\mathcal{W}}_0 = \widetilde{\mathcal{W}}_1 = \mathcal{W}_\emptyset$ that the agent accepts and, being indifferent,
 12 immediately recommends either execution or abandonment without performing
 13 any due diligence.

14 (ii) $U_1 \leq \mathcal{W}_1 < \mathcal{W}_\emptyset < U_0 < \mathcal{W}_0$. In this case, $\mathbb{E}_{\delta_\tau}[U(\mathcal{R}) - \mathcal{W}(\mathcal{R})] < 0$ (a.s.) and so
 15 $\mathcal{U}(\delta, W) < 0$, which contradicts the optimality of the contract.

1 (iii) $\mathcal{W}_1 < U_1 < \mathcal{W}_\emptyset < U_0 < \mathcal{W}_0$. The principal's expected payoff can be written as

$$\mathcal{U}(\delta, W) = \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \underline{\delta})] [(1 - \underline{\delta})(U_0 - \mathcal{W}_0) + \underline{\delta}(U_1 - \mathcal{W}_1)] - \mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \bar{\delta})] \mathcal{W}_\emptyset. \quad (20)$$

2 It is not hard to see that the term $\mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \underline{\delta})]$ is increasing in $\underline{\delta}$, while the
 3 term $\mathbb{E}_\delta[e^{-r\tau} \mathbb{1}(\delta_\tau = \bar{\delta})]$ is decreasing in $\underline{\delta}$. Using the representations of $\mathcal{W}_0$ and
 4 $\mathcal{W}_1$ in terms of $\underline{\delta}$ and $\bar{\delta}$, we next show that the factor $(1 - \underline{\delta})(U_0 - \mathcal{W}_0) + \underline{\delta}(U_1 - \mathcal{W}_1)$
 5 in the right-hand side is also increasing $\underline{\delta}$. These facts imply that $\mathcal{U}(\delta, W)$ is
 6 increasing $\underline{\delta}$, which contradicts the optimality of the contract W .

7 To this end, consider an arbitrary pair of thresholds $\underline{\delta}$ and $\bar{\delta}$ with $\underline{\delta} < \bar{\delta}$ and let
 8 $\mathcal{W}_0$ and $\mathcal{W}_1$ the corresponding payoff that induce cut-offs $\{\underline{\delta}, \bar{\delta}\}$. It follows that

$$\frac{\partial}{\partial \underline{\delta}} [(1 - \underline{\delta})(U_0 - \mathcal{W}_0) + \underline{\delta}(U_1 - \mathcal{W}_1)] = (U_1 - \mathcal{W}_1) - (U_0 - \mathcal{W}_0) - \left((1 - \underline{\delta}) \frac{\partial}{\partial \underline{\delta}} \mathcal{W}_0 + \underline{\delta} \frac{\partial}{\partial \underline{\delta}} \mathcal{W}_1 \right).$$

From equation (12) in Lemma 3 it also we have that $\mathcal{W}_0 + \widehat{\omega} = \mathcal{A}(\mathcal{W}_0 + \widehat{\omega})$ and
 $\mathcal{W}_1 + \widehat{\omega} = \mathcal{B}(\mathcal{W}_0 + \widehat{\omega})$. As a result, we get that

$$(1 - \underline{\delta}) \frac{\partial}{\partial \underline{\delta}} \mathcal{W}_0 + \underline{\delta} \frac{\partial}{\partial \underline{\delta}} \mathcal{W}_1 = \left((1 - \underline{\delta}) \frac{\partial}{\partial \underline{\delta}} \mathcal{A} + \underline{\delta} \frac{\partial}{\partial \underline{\delta}} \mathcal{B} \right) (\widehat{\omega} + \mathcal{W}_0) = 0,$$

where the last equality follows from (11). As a result, we get that

$$\frac{\partial}{\partial \underline{\delta}} [(1 - \underline{\delta})(U_0 - \mathcal{W}_0) + \underline{\delta}(U_1 - \mathcal{W}_1)] = (U_1 - \mathcal{W}_1) - (U_0 - \mathcal{W}_0) > 0$$

9 since we are considering the case $\mathcal{W}_1 < U_1 < \mathcal{W}_\emptyset < U_0 < \mathcal{W}_0$. We conclude that
 10 $\mathcal{U}(\delta, W)$ is increasing in $\underline{\delta}$ which contradicts its optimality. $\square$

