

When Are Two Heads Better than One? The Organizational Design of Innovation

Felix Z. Feng* Brett Green† Curtis R. Taylor‡ Mark Westerfield§

March 4, 2026

Abstract

We study the optimal organizational design in a dynamic model of discovery. The innovation process involves two complementary tasks: research and development, both of which are subject to a moral hazard problem. Moreover, agents observe private information about the quality of discoveries during the process. The principal must elicit this information to avoid costly mistakes. We study the question of how to allocate tasks to agents over the project’s life cycle. Assigning both tasks to a “solo innovator” requires dissuading double deviations (shirk and lie). As time passes without a discovery, such deviations become increasingly tempting. A team-based organizational structure eliminates the scope for double deviations, but is subject to a free-riding problem. We characterize when each structure is preferable and how these allocations shift over time. We provide conditions under which the organizational design changes over the project’s life cycle: the project begins with a solo innovator but transitions to a two-agent team after an endogenous deadline.

JEL Classification: D82, D86, D23, O31

Key Words: dynamic incentives, organizational design, contract theory, innovation

*Terry College of Business, University of Georgia, Email: ffeng@uga.edu

†Olin School of Business, Washington University in St. Louis, Email: b.green@wustl.edu

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

§Foster School of Business, University of Washington. Email: mwesterf@uw.edu

1 Introduction

Innovation projects routinely combine complementary activities—discovering promising ideas and turning them into working products or services—yet organizations differ markedly in how they allocate these tasks, and often reallocate them as learning unfolds. Some projects begin as a “one-person shop”—a founder (or lone engineer) who both explores the technological frontier and writes the first working version of the code or blueprint—and only later split into research, engineering, and product teams once a direction looks viable (Levy, 2011). Others impose a division of labor from day one, for example structurally separating exploratory work from execution-oriented development units with distinct goals and evaluation metrics (O’Reilly and Tushman, 2004). Still others rely on flexible forms of capacity: short-term consultants brought in to diagnose early setbacks (Brown, 2008), “tiger teams” activated after an initial prototype disappoints (Pavlak, 2004), or parallel groups that are temporarily spun up to pursue competing approaches before the organization consolidates around the most promising one (Sobek, Ward, and Liker, 1999). In this paper we ask how a principal should *dynamically* allocate research and development—whether to concentrate both tasks in a single agent or split them across multiple agents—when effort choices and intermediate information are privately held.

Formally, we consider a setting in which a project involves two tasks: *research*, which determines the quality of discoveries, and *development*, which generates those discoveries over time. Agents privately observe the quality of their discoveries and the principal must elicit this information to avoid costly mistakes. The central tension is that by consolidating tasks to a “solo innovator” the principal must prevent the agent from engaging in *double deviations*: shirking and then lying. As the agent’s continuation value declines, this double deviation becomes increasingly attractive. Separating the tasks across two agents eliminates the scope for double deviations but sacrifices economies of scope.

First, we characterize the optimal incentive contracts for both organizational structures and show that the cost of incentivizing a solo innovator is *history-dependent*—it increases as the agent’s continuation value falls—while the cost of a two-agent team is constant. This generates a time-varying tradeoff: early in the project, a solo innovator is cheaper; later, when continuation values are low, a team may be preferable.

Second, we characterize the optimal *dynamic* organizational design, which depends crucially on a key parameter K representing the difficulty of detecting flawed output. When K is small—meaning flawed output is easy to detect—the project is optimally run with a sequence of solo innovators: at any point, a single agent is responsible for both tasks, and when one is terminated, another takes his place. By contrast, when flaws are difficult to de-

tect, it is optimal to hire both agents simultaneously from the outset and assign each agent to a specialized task. For intermediate values of K , the optimal design features a staggered overlap: the project starts with a solo innovator, but before that agent is terminated, a second agent is brought in, transitioning the organization to a two-agent team. Thus, the optimal organizational structure shifts systematically with the difficulty of detecting flawed output.

Third, within the staggered overlap regime, we identify two distinct structures—the *Standby Developer* and the *External Consultant*—that differ in how tasks are allocated when the second agent joins. These structures arise endogenously from the solution to the principal’s problem and reflect the optimal way to preserve incentive capacity across agents.

Our work contributes to the theory of organizational design by showing how dynamic incentive capacity itself becomes a state variable that governs optimal task allocation over a project’s life cycle. Our mechanism is not driven by technological comparative advantage or learning-by-doing, but by how bundling versus separating tasks changes the set of profitable deviations under limited liability and private information. The model yields robust predictions about when organizations should expand from one agent to two, and how the marginal hire is used: environments with more opaque quality (harder ex post verification, weaker internal auditing, or greater scope for “passing off” lemons) should exhibit earlier specialization and more frequent transitions from generalists to teams, while environments with more transparent quality should exhibit sequential generalist structures. More broadly, the model provides a tractable framework for thinking about widely observed patterns in R&D and product development—from founder-led early stages to later-stage specialization, and from internal standby capacity to external consulting relationships in the form of mitigation experts—as optimal responses to evolving incentive constraints rather than purely technological considerations.

Related Literature. This paper contributes to the literature on team production and task allocation in agency settings. The most closely related paper is [Laux \(2001\)](#), which studies task bundling under limited liability when tasks are independent and their outcomes are separately observable. In that setting, assigning all tasks to a single agent is always optimal, because bundling relaxes the limited-liability constraint: outcomes from tasks on which the agent exerts effort can effectively “bond” the agent and deter shirking on other tasks. [Schmitz \(2005\)](#) studies a related two-stage model that favors separation through a different mechanism. As in our model, success in the first stage is not required for success in the second. When a single agent controls both stages, the most profitable deviation is therefore to shirk in the first stage to increase rent extraction in the second. Separating

control rights mitigates this distortion by limiting the second-stage agent’s ability to benefit from early failure. However, this mechanism relies on strong assumptions: the outcome of the first stage must be observable and contractible, and control rights must be allocated before first-stage outcomes are realized. If control rights could instead be made contingent on first-stage performance, it would be optimal to allow an agent to continue only following success—analogous to the sequential hiring structure we characterize in Section 5, and to the optimal organization in [Laux \(2001\)](#) when the contracting problem can be repeated after task failure.

In a similar vein, [Gromb and Martimort \(2007\)](#) study a setting in which a principal relies on experts to gather and report two private signals about a project’s value. They show that hiring two experts—each responsible for gathering and reporting one signal—can be cheaper than relying on a single agent to report both signals, because contracts can exploit cross-checking: the agent is paid only if the two reports agree. This ranking is reversed, however, when experts can collude, either horizontally among themselves or vertically with the principal. Related tradeoffs also appear in the literature on public–private partnerships ([Martimort and Pouyet, 2008](#); [Iossa and Martimort, 2012, 2015](#)), which studies analogous issues in procurement settings. A key distinction is that these models are static in their contracting structures and control rights. In contrast, we study a dynamic environment in which the optimal organization—one agent versus two—varies endogenously with agents’ performance histories, generating a richer set of organizational arrangements consistent with those observed in practice.

Our study is also related to the broader literature on team effort, including [Admati and Perry \(1991\)](#), [Keller, Rady, and Cripps \(2005\)](#), [Bonatti and Horner \(2011\)](#), [Cetemen, Hwang, and Kaya \(2020\)](#), [Yildirim \(2023\)](#), among others. This literature has traditionally focused on moral hazard—specifically, the underprovision of effort by team members—either across agents (e.g., free-riding) or over time (e.g., procrastination). A subset of studies emphasizes the benefits of teams. For example, [Bolton and Harris \(1999\)](#), [Yildirim \(2006\)](#), and [Georgiadis \(2015\)](#) analyze environments in which effort by one team member increases the productivity of others, often by accelerating the arrival of a critical event (such as project completion) at which all team members are rewarded. The latter two studies also examine how the project horizon interacts with team size, showing that encouragement effects dominate early in a project, while free-riding becomes more severe at later stages. Other work, such as [Marx and Matthews \(2000\)](#) and [Che and Yoo \(2001\)](#), highlights how teams may deter free-riding in dynamic settings when agents can collectively punish deviations by withholding future effort.

Our study differs from this literature in several important ways. First, all of the afore-

mentioned studies consider teams of fixed size.¹ They typically consider identical agents performing identical tasks. In contrast, we allow the principal to adjust the team size dynamically as the project unfolds and identify when the expansion of the team (from one to two agents) is optimal. Moreover, because we consider asymmetric tasks, our results generate additional implications for *how* the additional team member will be utilized, not only *when*. Finally, in most of the existing studies, there is typically no agency friction when the project is operated by a single agent. Instead, agency considerations such as free-riding or procrastination arise from team production itself. In our model, agency frictions exist even with a single agent in the form of both moral hazard and adverse selection. However, the nature of the main agency friction (whether it is under-provision of effort or misreporting of the outcome) determines whether team effort exacerbates or mitigates these frictions.

Finally, our model belongs to the dynamic contracting literature with both moral hazard and adverse selection, such as [Gershkov and Perry \(2012\)](#), [Halac, Kartik, and Liu \(2016\)](#), [Green and Taylor \(2016\)](#), [He, Wei, Yu, and Gao \(2017\)](#), [Liu and Lu \(2018\)](#), [Varas \(2018\)](#), [Che, Iossa, and Rey \(2021\)](#), etc. In all these studies, the principal only contracts with a single agent. Our model allows contracting with multiple agents, demonstrating how different organizational arrangements (solo innovator or an agent team) help reduce the cost of incentives for different constellations of agency frictions. Our model illustrates how moral hazard or adverse selection can endogenously become costly to incentivize relative to each other, both statically and dynamically, and how these shifting incentive costs generate optimal transitions across organizational forms.

Outline. Section 2 presents a static benchmark that provides intuition for why the principal may prefer two agents. Section 3 introduces the dynamic environment. Section 4 derives the incentive compatibility conditions under each organizational structure. Section 5 presents our main results on optimal organizational design. Section 6 discusses several model variations. Section 7 concludes. All proofs are in the Appendix.

2 Static Benchmark

2.1 Basic Environment

There are three dates $t \in \{0, 1, 2\}$ and three players: two agents and one principal. The principal contracts with one or more agents to complete a project that entails two tasks,

¹[Yildirim \(2006\)](#) and [Georgiadis \(2015\)](#) discuss optimal team sizes but team size must be chosen before project inception and cannot be changed over time. Thus, the costs and benefits of larger teams are derived through comparative statics.

which we label *research* and *development*. Both agents have identical skills and are capable of completing either task, or both. All players are risk neutral, and agents are protected by limited liability. We refer to the principal as “she” and to each agent as “he.”

At $t = 0$, the principal can invest in research at cost A . Research is carried out by an agent, who chooses an unobservable research effort $a \in \{s, w\}$. Action s corresponds to *shirking* and yields a private benefit $b < A$ to the agent. When the agent shirks, he produces a high-quality discovery, which we refer to as a *peach*, with probability $\lambda_0 < 1$; otherwise the discovery is low-quality, which we refer to as a *lemon*. The action w corresponds to *working*, which yields no private benefit but increases the probability of producing a peach to $\lambda_1 = \lambda_0 + \Delta_\lambda$, where $\Delta_\lambda > 0$.

At $t = 1$, the discovery quality is privately observed by the agent(s) and reported to the principal. Based on this report, the principal decides whether to develop the product, which requires an additional investment c . Development is carried out by an agent who chooses an unobservable effort $e \in \{s, w\}$. If the agent works, the project is completed at $t = 2$. If he shirks, he obtains a private benefit b and the project is not completed.

At $t = 2$, payoffs are realized. If the discovery quality is a peach and the project is completed, the principal receives payoff R . If the discovery is a lemon and the project is completed, the principal receives payoff R with probability $1 - q$ and $L < R$ with probability q . Thus q is the probability with which lemons can be identified ex post if they are produced. If the principal decides not to develop, or the agent shirks during the development phase, the principal’s payoff at $t = 2$ is zero.

To fix ideas, we impose the following assumptions.

Assumption 1 *Lemons should not be produced:*

$$(1 - q)R + qL - c < 0.$$

Assumption 2 *Research is socially optimal:*

$$\Delta_\lambda(R - c) \geq A.$$

2.2 Analysis

We now analyze the optimal contract for assigning both tasks to a single agent versus that for two agents, each assigned one task.

A. Solo Innovator

Suppose the principal contracts with a single agent to complete both tasks. If the agent reports a lemon after the research phase, he is paid γ and the project is terminated. If the agent reports a peach, the principal invests in the development phase and pays the agent a reward β if payoff R is realized, and zero otherwise. A contract is therefore a pair (β, γ) .

There are two binding constraints the principal must satisfy: (1) incentive compatibility in the research phase and (2) truthful reporting of lemons. We refer to the combination of shirking in the research phase and lying about lemons as the *double deviation*. Since the binding constraints are precisely those that make the double deviation unprofitable, this deviation effectively determines the structure of the optimal contract.

Inducing effort in the research phase requires:

$$\lambda_1 \beta + (1 - \lambda_1) \gamma \geq \lambda_0 \beta + (1 - \lambda_0) \gamma + b \iff \beta \geq \frac{b}{\Delta_\lambda} + \gamma \quad (1)$$

Because shirking on research leads to a higher probability of collecting γ , the cost of inducing research effort increases with γ .

To induce truthtelling, the reward must be better than (i) lying and then shirking, which yields b , or (ii) lying and then exerting effort in the development phase, which yields β with probability $1 - q$. Therefore, the truthtelling condition is

$$\gamma \geq \max \{b, (1 - q) \beta\} \quad (2)$$

There are two additional constraints to check: truthful reporting of peaches and inducing effort in the development phase. The former requires $\beta \geq \gamma$ while the latter requires $\beta \geq b$. Both are implied by (1). Setting (1) and (2) to equality and solving for β and γ , we get the following.