References

- V.F. Araman and R.A. Caldentey. Diffusion approximations for a class of sequential experimentation problems. *Management Science*, 68(8):5958–5979, 2022. 7, 12, 22
- Kenneth J. Arrow, David Blackwell, and Meyer A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3-4):213–244, 1949. 7
- Felipe Balmaceda. Optimal task assignments. *Games and Economic Behavior*, 98:1 – 18, 2016. ISSN 0899-8256. doi: <http://dx.doi.org/10.1016/j.geb.2016.05.002>. URL <http://www.sciencedirect.com/science/article/pii/S0899825616300288>. 8, 9
- Suren Basov and Svetlana Danilkina. Multitasking, multidimensional screening, and moral hazard with risk neutral agents. *The Economic Record*, 86(s1): 80–86, September 2010. doi: 10.1111/j.1475-4932.2010.00662.x. URL <https://ideas.repec.org/a/bla/ecorec/v86y2010is1p80-86.html>. 8, 9
- P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2): 349–374, Mar 1999. 7
- P. Bond and A. Gomes. Multitask principal–agent problems: Optimal contracts, fragility, and effort misallocation. *Journal of Economic Theory*, 144(1):175–211, January 2009. 8
- Olivier Compte and Philippe Jehiel. Gathering information before signing a contract: A screening perspective. *International Journal of Industrial Organization*, 26(1):206–212, 2008. 6
- Jacques Crémer and Fahad A. Khalil. Gathering information before signing a contract. *American Economic Review*, 82(3):566–578, 1992. 6

1 M. Dewatripont, I. Jewitt, and J. Tirole. Multitask agency problems: Focus and
2 task clustering. *European Economic Review*, 44(Issues 4-6):869–877, 2000. 8

3 Felix Zhiyu Feng, Curtis R. Taylor, Mark M. Westerfield, and Feifan Zhang.
4 Setbacks, shutdowns, and overruns. *Econometrica*, 92(3):815–847, 2024. 7

5 George Georgiadis and Balazs Szentes. Optimal monitoring design. *Economet-*
6 *rica*, 88(5):2075–2107, 2020. 7

7 Dino Gerardi and Lucas Maestri. A principal–agent model of sequential testing.
8 *Theoretical Economics*, 7(3):425–463, 2012. 7

9 Brett Green and Curtis R. Taylor. Breakthroughs, deadlines, and self-reported
10 progress: Contracting for multistage projects. *American Economic Review*,
11 106(12), 2016. 7

12 Emeric Henry and Marco Ottaviani. Research and the approval process: The or-
13 ganization of persuasion. *American Economic Review*, 109(3):911–955, 2019.
14 7

15 Bengt Holmström and Paul Milgrom. Multitask principal-agent analyses: In-
16 centive schemes, asset ownership and job design. *Journal of Law, Economics,*
17 *& Organization*, 7:24–52, 1991. 8

18 Godfrey Keller, Sven Rady, and Martin Cripps. Strategic experimentation with
19 exponential bandits. *Econometrica*, 73(1):39–68, 2005. 7

20 Nicolas Klein. The importance of being honest. *Theoretical Economics*, 11(3):
21 773–811, 2016. 7

22 Ruitian Lang. Try before you buy: A theory of dynamic information acquisition.
23 *Journal of Economic Theory*, 183:1057–1093, 2019. 8

- 1 Christian Laux. Limited liability and incentive contracting with multiple
2 projects. *RAND Journal of Economics*, 32:514–526, 2001. 8
- 3 Tracy R. Lewis and David E. M. Sappington. Information management in
4 incentive problems. *Journal of Political Economy*, 105(5):796–821, 1997. 6
- 5 Erik Madsen. Designing Deadlines. *American Economic Review*, 112(3):
6 963–997, March 2022. 7
- 7 Andrew McClellan. Experimentation and approval mechanisms. *Econometrica*,
8 90(5):2215–2247, 2022. 7
- 9 Giuseppe Moscarini and Lones Smith. The optimal level of experimentation.
10 *Econometrica*, 69(6):1629–1644, 2001. 7
- 11 Bernt Øksendal. *Stochastic differential equations: an introduction with applica-*
12 *tions*. Springer Science & Business Media, 2013. 28, 53, 57
- 13 Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary prob-*
14 *lems*. Springer, 2006. 10
- 15 Kevin Roberts and Martin L Weitzman. Funding criteria for research, develop-
16 ment, and exploration projects. *Econometrica*, 49(5):1261–1288, 1981. 7
- 17 Dezső Szalay. The economics of clear advice and extreme options. *The Review*
18 *of Economic Studies*, 72(4):1173–1198, 2005. 7
- 19 A. Wald and J. Wolfowitz. Optimum character of the sequential probability
20 ratio test. *Ann. Math. Stat.*, 19(3):326–339, 1948. 7
- 21 Abraham Wald. Sequential analysis. *John Wiley & Sons*, 1947. 7
- 22 Eyal Winter. Incentive reversal. *American Economic Journal: Microeconomics*,
23 1(2):133–147, 2009. 8

- 1 Eyal Winter. Transparency and incentives among peers. *RAND Journal of*
2 *Economics*, 41(3):504–523, 2010. 8, 9
- 3 Joachim Winter. Sequential task performance, incentives and coordination.
4 *Journal of Economic Behavior & Organization*, 61(1):116–132, 2006. 8, 9
- 5 Yu Fu Wong. Dynamic monitoring design. Technical report, SSRN,
6 June 2025. Available at SSRN: <https://ssrn.com/abstract=4466562> or
7 <http://dx.doi.org/10.2139/ssrn.4466562>. 8
- 8 Weijie Zhong. Optimal dynamic information acquisition. *Econometrica*, 90(4):
9 1537–1582, 2022. 8