Lemma 1 *With a solo innovator, the optimal contract is:*

$$(\beta^*, \gamma^*) = \begin{cases} \left(\frac{b}{\Delta_\lambda} + b, b \right), & \text{if } q \geq \frac{1}{1 + \Delta_\lambda}, \\ \left(\frac{b}{q \Delta_\lambda}, \frac{(1 - q) b}{q \Delta_\lambda} \right), & \text{if } q < \frac{1}{1 + \Delta_\lambda} \end{cases}$$

When q is large, the relevant deviation in (2) is to lie and shirk, which yields only b . In contrast, when q is small, the relevant deviation is to lie and then work, which requires $\gamma > b$, and therefore necessitates higher powered incentives to satisfy (1).

B. Two-Agent Team

Now suppose that the principal uses a separate agent for each task. Let (β_R, γ_R) be the contract for the agent in charge of research and (β_D, γ_D) denote the contract for the agent assigned to development. Assume that both agents perfectly observe project quality, so the principal need only elicit a truthful report from one of them.

The researcher's IC and truth-telling constraints are

$$\begin{aligned}\beta_R &\geq \frac{b}{\Delta_\lambda} + \gamma_R \\ \gamma_R &\geq (1 - q) \beta_R\end{aligned}$$

The developer's IC and truth-telling constraints are

$$\begin{aligned}\beta_D &\geq b \\ \gamma_D &\geq \max\{b, (1 - q)\beta_D\}\end{aligned}$$

For the researcher, the cost of motivating effort increases with the reward for reporting a lemon. Intuitively, this is because the researcher controls project quality so the temptation to shirk and produce more lemons is higher when the reward for producing a lemon is higher. In contrast, the cost of motivating the developer to exert effort is independent of the reward for reporting a lemon: if the developer shirks, he reduces the likelihood of completing the project but does not influence whether lemons are produced. For this reason, it is cheaper to elicit project quality from the developer.

Lemma 2 *With a two-agent team, the optimal contract elicits truth-telling from the developer (i.e., only rewards lemons reported by the developer) and involves*

$$\beta_D^* = \gamma_D^* = b, \quad \beta_R^* = \frac{b}{\Delta_\lambda}, \quad \gamma_R^* = 0.$$

2.3 Cost Comparison

In both cases, the actions of each agent and the investment decision of the principal are the same. Hence, to determine which organizational design is preferable, it is sufficient to compare their expected cost of compensation. With the solo innovator, the expected cost is:

$$\lambda_1 \beta^* + (1 - \lambda_1) \gamma^* = \lambda_1 \frac{b}{\Delta_\lambda} + \gamma^*$$

With two agents, the expected compensation cost is

$$\lambda_1(\beta_D^* + \beta_R^*) + (1 - \lambda_1)\gamma_D^* = \lambda_1 \frac{b}{\Delta_\lambda} + b$$

Proposition 1 *The principal strictly prefers two agents to a solo innovator if and only if $\gamma^* > b$, which is equivalent to*

$$q < \frac{1}{1 + \Delta_\lambda}.$$

The smaller is q , the harder it is for the principal to detect and punish lemons ex post. This, in turn, makes it more expensive to prevent the double deviation in which the agent shirks on research and then lies about quality. Only a solo innovator can engage in such a deviation. In the two-agent structure, the researcher can shirk but is not responsible for reporting. The developer can lie and then shirk on the development task, but does not extract any additional rent from lying. Thus, by unbundling the tasks of quality production and quality reporting, the principal can reduce total compensation costs precisely when double deviations are hard to detect. This insight will play an important role in the dynamic model.

Two additional remarks are worthwhile. First, private information is essential for the mechanism we have just described. If quality is observed by the principal, a single agent is always cheaper because the principal can set $\gamma = 0$, while the cost of two agents is unchanged. Second, while the static model provides intuition for why two agents can be preferable, it falls short of addressing our main question of interest: how does the optimal organizational design vary over the life of the project?

3 Dynamic Environment

We now extend the framework to a dynamic environment in which discoveries arrive over time, their quality depends on complementary tasks, and the principal must decide how to allocate these tasks as the project evolves.

3.1 Setting

Time is continuous. There are two tasks, which we continue to refer to as *research* and *development*. The development task requires a flow cost c and generates discoveries according to a Poisson process with intensity 1. The research task also requires cost c and improves the quality of each discovery as it arises. Without research, a discovery is a peach with

probability $1 - \theta$ and a lemon with probability θ ; with research, all discoveries are peaches. Without loss of generality, we normalize the Poisson arrival rate to $\lambda = 1$ by rescaling time.

Each agent can perform either task or both tasks, provided the principal supplies the necessary resources. Agents are technologically identical, so differences in assignment reflect incentives rather than technological comparative advantage. An agent assigned only to development is the *developer*, an agent assigned only to research is the *researcher*, and an agent performing both tasks is a *solo innovator*.

At any time t , the agent(s) simultaneously choose their actions. The developer chooses $e_t \in A^D = \{s, w\}$, where s denotes shirking and w denotes working. The researcher chooses $a_t \in A^R = \{s_p, s_a, w\}$, where s_p denotes passive shirking, i.e., providing no research, and s_a denotes active shirking, providing limited input.² A solo innovator chooses the pair $(e_t, a_t) \in A^D \times A^R$.

Each action profile maps into an arrival rate of discoveries and a distribution over their quality. Shirking on development eliminates the possibility of making a discovery and yields private flow benefit b . Shirking on research affects the quality distribution of discoveries. If the researcher shirks passively, discoveries follow the baseline distribution without research. If the researcher shirks actively, lemons never arise, but the arrival rate of peaches is $(1 - \theta) < 1$. Arrival rates of each type of discovery given any action profile appear in Table 1. For convenience, we assume that either type of shirking yields private benefit $b < c$. Allowing the private benefits to differ across types of shirking would not affect any qualitative results.

Development e_t	Research a_t	Peaches	Lemons
w	w	1	0
w	s_p	$(1 - \theta)$	θ
w	s_a	$(1 - \theta)$	0
s	$\dots$	0	0

Table 1: Arrival rates under feasible action profiles.

As in the static benchmark, when a peach is built and accepted by the principal it yields the payoff $R > 0$. If a lemon is built and accepted by the principal, it develops a flaw with probability q and pays $-L < 0$. With probability $(1 - q)$, no flaw develops and the lemon also yields R . Agency frictions arise because agent's efforts are private and the developer privately observes whether a discovery is a peach or a lemon. The principal observes when a product is discovered, but does not observe its quality. The principal only ex post observes the payoff R or $-L$ if the discovery is produced.

²We discuss our motivation for allowing two forms of shirking in Section 6.

3.2 First Best

Before proceeding, it is useful to establish the first-best benchmark of the dynamic version of the model:

Lemma 3 *In the first best when the principal sees the agent's effort and the quality of any discovery, if $\theta > 1/2$, then the principal finances research and development until a peach is discovered. Her first best expected payoff is given by*

$$\pi = R - 2c. \quad (3)$$

Intuitively, research is worthwhile to the principal if it improves the likelihood of discovering a peach substantially. The lower bound for θ is $1/2$ in the first best but is higher when the principal is subject to asymmetric information, given the additional rents to the agent needed than just incentivizing development. Because the question of whether one or two agents should be employed is only relevant when both tasks are desired, throughout the rest of paper we will impose a stronger condition to ensure that financing and incentivizing research is always preferred by the principal despite the agency cost:

Assumption 3 *Research is sufficiently valuable to the principal compared to its cost:*

$$\theta \geq 1 - \frac{c}{R} > \frac{1}{2}$$

In the Appendix (Section B.4), we derive a weaker sufficient condition that ensures research is always optimal. In what follows, we maintain this assumption so that the principal always invests in and incentivizes both tasks at all times until the project is terminated.

3.3 Contracts and Continuation Utilities

We index agents by $i \in \mathcal{I} = \{1, 2\}$ and tasks by $j \in \{R, D\}$. We work on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The principal observes reported quality and realized payoffs; her information filtration is $\{\hat{\mathcal{F}}_t\}_{t \geq 0}$. Each agent i observes his own actions, any quality signals he receives, and his compensation; his filtration is $\{\mathcal{F}_t^i\}_{t \geq 0}$.

A contract specifies, for each agent i , an assignment process $\chi_{it} \in \{0, 1\}^2$ and a cumulative compensation process C_t^i , both adapted to $\hat{\mathcal{F}}_t$. The first component of χ_{it} indicates assignment to research, the second to development. Thus, $\chi_{it} = (1, 0)$ indicates that agent i is the researcher, $\chi_{it} = (0, 1)$ indicates that agent i is the developer, and $\chi_{it} = (1, 1)$ indicates that agent i is the solo innovator. Let $B(\chi_{it}, a_t, e_t)$ denote the flow of private benefits that accrue to agent i at time t under the assignment rule χ_t and action profile (a_t, e_t) .

We refer to the date at which agent i is hired as $t_i = \inf\{t : \chi_{it} \neq (0, 0)\}$. Agent i is terminated at $\tau^i = \inf\{t : \chi_{is} = (0, 0) \text{ for all } s \geq t\}$. Let $\tau = \max\{\tau^1, \tau^2\}$ denote the project termination date.

Given a contract, agent i 's continuation utility at time t is

$$W_t^i = \mathbb{E} \left[\int_t^{\tau^i} B(\chi_{is}, a_s, e_s) ds + \int_t^{\tau^i} dC_s^i \middle| \mathcal{F}_t^i \right],$$

and the principal's continuation value is

$$F_t = \mathbb{E} \left[Y_\tau - 2c(\tau - t) - \int_t^\tau (dC_s^1 + dC_s^2) \middle| \hat{\mathcal{F}}_t \right], \quad (4)$$

where $Y_\tau = R$ if a peach is produced, $Y_\tau \in \{R, -L\}$ if a lemon is produced, and $Y_\tau = 0$ if the project ends without a product. A contract is incentive compatible if each agent finds it optimal to exert effort on each task to which he is assigned and report quality truthfully. A contract is optimal if it maximizes the principal's date-0 payoff, F_0 , over the set of contracts that are incentive compatible and deliver each agent non-negative initial utility: $W_{t_i}^i \geq 0$ for $i \in \{1, 2\}$.

4 The Agent(s) Problem

In this section we analyze the agent's IC constraints. First, we narrow the contracting space:

Lemma 4 *The optimal contract never induces shirking. It is without loss to defer all payments to date τ and the principal only develops a discovery that the agent reports to be a peach.*

The first part of the result arises because $b < c$, so any shirking can be strictly improved by skipping that portion of the contract and giving the agent compensation equal to his shirking benefits. The second part of the result arises because both the principal and the agents share the same discount factor (zero) so deferring payments to when the contract ends is costless. Furthermore, because the agent is risk-neutral, any payment that the agent could receive when the contract ends without a product and be moved to the state in which a product is generated. Finally, because $(1 - q)R < qL$, the principal never accepts a lemon if the agent reports so.

Let N_t denote the counting process of successful discoveries, adapted to $\hat{\mathcal{F}}_t$. For any

incentive compatible contract, there exists a predictable process β_t^i , adapted to $\hat{\mathcal{F}}_t$, such that

$$dW_t^i = \beta_t^i(dN_t - dt), \quad (5)$$

The evolution of the agent's continuation utility given in (5) is based on the standard martingale representation theorem and the fact that the agent does not discount future payoffs. Intuitively, β_t^i is the jump in W_t^i when a peach is produced. In the absence of discoveries, W_t^i drifts downward at rate β_t^i until it hits zero, at which point the agent is terminated.

Proposition 2 *Given the assignment χ_{it} , effort is incentive compatible if and only if*

$$\beta_t^i \geq \begin{cases} \beta_D \equiv b, & \chi_{it} = (0, 1) \text{ (developer)}, \\ \beta_R \equiv b/\theta, & \chi_{it} = (1, 0) \text{ (researcher)}, \\ \beta_S(W_t^i) \equiv \max\{2b, b/(\theta q) - W_t^i\}, & \chi_{it} = (1, 1) \text{ (solo innovator)}, \end{cases}$$

and all inequalities bind in the optimal contract.

Below, we derive the IC conditions in Proposition 2 given the agent's assignment, and elaborate their implications on the cost of hiring one versus two agents.

4.1 Solo Innovator

To incentive the solo innovator, the principal must guard against a set of possible deviations. The agent can shirk on either or both tasks and the agent can misreport lemons as peaches or vice versa. We can split the incentive compatibility condition into two sets of constraints. The first set of constraints ensures that agent exerts on both tasks effort conditional on truthful reporting. The second set of constraints ensures that agent reports the quality of discoveries truthfully.

Let β_t^S denote the solo innovator's reward for reporting a peach that generates R and let γ_t^S denote the reward for reporting a lemon. The agent's expected flow benefit from exerting effort on both tasks is β_t^S . Consider first the strategy of shirking on both tasks (*double shirking*), which generates flow private benefit $2b$ and no possibility of a discovery. Preventing this deviation requires $\beta_t^S \geq 2b$. If this constraint is satisfied, we can immediately rule out two other deviations: (1) shirking only on development, which also yields no possibility of discoveries but a private benefit of only b , and (2) working on development and a-shirking on research, which yields a flow benefit of $b + (1 - \theta)\beta_t^S \leq \beta_t^S$ since $\theta \geq 1/2$. Therefore, the

effort IC can be reduced to:

$$\beta_t^S \geq \max \left\{ \underbrace{2b}_{\text{double shirking}}, \underbrace{(1-\theta)\beta_t^S + \theta\gamma_t^S + b}_{\text{p-shirking}} \right\}. \quad (\text{IC Effort})$$

Next, consider the solo innovator's incentive to report truthfully. Reporting peaches truthfully is incentive compatible provided that $\beta_t^S \geq \gamma_t^S$, which as we will see, is not a binding constraint. If the agent misreports a lemon as a peach, then the principal develops the discovery. With probability $1 - q$, no flaw materializes, the principal realizes R , and the agent is rewarded β_t^S . With probability q , a flaw materializes, which reveals that the agent lied and the principal punishes the agent to the maximum extent possible. However, due to limited liability, the maximum punishment is simply to terminate the agent. Hence, truthfully reporting lemons requires that.

$$\gamma_t^S + W_t^S \geq (1 - q)(\beta_t^S + W_t^S) + q \cdot 0 \quad (6)$$

Rearranging terms, we get

$$\gamma_t^S \geq (1 - q)\beta_t^S - qW_t^S. \quad (\text{IC Report})$$

Observe that making either constraint (IC Effort) or (IC Report) slack will make the other constraint harder to satisfy. Therefore, both constraints will hold with equality at the optimum. Substituting the value for γ_t^S from (IC Report) into (IC Effort), we arrive at the expression in Proposition 2:

$$\beta_t^S = \beta_S(W_t^S) = \max \left\{ 2b, \frac{b}{\theta q} - W_t^S \right\}. \quad (7)$$

When W_t^S is high, the agent has substantial "skin in the game," so the threat of termination is powerful and the binding constraint is double shirking. When W_t^i is low, the agent is more tempted to misreport a lemon as a peach, so truthful reporting requires $\gamma_t^S > 0$, which in turn raises the attractiveness of p -shirking.³ Whether this temptation exceeds the benefit of double shirking depends on model parameters which can be summarized as follows.

³Producing lemons requires the agent's effort in development, so the agent cannot reap unlimited utility by continuously fabricating lemons.

Corollary 1 *If $\theta q < 1/2$, then*

$$\beta_S(W) = \begin{cases} 2b, & W \geq \underline{W}, \\ K - W & W < \underline{W}, \end{cases} \quad (8)$$

where

$$K \equiv \frac{b}{\theta q} \quad (9)$$

$$\underline{W} \equiv K(1 - 2\theta q) = \frac{b(1 - 2\theta q)}{\theta q} \quad (10)$$

If $\theta q \geq 1/2$, then $\beta_S(W) = 2b$ for all $W \geq 0$.

Recall that q corresponds to the likelihood that a lemon turns out to be flawed. The lower q is, the more difficult it is for the principal to distinguish the lemon from the peach *ex post* and punish the agent for lying. Similar to when the agent has a lower continuation value, this makes the double deviation more tempting, which requires stronger incentives to dissuade. $\underline{W}$ captures the continuation value at which the agent is indifferent between double shirking and the double deviation. The principal's cost of providing incentives is constant when $W_t^S > \underline{W}$ and increases as W_t^S declines below $\underline{W}$, reaching the highest level as $W_t^S \rightarrow 0$. Notice that $\underline{W}$ decreases with q . When q is sufficiently high, $\underline{W} \leq 0$. In this case, the principal only needs to compensate the agent for the private benefits from double shirking so β_t^S is always constant ($2b$).

4.2 Two Agents

We now derive the IC conditions when the two tasks are assigned to separate agents. The timing of the stage game is as follows. Each agent simultaneously and privately chooses whether to put in effort. If a discovery is made, the developer observes its quality and reports it to the principal. If the developer reports a peach, the principal produces the discovery and realizes R or $-L$. If the developer reports a lemon, the discovery is discarded. We require the equilibrium of the stage game to be subgame perfect.

Suppose the researcher exerts effort. For the developer, the expected flow benefit from working is β_t^D . If he shirks, he receives private flow benefit b and nothing is produced. Thus, incentivizing the developer to work requires:

$$\beta_t^D \geq b. \quad (11)$$

To ensure the developer truthfully reports peaches and lemons requires

$$\beta_t^D \geq \gamma_t^D, \quad (12)$$

$$\gamma_t^D \geq (1 - q)\beta_t^D - qW_t^D. \quad (13)$$

In equilibrium, the researcher will exert effort and thus lemons will not be produced on path. Therefore, γ_t^D does not enter the developer's IC constraint, which means the principal does not have to give information rents to the developer for reporting truthfully. Critically, this means the principal has no need to elicit information from the researcher and can use the developer's report to provide incentives to the researcher.

The IC condition for the researcher is

$$\beta_t^R \geq \max\{\underbrace{(1 - \theta)\beta_t^R + b}_{a\text{-shirking}}, \underbrace{(1 - \theta)\beta_t^R + b + \theta\gamma_t^R}_{p\text{-shirking}}\} \quad (14)$$

The principal can set $\gamma_t^R = -W_t^R$ to deter p -shirking. Importantly, as long as $\gamma_t^R \leq 0$, which is always feasible, a -shirking dominates p -shirking. That is, given that the developer will report lemons truthfully, the researcher cannot gain from generating lemons via p -shirking. Thus, if he decides to shirk on research, he is always better off shirking actively so as not to generate lemons. Hence, the a -shirking constraint will bind under the optimal contract and $\beta_t^R = \beta_R = \frac{b}{\theta}$.

4.3 Cost Comparison.

Comparing the IC conditions for two agents and that for a solo-innovator reveals the following result regarding the cost of either organization structure:

Corollary 2 *The principal's flow cost of incentivizing two separate agents is lower than the flow cost of a solo-innovator if q is sufficiently low and the continuation utility of the solo-innovator is sufficiently low, i.e.,*

$$\beta_D + \beta_R < \beta_S(W_t^S) \text{ if } q < \frac{1}{1 + \theta} \text{ and } W_t^S < K - \frac{b(1 + \theta)}{\theta} \quad (15)$$

Notice that the condition $q < \frac{1}{1 + \theta}$ is identical to the condition for two agents to be cheaper in Proposition 1, where θ is equivalent to Δ_λ as they both capture the incremental likelihood of peach if the research effort is spent. Thus, the intuition for why two agents can be cheaper than a solo innovator from the static model still applies.

Yet, the dynamic model also illustrates an additional insight: the cost of a solo innovator

is *time-varying* depending on how much continuation utility, or time, the solo innovator has left. The intuition of the static model is a special case for exactly *one attempt* at developing a high quality product, which corresponds to when $W_t^S \rightarrow 0$. The cost of incentives for the solo innovator is lower as the “number of attempts” grows, while the total cost of the team is constant.

In summary, the organizational structure determines which deviations are binding and, therefore, the total cost of incentive provision. By consolidating the two tasks to a solo innovator, the principal can leverage economies of scope and offer a lower total reward to motivate effort (i.e., $2b < b + \frac{b}{\theta}$). However, a solo innovator must be disciplined against double deviations, which is more expensive when the agent’s continuation value is low and when lemons are difficult to detect.

5 Optimal Contract

Having characterized the incentive compatibility conditions under each organizational regime and explored the trade-offs between a solo innovator and a two-agent team, we are now ready to address our primary question of interest: what is the optimal dynamic arrangement that maximizes the principal’s expected payoff at the outset of the contract. Because the two agents are identical, we will label the agent that is hired first as agent 1, and the second as agent 2. If both agents are hired at the same time, agent 1 refers to the agent initially assigned to development. Before proceeding, we impose a parametric assumption that ensures the principal would consider using a two-agent team.

Assumption 4

$$\pi > \frac{(1 + \theta)b}{\theta} = \beta_D + \beta_R \quad (16)$$

This ensures that the expected output exceeds the cost of providing incentives to the team for one “instant”. Without this assumption, the principal would never use a two-agent team structure. We do not make the analogous assumption for the solo innovator structure because it depends on our parameter of interest, q , which also plays an important role in the determination of the optimal organizational design.

5.1 When Both Are Hired

We start off by solving the task assignment problem *given* that both agents have been hired with exogenous continuation values W_1 for agent 1 and W_2 for agent 2. Let $V(W_1, W_2)$

denote the principal's value function given the continuation value of each agent. Thus, for each $\vec{W} = (W_1, W_2) \in \mathbb{R}_+^2$ the principal has four options:

1. Use Agent 1 as a solo innovator.
2. Use Agent 2 as a solo innovator.
3. Use Agent 1 as a researcher and Agent 2 as a developer.
4. Use Agent 1 as a developer and Agent 2 as a researcher.

The HJB for the principal's problem is

$$V(W_1, W_2) = \pi - (W_1 + W_2) - \min \left\{ \begin{array}{l} \beta_S(W_1)(1 + V_1), \\ \beta_S(W_2)(1 + V_2), \\ \beta_D(1 + V_1) + \beta_R(1 + V_2), \\ \beta_R(1 + V_1) + \beta_D(1 + V_2) \end{array} \right\} \quad (17)$$

where $V_i = \frac{\partial V}{\partial W_i}$. The boundary conditions for the principal's problem are

$$\begin{aligned} V(W_1, 0) &= F_1(W_1) \\ V(0, W_2) &= F_1(W_2) \end{aligned}$$

where F_1 is the principal's value function with a single agent assigned to both tasks (see Appendix A for derivation).

Because the problem is symmetric, it suffices to focus on either the region where $W_1 \geq W_2$ or where $W_2 \geq W_1$. We focus on the latter. As will be shown later, this is also where the equilibrium path lies (recall agent 1 is the first agent assigned to one or more tasks). Define $\mathcal{W}_2 = \{\vec{W} \in \mathbb{R}_+^2 : W_2 \geq W_1\}$.

Because V is symmetric and must be weakly concave in both of its arguments, it will be never be optimal to assign the agent with less continuation value to both tasks or to the more expensive task. Intuitively, it is better to assign the agent with the higher continuation value to both tasks or the task that requires stronger incentives. Hence, for $W_2 \geq W_1$, we only need to consider two of the four possible allocations: assign agent 2 to both tasks or assign agent 2 to research and assign agent 1 to development. The HJB for the principal's problem can be simplified to

$$V(\vec{W}) = \pi - (W_1 + W_2) - \min \{ \beta_S(W_2)V_2(\vec{W}), \beta_R V_2(\vec{W}) + \beta_D V_1(\vec{W}) \}.$$

Define W_S such that $\beta_S(W_S) = \beta_D + \beta_R$, which is the continuation value such that assigning one agent to both tasks requires the same reward as using one agent for each task. Direct calculation yields

$$W_S = \frac{b(1 - q(1 + \theta))}{\theta q} \quad (18)$$

which is non-negative if $q \leq 1/(1 + \theta)$. The *myopic solution*, which minimizes the total reward to both agents in each state, is to assign agent 2 to both tasks when $W_2 > W_S$ and each agent to one task when $W_2 < W_S$. The optimal task allocation differs from the myopic solution because it accounts for the influence the current assignment has on the cost of future assignments.

To understand where this difference arises, consider the point labeled A in Figure 1. Starting from point A , the myopic solution assigns agent 1 to development and agent 2 to research. In this case, both agents' continuation values decrease if the project does not succeed. Because agent 2 is assigned to the more expensive task, her continuation value decreases faster than agent 1's, i.e., $dW_2/dW_1 = \beta_R/\beta_D = \frac{1}{\theta}$. Hence, $\vec{W}_t$ travels downward along the line with slope $1/\theta$ until reaching the y-axis at which point agent 1 is terminated and agent 2 is assigned to both tasks.

In contrast, the optimal solution starting from point A assigns agent 2 to both tasks, so that $\vec{W}_t$ travels vertically downward until reaching $\mathcal{K}$, at which point each agent is assigned to a single task. In both the optimal and myopic solutions, the amount of time that each agent is assigned to one task is the same. The difference is that under the optimal solution, agent 2 is assigned to both tasks when her continuation value is relatively high, and hence when incentive costs are relatively low. As a result, the principal can increase the time agent 2 works on both tasks by front-loading this assignment, rather than postponing it until agent 1 is terminated, as in the myopic solution.

We now characterize the solution to the task assignment program. In order to do so, let $\mathcal{K} = \{\vec{W} \in [0, W_S]^2 : \theta W_1 \leq W_2 \leq W_1/\theta\}$ denote the “kite” shaped region that extends northeast from the origin. Let $\mathcal{A}_i = \{\vec{W} : W_i > \min\{W_j/\theta, W_S\}, W_j < W_S\}$ and let $\mathcal{E} = \{\vec{W} : W_1, W_2 > W_S\}$.

Proposition 3 *The solution to the task assignment problem is as follows:*

- If $\vec{W} \in \mathcal{K}$, it is optimal to assign one task to each agent.
- If $\vec{W} \in \mathcal{A}_i$, then it is optimal to assign both tasks to agent i .
- If $\vec{W} \in \mathcal{E}$, then allocating both tasks to either agent is optimal.




Figure 1: Illustration of the solution to the task assignment problem

5.2 When to Hire the Second Agent

Having characterized the solution to the task assignment problem and the principal's value function $V(W_1, W_2)$ once both agents are hired, we now determine *when* it is optimal for the principal to bring in the second agent. At this stage of the project, assume that only agent 1 is active and the principal can decide when to hire agent 2. The decision problem thus takes the form of a stopping problem: continue with a single solo-innovator or hire the second agent.

5.2.1 The Principal's Problem

If the principal hires the second agent when the state reaches $W_t = w$, she enters the two-agent regime characterized by Proposition 3. Conditional on hiring at that point, she chooses the initial promise to the new agent to maximize her continuation payoff:

$$G(w) = \max_{y \geq 0} V(w, y).$$

The function $G(w)$ gives the principal's value *from hiring agent 2 immediately* when agent 1's continuation value is w . Because the principal can randomize over continuation utilities for

agent 1 at the hiring date, the relevant object for the stopping problem is the *concave envelope* of G ,

$$\bar{G}(w) = \sup\{g \mid (w, g) \in \text{co}(G)\},$$

which represents the best payoff attainable from possibly mixing over agent's continuation values at the time of hire.⁴ We refer to $\bar{G}$ as the *obstacle*.

Starting with a single agent, the principal chooses the initial promise W_0 and a stopping time τ at which to hire the second agent to maximize her payoff:

$$\Pi = \sup_{(\tau, W_0)} \int_0^\tau e^{-t} (\pi - W_t - \beta_S(W_t)) dt + e^{-\tau} \bar{G}(W_\tau), \quad dW_t = -\beta_S(W_t) dt. \quad (19)$$

The first term is the expected payoff from continuing with one agent until τ , while the second term is the continuation value if agent 2 is hired at that point. Let (τ^*, W_0^*) denote the solution to the problem and let $\hat{W}_1 = W_{\tau^*}$.

Differentiating (19) with respect to τ gives the marginal value of waiting an additional instant before hiring the second agent:

$$\frac{d\Pi}{d\tau} = e^{-\tau} \Phi(W_\tau), \quad \Phi(W) \equiv \pi - W - \beta_S(W)(1 + \bar{G}'(W)) - \bar{G}(W).$$

The function $\Phi(W)$ therefore captures the *benefit of delay*: if $\Phi(W) > 0$, it is optimal to wait longer before hiring agent 2, while if $\Phi(W) < 0$, the principal should hire agent 2 immediately. An interior solution $\hat{W}_1$ must satisfy the first-order condition

$$\Phi(\hat{W}_1) = 0, \quad (20)$$

which corresponds to a smooth-pasting condition between the single-agent and two-agent value functions. Observe that for all $W \in [0, W_S]$,

$$\Phi'(W) = -(K - W)\bar{G}''(W) \geq 0.$$

Thus, Φ satisfies the (weak) single-crossing property.

Connection to the QVI formulation. The stopping problem in (19) is equivalent to a one-obstacle quasi-variational inequality (QVI) for the principal's value function $F(W)$

⁴For any point $(w, g) \in \text{co}(G)$, there exists a distribution μ over continuation values for agent 1 such that $E_\mu W_1 = w$ and $E_\mu G(W_1) = g$. Therefore, conditional on hiring agent 2 when agent 1's continuation value is W_1 , the principal can achieve a payoff of $\bar{G}(W_1)$.

when she has the option to hire a second agent. Let the solo innovator $\mathcal{A}$ act on F as

$$\mathcal{A}F(W) = (\pi - W - \beta_S(W)) - \beta_S(W)F'(W) - F(W)$$

The principal's value function then satisfies

$$\max \{ \mathcal{A}F(W), \bar{G}(W) - F(W) \} = 0, \quad F(0) = \bar{G}(0). \quad (21)$$

In the continuation region ($W > \hat{W}_1$) the first term binds, while in the switching region ($W \leq \hat{W}_1$) we have $F(W) = \bar{G}(W)$. At an interior threshold $\hat{W}_1$, the free-boundary conditions

$$F(\hat{W}_1) = \bar{G}(\hat{W}_1), \quad \mathcal{A}F(\hat{W}_1) = 0$$

collapse to the first-order condition $\Phi(\hat{W}_1) = 0$ defined above. Hence, the QVI yields the solution to the stopping problem for any initial W_0 : $\Phi(W)$ is simply the HJB operator evaluated at the obstacle $\bar{G}$. Once we have solved (21) for F , we can then turn to the choice of the optimal W_0^* .

We proceed with two useful lemmas and then state our main result.

Lemma 5 *Let $W_2^*(w) \in \arg \max_{y \geq 0} V(w, y)$. If $(w, W_2^*(w)) \in \mathcal{A}_1 \cup \mathcal{E}$ then $\Phi(w) \geq 0$. In words: if, conditional on hiring agent 2, it is optimal to assign agent 1 to both tasks, then it is optimal to wait to hire agent 2.*

The lemma simply says there is no reason to “hire” agent 2 unless it is strictly optimal for the principal to put agent 2 to work conditional on being hired. This result is intuitive. Nevertheless, it is useful in restricting the domain over which it is necessary to look for an interior solution. Let $\mathcal{W}_1 = \{w \in \mathcal{R}_+ : (w, W_2^*(w)) \in \mathcal{K} \cup \mathcal{A}_2\}$ denote the relevant domain.

Next, we decompose $\mathcal{W}_1$ into two sub-regions: the stopping region ($\mathcal{S}$) and the continuation region ($\mathcal{C}$). We include K as an argument for each sub-region in order to highlight the sub-regions dependence on our key parameter of interest.

$$\mathcal{C}(K) \equiv \{W \in \mathcal{W}_1 : \Phi(W; K) > 0\}$$

$$\mathcal{S}(K) \equiv \{W \in \mathcal{W}_1 : \Phi(W; K) \leq 0\}$$

In words, $\mathcal{C}(K)$ is the region where the marginal benefit of delay is positive and therefore it is optimal for the principal to delay hiring agent 2. Whereas $\mathcal{S}(K)$ is the stopping region where immediately hiring agent 2 is optimal. Since Φ is weakly increasing in W , both $\mathcal{C}$ and $\mathcal{S}$ are intervals, which can be characterized by $\hat{W}_1(K)$, which is the largest value in $\mathcal{W}_1$ such

that $\Phi(\hat{W}_1) \leq 0$. We let $\hat{W}_1(K) = 0$ if Φ is strictly positive on $\mathcal{W}_1$. Thus, $\mathcal{S}(K) = [0, \hat{W}_1(K)]$ and $\mathcal{C}(K) = (\hat{W}_1(K), \max(\mathcal{W}_1)]$.

Lemma 6 *There exists positive and finite $K_{\text{seq}} < K_{\text{sim}}$ such that*

1. *For all $K \leq K_{\text{seq}}$, $\hat{W}_1(K) = 0$.*
2. *For all $K \in (K_{\text{seq}}, K_{\text{sim}})$, $\hat{W}_1(K) \in (0, \max(\mathcal{W}_1))$.*
3. *For all $K \geq K_{\text{sim}}$, $\hat{W}_1(K) = \max(\mathcal{W}_1)$.*

Hiring the second agent involves making a fixed cost investment (i.e., $W_2^*(W_1)$) that the principal would like to avoid. However, as W_1 decreases, it becomes more expensive to dissuade double deviations. For intermediate K , the principal resolves the tradeoff by waiting until W_1 reaches the critical level $\hat{W}_1(K)$ before adding the second agent to the team. When K is small, the cost of dissuading double deviations is low, and the optimal policy is to wait until $W_1 = 0$ (and agent 1 is terminated) before hiring agent 2. When K is large, double deviations are very costly to dissuade, making early hiring of the second agent optimal despite the fixed cost.

5.3 Organizational Regimes

Given any initial utility for Agent 1, Lemma 6 characterizes the optimal policy for when to hire the second agent $\tau = \inf\{t : W_t \leq \hat{W}_1(K)\}$. All that remains is to determine the optimal initial utility for the first agent. There are two possibilities. Either $W_0^* > \hat{W}_1$ in which case agent 1 works as a solo-innovator before agent 2 is hired, or $W_0^* = \hat{W}_1$, in which case the two agents are hired simultaneously. In this latter case, the initial starting point is not uniquely determined because the principal's value function is flat in the interior of the kite region. However, if there is an arbitrarily small cost of reallocating tasks to agents, then the solution is unique (up to the choice between two symmetric optimal paths). In what follows, we apply this selection criterion.

If $W_0^* > \hat{W}_1$, then the principal hires agent 1 before agent 2. Effectively, the principal “attaches” the single agent value function to $\bar{G}$. Hence, the principal's value function is linear on $[\hat{W}_1, \underline{W}]$ with slope A that is determined by the boundary condition

$$F(\hat{W}_1) = A\hat{W}_1 + \pi - (1 + A)K = \bar{G}(\hat{W}_1)$$

Rearranging terms, we see that

$$A(\hat{W}_1) = \frac{\pi - K - \bar{G}(\hat{W}_1)}{K - \hat{W}_1} \quad (22)$$

For intuition, we can note that $A(\hat{W}_1)$ corresponds to $d\Pi/dW_0$ evaluated at $W_0 = \hat{W}_1$. Thus, if $A(\hat{W}_1)$ is positive (negative), then the solution must involve $W_0^* > \hat{W}_1$ ($W_0^* = \hat{W}_1$).

Lemma 7 *Agent 1 is hired strictly before Agent 2 if and only if $K < K_{\text{sim}}$.*

From Lemmas 6 and 7, we can conclude that the optimal organizational design can take one of three possible forms as summarized in the following proposition.

Proposition 4 (Organizational Regimes) *The optimal organizational regime is as follows:*

1. **Sequential Generalist:** For $K \leq K_{\text{seq}}$: the project starts with agent 1 as the solo innovator. The principal hires agent 2 when agent 1 is terminated. Each agent is assigned to both tasks during their employment.
2. **Simultaneous Specialist:** For $K \geq K_{\text{sim}}$, the principal hires both agents at the same time. Each agent is assigned to one task during their employment. Both agents are terminated at the same time.
3. **Staggered Overlap:** For $K \in (K_{\text{seq}}, K_{\text{sim}})$, the project starts with agent 1 as the solo innovator. Agent 2 is hired before agent 1 is terminated.

In the sequential-generalist regime, it is optimal for a single agent to complete both tasks; effectively, whichever agent is employed acts as the solo innovator. By contrast, in the simultaneous-specialist regime, two heads are preferred to one, and each agent is assigned to a distinct task for the entire duration of the project. Importantly, these “specialists” do not arise from inherent differences in ability, but rather from the principal’s need to control agency rents. The research task is more costly to incentivize—due to the free-rider problem—so the agent assigned to research must be granted higher rents per unit of time than the developer. While the principal could, in principle, equalize the agents’ initial utilities, assign agent 1 to research and agent 2 to development, and then rotate their roles midway through the project, doing so would require incurring (albeit small) reallocation costs that she prefers to avoid. The optimal policy is therefore to grant one agent a higher initial utility and have him specialize in research for the duration of the project.

In the staggered overlap structure, agent 1 is assigned to both tasks for a period of time before agent 2 is hired. If agent 1 does not succeed by the time her continuation value reaches $\hat{W}_1$, then agent 2 is hired. How tasks are allocated to agents once agent 2 is hired can subsequently take one of two forms.

1. *Standby Developer*: This occurs when $(\hat{W}_1, W_2^*(\hat{W}_1)) \in \mathcal{A}_2$. In this case, agent 2 takes over both tasks upon being hired, while agent 1 is put on "standby." If agent 2 does not succeed for a while then eventually, the state reaches the kite region. At that point, agent 1 takes over development, while agent 2 continues research.
2. *External Consultant*: This occurs when $(\hat{W}_1, W_2^*(\hat{W}_1)) \in \mathcal{K}$. In this case, agent 2 is brought in as an "external consultant" in charge of research for the remaining duration of the project, while agent 1 continues to develop.

Proposition 5 *For all $K \in (K_{\text{seq}}, \tilde{K})$, the Standby Developer structure is optimal. For all $K \in (\tilde{K}, K_{\text{sim}})$, the External Consultant structure is optimal.*

To provide intuition, we leverage continuity of the solution with respect to K . At the boundaries:

- When $K \leq K_{\text{seq}}$: $\hat{W}_1(K) = 0$, so agent 2 is hired only after agent 1 is terminated (sequential hiring).
- When $K \geq K_{\text{sim}}$: $\hat{W}_1(K) = W_0^*$, so both agents are hired simultaneously.

For intermediate values $K \in (K_{\text{seq}}, K_{\text{sim}})$, the hiring threshold $\hat{W}_1(K)$ varies continuously between these extremes.

For K just above K_{seq} , the hiring threshold $\hat{W}_1(K)$ is close to zero—agent 1 has little continuation value remaining when agent 2 is hired. At this point, agent 1 is expensive to incentivize, while the freshly hired agent 2 has high continuation value and is cheap to incentivize. The optimal response is to hire agent 2 in $\mathcal{A}_2$ and let her operate as a solo innovator for a while. Agent 1's remaining continuation value is *preserved* for when it is most valuable: if agent 2 does not succeed and the state eventually enters the kite region $\mathcal{K}$. At that point, agent 1 rejoins to handle development while agent 2 continues research.

For K closer to K_{sim} , the hiring threshold $\hat{W}_1(K)$ is larger—agent 1 still has substantial continuation value when agent 2 is hired. Both agents now have enough skin in the game that splitting tasks is immediately optimal. Agent 2 enters as an "external consultant" for research, while agent 1 continues development. The agents work in parallel from the outset. The threshold $\tilde{K}$ is the unique value of K at which the optimal hiring point $(\hat{W}_1(K), W_2^*)$ lies on the boundary between $\mathcal{A}_2$ and $\mathcal{K}$.

6 Discussion of Assumptions and Extensions

In this section, we discuss several of our model assumptions and the implications of relaxing them.

6.1 Lemons on Path

When both agents exert effort, lemons are not produced on path. This assumption is convenient and simplifies the analysis because the developer's bump in utility from reporting a lemon does not materialize. Suppose instead that lemons are produced on path with probability $\epsilon > 0$, either when both agents exert effort or when the developer exerts effort and the researcher actively shirks. The arrival rate of lemons and peaches in this extension of the model is:

Development e_t	Research a_t	Peaches	Lemons
w	w	$(1 - \epsilon)$	ϵ
w	s_p	$(1 - \theta)$	θ
w	s_a	$(1 - \theta - \epsilon)$	ϵ
s	$\dots$	0	0

Table 2: Arrival rates under feasible action profiles with lemons on path.

In this case, the developer's continuation value increases following the report of a lemon, which makes characterizing the full solution to the principal's problem more involved. But the underlying economics remain unchanged.

Proposition 6 *Suppose lemons arise on path as in Table 2 and fix $\epsilon \in (0, \theta)$. Then the following contracts implement effort on both tasks and truthful reporting, and are optimal within their respective organizational regimes.*

1. *Solo Innovator:*

$$\beta_\epsilon^S(W^S) = \max\{K_\epsilon - W^S, \bar{\beta}_\epsilon\}, \quad \gamma_\epsilon^S(W^S) = \max\{(1 - q)K_\epsilon - W^S, \bar{\gamma}_\epsilon\},$$

$$\text{where } K_\epsilon := \frac{b}{q(\theta - \epsilon)}, \bar{\beta}_\epsilon := 2b + \epsilon \frac{b}{\theta - \epsilon}, \bar{\gamma}_\epsilon := 2b - \frac{(1 - \epsilon)b}{\theta - \epsilon}.$$

2. *Two-Agent Team:*

$$\beta_\epsilon^D = \gamma_\epsilon^D = b, \quad \beta_\epsilon^R = \frac{b}{\theta - \epsilon}, \quad \gamma_\epsilon^R = 0.$$

This structure closely mirrors Proposition 2 (the benchmark with $\epsilon = 0$). A new feature when lemons occur on path is that the principal may optimally reward the solo innovator for lemons even when the agent's continuation value is high. In particular,

$$\bar{\gamma}_\epsilon > 0 \iff \theta > \frac{1 + \epsilon}{2}.$$

The intuition comes from the double-shirking constraint,

$$(1 - \epsilon)\beta + \epsilon\gamma \geq 2b,$$

which requires the principal to provide sufficient incentives in expectation across peach and lemon realizations. While the principal would like to keep both instruments small, β is paid as a transfer upon a produced peach whereas γ can be delivered by increasing promised continuation utility following a reported lemon. Shifting compensation weight from β toward γ reduces the principal's expected transfer burden, making γ a relatively cheap way to satisfy double-shirking when lemons arrive with positive probability. This substitution is limited by the other IC constraints—most importantly, passive shirking and truthful reporting—which limit the extent to which the principal can rely on rewarding lemons to provide incentives.

Given that agents are now compensated for both peaches and lemons, comparing the total reward for peaches alone (as we did in Corollary 2) is no longer sufficient to characterize the total incentive cost of each organizational structure. Nevertheless, we can still provide an analogous and equally sharp characterization of when each regime is more cost-effective.

Corollary 3 *Define*

$$\bar{W}_\epsilon := (1 - q)K_\epsilon - b = \frac{(1 - q)b}{q(\theta - \epsilon)} - b,$$

which is decreasing in q and strictly positive if and only if $q < \frac{1}{1 + \theta - \epsilon}$. The principal's flow cost of incentivizing a solo innovator is lower than that of the two-agent team, componentwise, if and only if $W^S > \bar{W}_\epsilon$, i.e.,

$$\beta_\epsilon^S(W^S) < \beta_\epsilon^D + \beta_\epsilon^R \quad \text{and} \quad \gamma_\epsilon^S(W^S) < \gamma_\epsilon^D + \gamma_\epsilon^R \iff W^S > \bar{W}_\epsilon.$$

Both inequalities reverse if $W^S < \bar{W}_\epsilon$.

6.2 Forms of Shirking on Research

We allow two forms of shirking for the researcher in the dynamic model. To see why, consider an alternative version of the model in which the researcher can only shirk passively. In this

case, the researcher cannot obtain private benefits without producing lemons. The principal can then use the developer's report of a lemon to induce the researcher to exert effort. In particular, by setting $\gamma_t^R = -W_t^R$, the IC condition reduces to

$$\beta_t^R = \max \left\{ 0, \frac{b}{\theta} - W_t^R \right\}.$$

Thus, incentivizing the researcher is “free” when her continuation value is sufficiently high.⁵ In this case, the total cost of incentivizing a two-agent team also depends on agents' continuation values, making the comparison between the two organizational structures more nuanced. However, several critical properties persist. Namely, the cost of incentivizing the two-agent team is independent of q . Meanwhile, the cost of incentivizing the solo innovator remains unchanged (recall that passive shirking was the binding constraint for the solo innovator). This implies that the *relative cost* of the solo innovator compared to a team still increases as lemons become harder to detect. As a result, when q is sufficiently small, the principal prefers team-based innovation. Whereas for q sufficiently large, a solo innovator is more attractive.

6.3 Development Only

Assumption 3 ensures that research is optimal. When it fails, the principal may find it optimal to save on the investment and incentive costs of the researcher and hire a single agent responsible for development and truthfully reporting.

Lemma 8 *Suppose the principal does not invest in research. Then the optimal contract that induces effort on development and truthful reporting is $\beta = \gamma = b$.*

The optimal contract with a solo developer is remarkably simple. Because the solo developer only controls the arrival of discoveries, there is no need to differentiate the incentives between lemons and peaches. In fact, under this regime, holding total compensation fixed, the principal prefers rewarding the agent with γ (promised utility) rather than β (transfers) for the same reason discussed in Section 6.1. Of course, truthfully reporting peaches requires that $\beta \geq \gamma$, which binds at the optimum.

As we show in the appendix, when θ is sufficiently high, a solo developer is dominated by a solo innovator for all continuation values and thus not part of the optimal arrangement.

⁵This conclusion hinges on allowing the principal to punish one agent for a non-verifiable report made by another agent. If we impose the restriction that $\gamma_t^R \geq 0$, then this “simplified” version of the model collapses to our main specification.

The converse is also true when θ is sufficiently small. A complete analysis of the optimal contract when the solo developer is neither dominant nor dominated is left for future work.

7 Conclusion

This paper studies how a principal should organize innovation when progress requires two complementary tasks—research and development—and when agents both choose unobservable effort and privately observe quality-relevant information during the process. Concentrating tasks in a solo innovator economizes on incentive payments when the agent has substantial continuation value, but it also creates a dynamic vulnerability to double deviations in which the agent both reduces effort and misreports quality. Separating tasks across two agents eliminates these double deviations but introduces free riding, making research effort more expensive in a way that is constant over time. The resulting comparison yields a simple organizing principle: because the solo innovator’s incentive cost rises as continuation value declines (especially when flawed output is hard to detect), the optimal organizational form can shift systematically over the project’s life cycle.

Our main finding is that optimal organizational design can be *history dependent* even when the production technology is stationary. When flaws are easy to detect, the principal optimally relies on sequential generalists—a sequence of solo innovators responsible for both tasks. When flaws are difficult to detect, she instead hires a two-agent team immediately and sustains specialization throughout. For intermediate environments, the project begins with a solo innovator and then transitions to a two-agent structure via staggered overlap; depending on parameters, the second hire appears either as a temporary takeover with the incumbent on standby or as an external-consultant-style research specialist who joins an ongoing development effort. These regimes provide a unified explanation for common organizational patterns observed in innovation-intensive sectors.

Several avenues for future work are natural. Expanding the team beyond two agents or allowing endogenous monitoring and cross-checking technologies would let the principal trade off organizational design against investments that improve verifiability. Related, incorporating heterogeneity in agent abilities, learning about project difficulty over time, or explicit switching/reallocation costs would all further inform empirical predictions about who is hired, when expansions occur, and how task assignments evolve within the team. We view these directions as promising ways to build on the paper’s core insights and engage additional institutional features of innovation focused organizations.

Appendix

A Value Functions

This appendix derives the principal's value functions for a single agent and a two-agent team hired simultaneously. We begin with a useful observation: an ODE of the form

$$x_1 F'(W) + F(W) + W = x_0 \quad (\text{A.1})$$

has the general solution

$$F(W) = C \exp\left(-\frac{W}{x_1}\right) - W + (x_0 - x_1) \quad (\text{A.2})$$

where the constant C is determined by a boundary condition.

A.1 One Agent

Let $F_1(W)$ denote the principal's value function with one agent assigned to both tasks. Define $K \equiv \frac{b}{\theta q}$.

Region $W \leq \underline{W}$. In this region, the binding IC constraint is the double deviation, so $\beta_S(W) = K - W$. The continuation utility evolves as $dW_t = (K - W_t)(dN_t - dt)$. The HJB equation is:

$$0 = -2c - (K - W)F_1'(W) + [R - K - F_1(W)] \quad (\text{A.3})$$

Rearranging: $-(K - W)F_1'(W) - F_1(W) = \pi - K$, which has the linear solution

$$F_1(W) = AW + \pi - (1 + A)K$$

The boundary condition $F_1(0) = 0$ yields $A = \frac{\pi}{K} - 1 = \left(\frac{\theta q}{b}\right) \pi - 1$.

Region $W > \underline{W}$. Here $\beta_S = 2b$ is constant, so the HJB becomes:

$$\begin{aligned} 0 &= -2c - \beta_S F_1'(W) + [R - W - \beta_S - F_1(W)] \\ &= -2(c + b) - 2b F_1'(W) + [R - W - F_1(W)] \end{aligned} \quad (\text{A.4})$$

Rearranging into form (A.1): $2b F_1'(W) + F_1(W) + W = \pi$. Using $x_1 = 2b$ and $x_0 = \pi$ in (A.2):

$$F_1(W) = C_1 \exp(-\rho_1 W) - W + \pi, \quad \rho_1 = \frac{1}{2b} \quad (\text{A.5})$$

The boundary condition $F_1(\underline{W}) = A\underline{W} + \pi - (1 + A)K$ (continuity at $\underline{W}$) determines:

$$C_1 = \exp(\rho_1 \underline{W})(1 + A)(\underline{W} - K) \quad (\text{A.6})$$

Optimal starting utility. The optimal initial utility satisfies $F_1'(W_1^*) = 0$, yielding:

$$W_1^* = \frac{\ln(-\rho_1 C_1)}{\rho_1}, \quad F_1(W_1^*) = \pi - \frac{1}{\rho_1} - W_1^* \quad (\text{A.7})$$

Remark 1 *If the contract uses randomization instead of rewarding γ for reporting a lemon (keeping $W \geq \bar{W}$ with $\bar{W} = \frac{2b(1-q)}{q}$), then $F_1'(0) = (\frac{q}{2b})\pi - 1$. Rewarding for lemon reports dominates randomization if $\theta > 1/2$, which is Assumption 3.*

A.2 Two Agents Hired Simultaneously

With two agents hired simultaneously (one for each task), define $W = W_1 + W_2$ as total promised utility. The evolution is $dW_t = (\beta_D + \beta_R)(dN_t - dt)$ where $\beta_D + \beta_R = b(1 + 1/\theta)$. The HJB is:

$$0 = -2c - b \left(1 + \frac{1}{\theta}\right) F_2'(W) + \left[R - W - b \left(1 + \frac{1}{\theta}\right) - F_2(W)\right] \quad (\text{A.8})$$

Rearranging into form (A.1) with $x_1 = \frac{b(1+\theta)}{\theta}$ and $x_0 = \pi$. With boundary condition $F_2(0) = 0$, substituting into (A.2) gives $C = -(x_0 - x_1) = -(\pi - 1/\rho_2)$:

$$F_2(W) = \pi(1 - \exp(-\rho_2 W)) - W, \quad \rho_2 = \frac{\theta}{b(1 + \theta)} < \rho_1 \quad (\text{A.9})$$

Under Assumption 4, $F_2'(0) = \pi\rho_2 - 1 > 0$, so the optimal starting utility is interior:

$$W_{sim}^* = \frac{1}{\rho_2} \log(\pi\rho_2), \quad F_2(W_{sim}^*) = \pi - \frac{1}{\rho_2} - W_{sim}^*$$

B Proofs

B.1 Proof of Proposition 3

The proof of Proposition 3 works as follows. Through a series of lemmas, which partition the state space into regions, we construct the principal's value function under the conjectured optimal policy and then verify that it satisfies the HJB. In doing so, we illustrate the dynamics of how the continuation value of each agent evolves.

To simplify the mathematical expressions, we will solve for the optimal social welfare function $S(W_1, W_2) = V(W_1, W_2) + W_1 + W_2$, which for all $\vec{W} \in \mathcal{W}_2$ solves

$$S(\vec{W}) = \pi - \min \{ \beta_S(W_2)S_2(\vec{W}), \beta_R S_2(\vec{W}) + \beta_D S_1(\vec{W}) \}$$

subject to boundary condition $S(0, W_2) = F_1(W_2) + W_2, S(W_1, 0) = F_1(W_1) + W_1$, where $S_j \equiv \frac{\partial S(W_1, W_2)}{\partial W_j}$.

Lemma A.1 *For $\vec{W}_t \in \mathcal{K}$, it is optimal to assign one agent to each task. The social value function is given by*

$$S(\vec{W}_t) = \left[1 - \exp \left(-\tau(\vec{W}_t) \right) \right] \pi \quad (\text{A.10})$$

where $\tau(\vec{W}) \equiv \frac{W_1 + W_2}{\beta_R + \beta_D}$. Any path such that one agent is assigned to each task and $\vec{W}_{t+s} \in \mathcal{K}$ for all $s \leq \tau(\vec{W})$ achieves this payoff.

Proof of Lemma A.1. The expression in (A.10) is derived from the conjecture that $\beta_R S_2 + \beta_D S_1 = \pi - S$ and the boundary condition $S(0, 0) = 0$. Note that $\tau(\vec{W})$, corresponds to the length of time before $\vec{W}_t$ reaches the origin. To verify that assigning one task to each agent is optimal, it suffices to check that for any $\vec{W} \in \mathcal{K}$ such that $W_2 \geq W_1$,

$$\beta_S(W_2)S_2 \geq \beta_R S_2 + \beta_D S_1 = \pi - S$$

Note that $S_2 = S_1 = \frac{e^{-\tau}\pi}{\beta_R + \beta_D}$. Therefore,

$$\beta_S(W_2)S_2 = \frac{\beta_S(W_2)}{\beta_R + \beta_D} e^{-\tau}\pi > \pi - S \iff \beta_S(W_2) > \beta_R + \beta_D$$

which holds for $W_2 < W_S$, which is necessary for $\vec{W} \in \mathcal{K}$. Hence, the exact assignment is payoff irrelevant as long as the path does not exit $\mathcal{K}$. ■

Lemma A.2 *For any $\vec{W} \in \mathcal{A}_2$, it is optimal to assign both tasks to Agent 2 and*

1. *If $W_2 \leq \underline{W}$, the social value function is given by*

$$S(\vec{W}) = \pi - (K - W_2)c(W_1) \quad (\text{A.11})$$

where

$$c(W_1) = \begin{cases} \left(\frac{1}{K - W_1/\theta} \right) \exp \left(-\frac{W_1}{\beta_D} \right) \pi & \text{if } W_1 \leq \theta W_S \end{cases} \quad (\text{A.12})$$

$$\begin{cases} \left(\frac{1}{\beta_D + \beta_R} \right) \exp \left(-\frac{W_1 + W_S}{\beta_D + \beta_R} \right) \pi & \text{if } W_1 \in (\theta W_S, W_S] \end{cases} \quad (\text{A.13})$$

2. *If $W_2 > \underline{W}$, the social value function is given by*

$$S(\vec{W}) = \pi - k(W_1)e^{-\rho_1 W_2} \quad (\text{A.14})$$

where $k(W_1) = \frac{c(W_1)}{\rho_1} e^{\rho_1 \underline{W}}$.

Proof. Conjecture that the solution involves assigning Agent 2 to both tasks in this region.

1. If $W_2 \leq \underline{W}$, $\beta_S(W_2) = K - W_2$ and the value function takes the form

$$S(\vec{W}) = \pi - \beta_S(W_2)S_2 \implies S = \pi - (K - W_2)c(W_1) \quad (\text{A.15})$$

where $c(W_1)$ is a function of only W_1 that is determined from the boundary condition. Assigning Agent 2 to both tasks means moving vertically downward until hitting the boundary of $\mathcal{K}$. If $W_1 \leq \theta W_S$, this occurs when $W_2 = W_1/\theta$ at which point $S(W_1, W_1/\theta) = (1 - e^{-\frac{W_1}{\beta_D}})\pi$ (from (A.10)). Thus, the boundary condition is that

$$S\left(W_1, \frac{W_1}{\theta}\right) = \pi - \left(K - \frac{W_1}{\theta}\right)c(W_1) = \left(1 - e^{-\frac{W_1}{\beta_D}}\right)\pi \quad (\text{A.16})$$

Solving for $c(W_1)$, we get the expression in (A.12). We can verify that the conjectured policy is optimal by noting that for all $W_2 \geq \frac{W_1}{\theta}$

$$\beta_R S_2 + \beta_D S_1 - \beta_S(W_2)S_2 = \frac{b \exp(-W_1 \beta_D)(\theta W_2 - W_1)\pi}{(W_1 - K\theta)^2} \geq 0. \quad (\text{A.17})$$

If $W_1 \in (\theta W_S, W_S]$ then the boundary condition applies at the top border of the kite region where $W_2 = W_S$, and

$$S(W_1, W_S) = \pi - (K - W_S)c(W_1) = \left[1 - \exp\left(-\frac{W_1 + W_S}{\beta_D + \beta_R}\right)\right]\pi. \quad (\text{A.18})$$

Solving for $c(W_1)$, we get the expression in (A.13). Again, we can verify the conjecture is correct by noting that

$$\beta_R S_2 + \beta_D S_1 - \beta_S(W_2)S_2 = \frac{e^{-\frac{(W_1 + W_S)}{\beta_R + \beta_D}} \pi \beta_R}{(\beta_D + \beta_R)^2} (\beta_D + \beta_R - \beta_S(W_2)) \geq 0, \quad (\text{A.19})$$

where the inequality holds because $W_2 \geq W_S$.

2. If $W_2 > \underline{W}$, $\beta_S(W_2) = 2b = \frac{1}{\rho_1}$. Therefore,

$$S = \pi - \beta_S(W_2)S_2 = \pi - \frac{S_2}{\rho_1} \implies S = \pi - k(W_1)e^{-\rho_1 W_2} \quad (\text{A.20})$$

Under the conjectured solution, the path moves vertically downward to the boundary $W_2 = \underline{W}$. Therefore, using (A.11), the boundary condition is $S(W_1, \underline{W}) = \pi - (K - \underline{W})c(W_1)$. Solving for $k(W_1)$ yields $k(W_1) = \frac{c(W_1)}{\rho_1} e^{\rho_1 \underline{W}}$. To verify the candidate satisfies the HJB, we need to show that

$$\beta_S(W_2)S_2 = k(W_1)e^{-\rho_1 W_2} \leq e^{-\rho_1 W_2} \left[\frac{1}{2\theta} k(W_1) - \beta_D k'(W_1) \right] = \beta_D S_1 + \beta_R S_2.$$

Canceling $e^{-\rho_1 W_2}$, rearranging, and using the definition of $k(W_1)$, the above inequality

is equivalent to

$$\beta_D \frac{c'(W_1)}{c(W_1)} \leq \left(\frac{1}{2\theta} - 1 \right)$$

For $W_1 < \theta W_S$, $c(W_1) = \frac{\pi}{K - \frac{W_1}{\theta}} \exp\left(-\frac{W_1}{\beta_D}\right)$ and thus

$$\frac{c'(W_1)}{c(W_1)} = \frac{1}{\theta} \left(K - \frac{W_1}{\theta} \right)^{-1} - \frac{1}{\beta_D}.$$

Thus, it suffices to show that $\frac{\beta_D}{\theta} (K - W_1/\theta)^{-1} \leq \frac{1}{2\theta}$, which holds iff $K - W_1/\theta \geq 2\beta_D$. Since $W_1 < \theta W_S$, $K - W_1/\theta > K - W_S = \beta_D + \beta_R > 2\beta_D$.

For $W_1 \in (\theta W_S, W_S]$, $c(W_1) = \frac{\pi}{\beta_D + \beta_R} \exp\left(-\frac{W_1 - W_S}{\beta_D + \beta_R}\right)$, and thus

$$\beta_D \frac{c'(W_1)}{c(W_1)} = -\frac{\beta_D}{\beta_D + \beta_R} = -\frac{\theta}{1 + \theta}.$$

Noting that $-\frac{\theta}{1+\theta} < \frac{1}{2\theta} - 1$ for all $\theta < 1$ completes the verification step.

■

For $\vec{W} \in \mathcal{E}$, we break the proof up into three sub-regions of the space because the value function takes a different form in each sub-region.

Lemma A.3 *If both $W_1, W_2 \in [W_S, W]$, it is optimal to assign both tasks to one agent but the exact assignment (both to Agent 1 or 2) is payoff irrelevant. The social value function is*

$$S(\vec{W}) = \pi - (K - W_2)(K - W_1)\rho_2 c(W_S) \quad (\text{A.21})$$

where $c(W_S)$ is defined in (A.13).

Proof. Suppose both tasks are assigned to agent 1. The social value function is then

$$S(W_1, W_2) = \pi - \beta_S(W_1)S_1 \implies S = \pi - (K - W_1)\hat{c}(W_2) \quad (\text{A.22})$$

for some function $\hat{c}$. The boundary condition is $W_1 = W_S$, where

$$S(W_S, W_2) = \pi - (K - W_S)\hat{c}(W_2) = \pi - (K - W_2)c(W_S) \quad (\text{A.23})$$

where $c(W_S)$ is given by (A.13) with $W_1 = W_S$. Thus,

$$\hat{c}(W_2) = \left(\frac{K - W_2}{K - W_S} \right) c(W_S) = (K - W_2)\rho_2 c(W_S) \quad (\text{A.24})$$

which then yields (A.21). Notice that $\beta_S(W_1)S_1 = \beta_S(W_2)S_2$. Therefore, by the symmetry of S , the principal achieves the same payoff by assigning both tasks to agent 2 until reaching the boundary where $W_2 = W_S$. To verify the principal cannot do better by assigning

one task to each agent, it suffices to consider the case where $W_1 \leq W_2$ and show that $\beta_S(W_2)S_2 \leq \beta_R S_2 + \beta_D S_1$, which is equivalent to

$$(K - W_2)(K - W_1) \leq \beta_R(K - W_1) + \beta_D(K - W_2)$$

Dividing both sides by $K - W_2$, we have

$$K - W_1 \leq \beta_R \left(\frac{K - W_1}{K - W_2} \right) + \beta_D.$$

Observing that the left-hand side is less than $\beta_D + \beta_R$ (since $W_1 \geq W_S$), whereas the right-hand side is bigger than $\beta_R + \beta_D$ completes the verification step. ■

Lemma A.4 *Suppose $W_2 > \underline{W}$ and $W_1 \in [W_S, \underline{W}]$, then assigning both tasks to either agent is optimal. The social value function is given by*

$$S(W_1, W_2) = \pi - \rho_2(K - W_1) \frac{c(W_S)}{\rho_1} e^{-\rho_1(W_2 - \underline{W})} \quad (\text{A.25})$$

Proof. Suppose the principal assigns both tasks to Agent 2. So,

$$S = \pi - 2bS_2 \implies S = \pi - h(W_1)e^{-\rho_1 W_2} \quad (\text{A.26})$$

for some $h(W_1)$ to be determined. The optimal path moves vertically downward to the boundary $W_2 = \underline{W}$. From (A.21), the boundary condition is $S(W_1, \underline{W}) = \pi - (K - W_1)(K - \underline{W})\rho_2 c(W_S)$. Solving for $h(W_1)$ yields the expression stated in the lemma.

We next verify that the principal is indifferent between assigning both tasks to agent 1 or agent 2 (as conjectured above). If we assign both tasks to Agent 1, then

$$\beta_S(W_1)S_1 = (K - W_1)(-h'(W_1)e^{-\rho_1 W_2}).$$

For indifference, we require $(K - W_1)(-h'(W_1)) = h(W_1)$. Since $h(W_1) = (K - W_1)\rho_2 c(W_S) \frac{e^{\rho_1 \underline{W}}}{\rho_1}$, we have $h'(W_1) = -\rho_2 c(W_S) \frac{e^{\rho_1 \underline{W}}}{\rho_1}$. Thus,

$$-(K - W_1)h'(W_1) = (K - W_1)\rho_2 c(W_S) \frac{e^{\rho_1 \underline{W}}}{\rho_1} = h(W_1),$$

which confirms that $\beta_S(W_1)S_1 = \beta_S(W_2)S_2$. Finally, following the same steps as in the proof of Lemma A.3, shows that $\beta_S(W_2)S_2 \leq \beta_R S_2 + \beta_D S_1 \iff \theta \leq 1$, which completes the verification. ■

Lemma A.5 *If $W_1, W_2 \geq \underline{W}$, assigning both tasks to either agent is optimal. The social value function is given by*

$$S(W_1, W_2) = \pi - C e^{-\rho_1(W_1 + W_2)} \quad (\text{A.27})$$

where

$$C = \frac{\rho_2 \exp(2\rho_1 \underline{W}) c(W_S)}{\rho_1^2} \quad (\text{A.28})$$

Proof. Conjecture that both tasks are assigned to agent 2. Since, $\beta_S(W_2) = 2b$. The conjecture implies that S takes the form $S(W_1, W_2) = \pi - C(W_1)e^{-\rho_1 W_2}$. Using symmetry and the value function in Lemma A.4, the boundary condition at $W_2 = \underline{W}$ implies that $C(W_1) = k(\underline{W})e^{-\rho_1 W_1} e^{\rho_1 \underline{W}}$. Thus, $S = \pi - C e^{-\rho_1(W_1 + W_2)}$, where

$$C = k(\underline{W})e^{\rho_1 \underline{W}} = \frac{\rho_2 \exp(2\rho_1 \underline{W}) c(W_S)}{\rho_1^2}. \quad (\text{A.29})$$

Noting that $S_1 = S_2$ and $\beta_S(W_1) = \beta_S(W_2)$, whether both tasks are assigned to agent 1 or agent 2 is payoff irrelevant in this region. Moreover, since $\beta_S(W_1) = \beta_S(W_2) < \beta_D + \beta_R$, the verification step in this region is immediate. ■

B.2 Characterization of G and W_2^*

Suppose that agent 1's continuation value is W_1 when agent 2 is hired. How much starting continuation value should that principal give to agent 2 and to what tasks should each agent be assigned? To answer these questions, recall that the principal's value function is denoted by $V(W_1, W_2)$. Let $W_2^*(W_1)$ denote the continuation value for agent 2 that solves

$$W_2^*(W_1) \in \arg \max_{W_2} V(W_1, W_2) \quad (\text{A.30})$$

Denote $G(W_1) = V(W_1, W_2^*(W_1))$ as the principal's expected payoff conditional on hiring Agent 2 at $W_2 = W_2^*(W_1)$ when agent 1's utility is W_1 .

To characterize W_2^* and G we will make use of several properties of V . The first property is that the principal's value function is linear in W_2 on $\underline{\mathcal{A}}_2 = \{\vec{W} \in \mathcal{A}_2 \cap W_2 < \underline{W}\}$ with slope $c(W_1)$ that is decreasing in W_1 . This property is analogous to the single agent value function below $\underline{W}$. Hence for any W_1 , such that $c(W_1) > 1$, $W_2^*(W_1) > \underline{W}$. Let W_a be such that $c(W_a) = 1$. One can show that $W_a > 0$ if and only if $\frac{\pi}{K} - 1 > 0$, which is equivalent to the project being worth initiating if there were only a single agent.

Lemma A.6 *If $K < \pi$, then there exists a $W_a > 0$ such that for all $W_1 \in [0, W_a)$, $(W_1, W_2^*(W_1)) \in \mathcal{A}_2$, while $(W_1, W_2^*(W_1)) \in \mathcal{K}$ for all $W_1 \geq W_a \in \mathcal{W}_1$. If $K \geq \pi$, then $(W_1, W_2^*(W_1)) \in \mathcal{K}$ for all $W_1 \in \mathcal{W}_1$.*

Lemma A.7 *Suppose that $K < \pi$, then for all $W_1 \in [0, W_a)$*

$$W_2^*(W_1) = \frac{\ln(\rho_1 k(W_1))}{\rho_1} = \underline{W} + \frac{\ln(c(W_1))}{\rho_1} > \underline{W} \quad (\text{A.31})$$

and

$$G(W_1) = \pi - W_1 - K - \frac{\ln(c(W_1))}{\rho_1} \quad (\text{A.32})$$

Note that in this region, G is (1) decreasing in K (2) concave in W for $W \in [0, \theta W_S]$, and (3) linear and decreasing for $W \in [\theta W_S, W_S]$. These will be useful properties for characterizing the optimal contract.

Proof. Because the value function is concave in W_2 and since $V_2 = c(W_1) - 1 > 0$ in the linear region, it is also strictly positive in the kite region. Hence, the only possible solution to (A.30) involves $W_2 > \underline{W}$. The first-order condition in the concave region (of W_2) yields the expression for W_2^* . Plugging the solution back into V yields the expression for G . ■

Lemma A.8 Suppose $K < \pi$ and consider any $W_1 \geq W_a$

1. If $W_a < \theta W_S$, then

$$W_2^*(W_1) = \begin{cases} W_1/\theta & W_1 \in (W_a, \theta W_{sim}^*/(1+\theta)) \\ W_{sim}^* - W_1 & W_1 \in (\theta W_{sim}^*/(1+\theta), W_{sim}^*/(1+\theta)) \end{cases} \quad (\text{A.33})$$

and therefore

$$G(W_1) = \begin{cases} F_2(W_1(1+1/\theta)) & W_1 \in (W_a, \theta W_{sim}^*/(1+\theta)) \\ F_2(W_{sim}^*) & W_1 \in (\theta W_{sim}^*/(1+\theta), W_{sim}^*/(1+\theta)) \end{cases} \quad (\text{A.34})$$

2. If $W_a \in [\theta W_S, W_S]$, then for all $W_1 \in (W_a, W_{sim}^*/(1+\theta))$

$$W_2^*(W_1) = W_{sim}^* - W_1 \quad (\text{A.35})$$

$$G(W_1) = F_2(W_{sim}^*) \quad (\text{A.36})$$

Proof. First, for W_1 such that $c(W_1) < 1$, the solution for $W_2^*(W_1)$ cannot lie above the kite region since the derivative is negative in both the linear strip and the concave strip. Hence, the solution must lie in $\mathcal{K}$ or below it.

For $W_1 < W_{sim}^*/(1+\theta)$, $V_2(W_1, \theta W_1) = F_2'(W_1(1+\theta)) = \pi \rho_2 e^{-\rho_2 W_1(1+\theta)} - 1 > \pi \rho_2 e^{-\rho_2 W_{sim}^*} - 1 = 0$ and hence $W_2^*(W_1) > \theta W_1$ implying the solution must lie in $\mathcal{K}$. Within $\mathcal{K}$, $V(W_1, W_2) = F_2(W_1 + W_2)$ which is increasing (decreasing) in W_2 for $W_1 + W_2 < (>) W_{sim}^*$. Hence, either (1) $W_2^*(W_1)$ lies on the upper boundary (i.e., $W_2^*(W_1) = \max\{W_2 : (W_1, W_2) \in \mathcal{K}\}$), which occurs when $W_1 + \max\{W_2 : (W_1, W_2) \in \mathcal{K}\} \leq W_{sim}^*$ or (2) $W_1 + W_2^*(W_1) = W_{sim}^*$. To complete the proof, it suffices to show that $W_1 + \max\{W_2 : (W_1, W_2) \in \mathcal{K}\} < W_{sim}^*$ if and only if $W_a < \theta W_S$ and $W_1 \in (W_a, \theta W_{sim}^*/(1+\theta))$.

Suppose that $W_a < \theta W_S$, then from (A.13), we have

$$c(\theta W_S) = \pi \rho_2 e^{-\rho_2(1+\theta)W_S} < 1 \implies \frac{1}{\rho_2} \ln(\rho_2 \pi) = W_{sim}^* < (1+\theta)W_S \quad (\text{A.37})$$

In other words, $W_a < \theta W_S$ is equivalent to the ray $W_2 = W_{sim}^* - W_1$ intersecting the kite region on its left boundary (i.e., where $\max\{W_2 : (W_1, W_2) \in \mathcal{K}\} = W_1/\theta$). Further, since S has a downward kink along this boundary, the point of intersection must occur at $\theta W_{sim}^*/(1+\theta) > W_a$. Hence, for all $W_1 \in (W_a, \theta W_{sim}^*/(1+\theta))$, $W_1 + \max\{W_2 : (W_1, W_2) \in \mathcal{K}\} = W_1 + W_1/\theta < W_{sim}^*$. Whereas, for $W_1 \in (\theta W_{sim}^*/(1+\theta), W_{sim}^*/(1+\theta))$, $\max\{W_2 : (W_1, W_2) \in \mathcal{K}\} \geq W_{sim}^*/(1+\theta)$.

Conversely, if $W_a \geq \theta W_S$ then $W_{sim}^* \geq (1 + \theta)W_S$ implying the ray $W_2 = W_{sim}^* - W_1$ intersects the kite region along the top boundary (i.e., where $\max\{W_2 : (W_1, W_2) \in \mathcal{K}\} = W_S$). Because S is differentiable along the top boundary, $S_2(W_a, W_S) = c(W_a) = 1 \implies W_a + W_S = W_{sim}^*$ and hence for all $W_1 \in (W_a, W_{sim}^*/(1 + \theta))$, $W_1 + \max\{W_2 : (W_1, W_2) \in \mathcal{K}\} = W_1 + W_S > W_{sim}^*$. ■

B.3 Proofs for Section 5.2

Proof of Lemma 5. If the active branch of V at $(w, W_2^*(w))$ is solo-1, then

$$V(w, W_2^*(w)) = \pi - (w + W_2^*(w)) - \beta_S(w)(1 + V_1(w, W_2^*(w))).$$

At a point contact with differentiable $\bar{G}$, we have $\bar{G}(w) = V(w, W_2^*(w))$ and, by the envelope theorem, $\bar{G}'(w) = V_1(w, W_2^*(w))$. Substituting into Φ yields

$$\Phi(w) = \pi - w - \beta_S(w)(1 + V_1) - [\pi - (w + W_2^*(w)) - \beta_S(w)(1 + V_1)] = W_2^*(w) \geq 0.$$

Thus w lies in the waiting region. Suppose $\bar{G}$ is linear on $[a, b]$. If the active branch at both $(a, W_2^*(a))$ and $(b, W_2^*(b))$ is solo-1, part (1) applied at a and b gives $\Phi(a) \geq 0$ and $\Phi(b) \geq 0$. By constancy, $\Phi(w) = \Phi(a) = \Phi(b) \geq 0$ for all $w \in [a, b]$. Hence every $w \in [a, b]$ is in the waiting region. ■

Lemma A.9 Fix $K > 2b$ so that $\beta_S(W; K) = K - W$ on $W \in [0, W_S]$. Then, on $\mathcal{W}_1$:

1. $\Phi(W; K)$ is continuous and weakly increasing in W , and strictly increasing wherever $\bar{G}$ is strictly concave.
2. $\Phi(W; K)$ is continuous and strictly decreasing in K for every fixed W .

Proof of Lemma A.9. Recall

$$\Phi(W; K) = \pi - W - (K - W)(1 + \bar{G}_W(W; K)) - \bar{G}(W; K),$$

with $\bar{G}$ the concave envelope of $G(w) = \max_{y \geq 0} V(w, y)$. Because $\bar{G}$ is concave and piecewise C^1 , each point is either: (i) a *point contact* with G (so locally $\bar{G} = G$ and $\bar{G}_W = V_1(\cdot, W_2^*(\cdot; K))$), or (ii) lies on a *linear segment* $[a, b]$ with $\bar{G}(W; K) = c_0(K) + s(K)W$ and constant slope $s \in \partial \bar{G}$ given by the chord between two supporting endpoints where $\bar{G} = G$.

Weak monotonicity in W . Differentiate w.r.t. W wherever $\bar{G}$ is C^2 (point contact), and use the one-sided derivative on linear pieces:

$$\Phi_W(W; K) = -(K - W)\bar{G}_{WW}(W; K).$$

Since $\bar{G}$ is concave, $\bar{G}_{WW} \leq 0$, hence $\Phi_W \geq 0$, with *strict* inequality wherever $\bar{G}_{WW} < 0$. On any linear segment, $\bar{G}_{WW} \equiv 0$, so $\Phi_W \equiv 0$ there. Continuity in W follows because $\bar{G}$ and $\bar{G}_W$ match continuously at junctions of concave arcs and linear pieces.

Strict monotonicity in K . We show $\Phi_K(W; K) < 0$ for every $W \in \mathcal{W}_1$.

Case A: point contact at (W, K) . Here $\bar{G} = G$ locally and

$$\Phi_K(W; K) = -(1 + \bar{G}_W) - (K - W) \bar{G}_{WK} - \bar{G}_K. \quad (*)$$

By the envelope, $\bar{G}_W = V_1(W, W_2^*(W; K))$, and the two-agent HJB yields the slope bound $1 + V_1 > 0$ at every finite state, so the first term in $(*)$ is *strictly negative*. If the supporting region is $\mathcal{K}$ (“kite”), then $\bar{G}$ is locally independent of K , so $\bar{G}_K = \bar{G}_{WK} = 0$ and hence $\Phi_K = -(1 + \bar{G}_W) < 0$. If the supporting region is $\mathcal{A}_2$, direct differentiation of the closed form gives

$$\Phi_K(W; K) = -\left(1 + \frac{2b(W + K\theta^2 - 2\theta W)}{(W - K\theta)^2}\right) < 0.$$

Thus $\Phi_K < 0$ at all point-contact states in $\mathcal{K} \cup \mathcal{A}_2$.

Case B: linear segment at W . Let $[a(K), b(K)]$ be the supporting chord with $\bar{G}(W; K) = c_0(K) + s(K)W$ for $W \in [a, b]$. Using $\beta_S(W; K) = K - W$, we have

$$\Phi(W; K) = \pi - W - (K - W)(1 + s(K)) - (c_0(K) + s(K)W) = \pi - K - c_0(K) - s(K)K,$$

so $\Phi(\cdot; K)$ is *constant in W* on $[a, b]$. In particular, for any interior $W \in (a, b)$,

$$\Phi_K(W; K) = \Phi_K(a(K); K),$$

because $\Phi_W \equiv 0$ on the segment and the endpoints move continuously in K . At the endpoint $a(K)$ we have *point contact* ($\bar{G} = G$), and by construction of $\mathcal{W}_1$ the supporting endpoints of a segment used at $W \in \mathcal{W}_1$ lie in $\mathcal{K} \cup \mathcal{A}_2$. Therefore $\Phi_K(a(K); K) < 0$ by Case A, implying $\Phi_K(W; K) < 0$ throughout the linear segment.

Combining Cases A and B yields $\Phi_K(W; K) < 0$ for every $W \in \mathcal{W}_1$. ■

Proof of Lemma 6. We first claim that it is sufficient to show that (i) there exists a K_1 such that $\Phi(0; K_1) > 0$ and (ii) there exists K_2 such that $\Phi(\max(\mathcal{W}_1); K_2) < 0$. We then prove (i) and (ii).

Suppose that (i) and (ii) hold.

- For (1), $\Phi(0; K_1) > 0$ (by assumption) and $\Phi(0; K_2) < 0$ because we found K_2 such that $\Phi(\max(\mathcal{W}_1); K_2) < 0$ and $\Phi_W \geq 0 \Rightarrow \Phi(\max(\mathcal{W}_1)) \geq \Phi(0)$. Since Φ is continuous in K , the intermediate value theorem implies that there exists K_{seq} such that $\Phi(0; K_{\text{seq}}) = 0$. Because Φ is decreasing in K , $\Phi(0; K) > 0$ for all $K < K_{\text{seq}}$.
- For (3), by monotonicity in W , $\Phi(\max(\mathcal{W}_1); K_1) \geq \Phi(0; K_1) > 0 > \Phi(\max(\mathcal{W}_1); K_2)$. The intermediate value theorem guarantees the existence of K_{sim} . Monotonicity in K and W ensure that Φ is negative on $\mathcal{W}_1$ for all $K > K_{\text{sim}}$.
- For (2), for all $K \in (K_{\text{seq}}, K_{\text{sim}})$, $\Phi(0; K) < 0$ and $\Phi(\max(\mathcal{W}_1); K) > 0$. Continuity and monotonicity of Φ with respect to W_1 ensure that $\hat{W}_1$ is interior for all such K .

To prove (i), notice that if $K \leq \beta_R + \beta_D$, then $W_S = 0$ and $\mathcal{K} = \mathcal{A}_i = \emptyset$, so $(0, W_2^*(0)) \in \mathcal{E}$. By Lemma 5, $\Phi(0) = W_2^*(0) > 0$, where the strict inequality follows from Assumption 4, which ensures the project is worth initiating with a single agent since $K \leq \beta_R + \beta_D < \pi$. To prove (ii), since Φ is weakly increasing in W , it is sufficient to check that $\Phi(W) \leq 0$ in the interior region of $\mathcal{K}$. Since $\bar{G} = F_2^{\max}$ for all $W \in \text{int}(\mathcal{K})$, we have that $\Phi(W) = \pi - K - F_2^{\max}$. Hence, any $K > \pi - F_2^{\max}$ implies Φ is strictly negative on $\mathcal{W}_1$. ■

Proof of Proposition 5. On $(K_{\text{seq}}, K_{\text{sim}})$, Lemma 6 implies that the hiring threshold $\hat{W}_1(K) \in (0, \max(\mathcal{W}_1))$ is interior and uniquely defined. Moreover, at any interior solution $\Phi(\hat{W}_1(K)) = 0$. By the implicit function theorem

$$\hat{W}_1'(K) = -\frac{\Phi_K}{\Phi_W}|_{W=\hat{W}_1(K)} > 0.$$

By Lemma A.8, there exists a boundary point $\bar{W}_1 \in (0, \max(\mathcal{W}_1))$ such that

$$(W_1, W_2^*(W_1)) \in \mathcal{A}_2 \text{ for } W_1 < \bar{W}_1, \quad (W_1, W_2^*(W_1)) \in \mathcal{K} \text{ for } W_1 > \bar{W}_1.$$

Hence, for any $K \in (K_{\text{seq}}, K_{\text{sim}})$:

$$\text{Standby Developer} \iff \hat{W}_1(K) < \bar{W}_1, \quad \text{External Consultant} \iff \hat{W}_1(K) > \bar{W}_1.$$

By continuity of $\hat{W}_1(K)$ and the behavior near the endpoints of the staggered region, we know that for K just above K_{seq} we have $\hat{W}_1(K) < \bar{W}_1$ (Standby Developer), and for K just below K_{sim} we have $\hat{W}_1(K) > \bar{W}_1$ (External Consultant). Since $\hat{W}_1(K)$ is weakly increasing in K , the set

$$\{K \in (K_{\text{seq}}, K_{\text{sim}}) : \hat{W}_1(K) < \bar{W}_1\}$$

is an interval of the form $(K_{\text{seq}}, \tilde{K})$, and the set

$$\{K \in (K_{\text{seq}}, K_{\text{sim}}) : \hat{W}_1(K) > \bar{W}_1\}$$

is an interval of the form $(\tilde{K}, K_{\text{sim}})$, for some $\tilde{K} \in (K_{\text{seq}}, K_{\text{sim}})$.

For $K \in (K_{\text{seq}}, \tilde{K})$ we therefore have $(\hat{W}_1(K), W_2^*(\hat{W}_1(K))) \in \mathcal{A}_2$, so the Standby Developer structure is optimal. For $K \in (\tilde{K}, K_{\text{sim}})$ we have $(\hat{W}_1(K), W_2^*(\hat{W}_1(K))) \in \mathcal{K}$, so the External Consultant structure is optimal. This proves the proposition. ■

Proof of Proposition 6. Fix $\epsilon \in (0, \theta)$. We prove optimality within each organizational regime using the principal's HJB and the incentive constraints.

Solo innovator. Let W^S denote the promised continuation utility to the solo innovator.

To implement full effort and truthful reporting, it is sufficient that (β, γ) satisfy:

$$\begin{aligned} (1 - \epsilon)\beta + \epsilon\gamma &\geq 2b && \text{(double shirking),} \\ (\theta - \epsilon)(\beta - \gamma) &\geq b && \text{(passive shirking),} \\ \gamma &\geq (1 - q)\beta - qW^S && \text{(truthful lemons),} \\ \beta &\geq \gamma && \text{(truthful peaches).} \end{aligned}$$

Define $K_\epsilon := \frac{b}{q(\theta - \epsilon)}$ and $(\bar{\beta}_\epsilon, \bar{\gamma}_\epsilon)$ as in the statement, and set

$$\bar{W}_\epsilon := K_\epsilon - \bar{\beta}_\epsilon = \frac{(1 - \epsilon q)b}{q(\theta - \epsilon)} - 2b.$$

To justify which constraints bind, consider the principal's HJB under on-path effort, written in terms of the solo-innovator value F :

$$F(W^S) = \max_{(\beta, \gamma) \in \Gamma(W^S)} \left\{ -2c - ((1 - \epsilon)\beta + \epsilon\gamma) F'(W^S) + (1 - \epsilon)(R - W^S - \beta) + \epsilon F(W^S + \gamma) \right\},$$

where $\Gamma(W^S)$ is the set defined by the IC constraints above. Assuming F is concave and $1 + F'(W) \geq 0$ (total surplus increases with W), the objective is weakly decreasing in β holding γ fixed (since $\partial/\partial\beta = -(1 - \epsilon)(1 + F'(W^S)) \leq 0$), and weakly decreasing in γ holding β fixed (since $\partial/\partial\gamma = \epsilon(F'(W^S + \gamma) - F'(W^S)) \leq 0$ by concavity). Thus the principal wants to total incentives as small as possible. When the double shirking constraint binds, the principal can trade off β for γ while keeping $(1 - \epsilon)\beta + \epsilon\gamma = 2b$; along that locus, the principal prefers shifting weight toward γ (since marginal cost of increasing β is -1 while the marginal cost of increasing γ is $F'(W^S + \gamma) \geq -1$).

Low W^S ($W^S < \bar{W}_\epsilon$). When W^S is low, truthful lemons must bind at the optimum; otherwise one could reduce γ without violating any other constraint and strictly relax the feasible set. With $\gamma = (1 - q)\beta - qW^S$, passive-shirking becomes $(\theta - \epsilon)(\beta - (1 - q)\beta + qW^S) = q(\theta - \epsilon)(\beta + W^S) \geq b$, so the smallest feasible β is $\beta = K_\epsilon - W^S$ and $\gamma = (1 - q)K_\epsilon - W^S$.

High W^S ($W^S \geq \bar{W}_\epsilon$). When W^S is high, truthful lemons can be slack. Then the HJB logic above implies that the optimal contract minimizes incentives subject to double shirking and passive shirking, with the tradeoff between β and γ governed by these constraints. In particular, the optimal pair has $(1 - \epsilon)\beta + \epsilon\gamma = 2b$ and $\beta - \gamma = b/(\theta - \epsilon)$ (double shirking and passive shirking bind), which solves to the constant pair $(\bar{\beta}_\epsilon, \bar{\gamma}_\epsilon)$. Truthful lemons is slack precisely when $W^S \geq \bar{W}_\epsilon$, i.e., $\bar{\gamma}_\epsilon \geq (1 - q)\bar{\beta}_\epsilon - qW^S$. Combining the two regions gives the max-formulas stated in the proposition.

Two-agent team. Let (W_D, W_R) be promised utilities. The researcher's passive-shirking IC is $(\theta - \epsilon)(\beta_\epsilon^R - \gamma_\epsilon^R) \geq b$. From the two-agent HJB, concavity implies the principal weakly prefers smaller continuation-value rewards for the researcher, so the optimal choice sets $\gamma_\epsilon^R = 0$ and makes the constraint bind: $\beta_\epsilon^R = b/(\theta - \epsilon)$.

For the developer, effort requires $(1 - \epsilon)\beta_\epsilon^D + \epsilon\gamma_\epsilon^D \geq b$ and truthful peaches requires $\beta_\epsilon^D \geq \gamma_\epsilon^D$. Along a binding effort constraint, the HJB implies the principal prefers shifting weight from transfers toward continuation utility, but $\beta_\epsilon^D \geq \gamma_\epsilon^D$ limits this substitution, so the optimum sets $\beta_\epsilon^D = \gamma_\epsilon^D = b$. Truthful lemons is then slack for all $W_D \geq 0$ because

$\gamma_\epsilon^D = b \geq (1-q)b - qW_D$. Thus the stated two-agent contract implements effort and truthful reporting and is optimal within the two-agent regime. ■

B.4 Sufficient Conditions for Research

In this section, we derive a sufficient condition under which the principal always finds it optimal to implement effort on research. We begin with the proof of Lemma 8, which characterizes the optimal contract that implements effort on development and truthful reporting. **Proof of Lemma 8.** Ignoring terms independent of (β, γ) , the principal solves

$$\max_{\beta, \gamma} -(1-\theta)\beta F'(W) - \theta\gamma F'(W) - (1-\theta)\beta + \theta F'(W + \gamma)$$

subject to

$$(1-\theta)\beta + \theta\gamma \geq b, \quad \beta \geq \gamma, \quad \gamma \geq (1-q)\beta - qW.$$

Let $\delta = (1-\theta)\beta + \theta\gamma$. Since $F'(W) \geq -1$, the objective is increasing in δ , so $\delta = b$. Holding δ fixed, the objective is increasing in γ , giving $\gamma = \delta = b$, and the remaining constraint implies $\beta = \gamma = b$. ■

Allowing for a solo developer, the principal's HJB in any state becomes:

$$V(w_1, w_2) = \max\{S_1, S_2, RD, DR, D_1, D_2\}(w_i, w_j),$$

where S_i denotes the *solo-innovator* branch for agent i , D_i denotes the *solo-developer* branch for agent i , and RD (DR) denotes the branch in which agent 1 is the researcher (developer) and agent 2 is the developer (researcher). Our goal is to show that for sufficiently high θ , the principal never finds it optimal to use a solo-developer branch. It suffices to prove that S_i dominates D_i for all (w_i, w_j) . A similar comparison using the two-agent team (RD or DR) yields the weaker condition $\theta \geq 1 - c/R$. Since the solo-innovator bound derived below implies a strictly smaller threshold, the team comparison does not tighten the sufficient condition.

Lemma A.10 *Suppose V is concave in each argument. Then the solo-innovator branch (S_i) dominates the solo-developer branch (D_i) at every state if*

$$\theta \geq 1 - \frac{c}{R - \frac{b}{\beta_0}(R - 2c)}.$$

Proof. Fix w_j and write $v(w_i) = V(w_i, w_j)$. From S_i ,

$$v(w_i) = (R - 2c) - (w_i + w_j) - \beta_S(w_i)(1 + v'(w_i)).$$

Evaluating at $w_i = 0$ and using $v(0) = F_1(w_j) \geq 0$,

$$\beta_0(1 + v'_+(0, w_j)) = (R - 2c) - (w_j + F_1(w_j)) \leq R - 2c,$$

so

$$1 + V_i(w_i, w_j) = 1 + v'(w_i) \leq 1 + v'_+(0, w_j) \leq \frac{R - 2c}{\beta_0}.$$

In the solo-developer branch (Lemma 8),

$$D_i = (1 - \theta)(R - (w_i + w_j + b)) - c - bV_i(w_i, w_j) + \theta V(w_i + b, w_j).$$

By concavity in w_i ,

$$V(w_i + b, w_j) \leq V(w_i, w_j) + bV_i(w_i, w_j),$$

so

$$D_i \leq V(w_i, w_j) + (1 - \theta)\tilde{D}_i - c, \quad \tilde{D}_i := (R - (w_i + w_j + b)) - V - bV_i.$$

Using the solo-operator identity $V + \beta_S(w_i)V_i = R - 2c - (w_i + w_j) - \beta_S(w_i)$, we obtain

$$\tilde{D}_i = 2c + (\beta_S(w_i) - b)(1 + V_i(w_i, w_j)) \leq 2c + (\beta_0 - b)\frac{R - 2c}{\beta_0} = R - \frac{b}{\beta_0}(R - 2c) =: D_{\max}.$$

Thus for every (w_i, w_j) , $D_i(w_i, w_j) \leq V(w_i, w_j) + (1 - \theta)D_{\max} - c$, so $S_i \geq D_i$ everywhere whenever $(1 - \theta)D_{\max} \leq c$, i.e.

$$\theta \geq 1 - \frac{c}{D_{\max}},$$

which is the claimed condition. ■

References

- Admati, A. R. and M. Perry (1991). Joint projects without commitment. *Review of Economic Studies* 58(2), 259–276.
- Bolton, P. and C. Harris (1999). Strategic experimentation. *Econometrica* 67(2), 349–374.
- Bonatti, A. and J. Horner (2011). Collaborating. *American Economic Review* 101(2), 632–663.
- Brown, T. (2008). Design thinking. *Harvard Business Review* 86(9), 62–72.
- Cetemen, D., I. Hwang, and A. Kaya (2020). Uncertainty-driven cooperation. *Theoretical Economics* 15(3), 1023–1058.
- Che, Y.-K., E. Iossa, and P. Rey (2021). Prizes versus contracts as incentives for innovation. *Review of Economic Studies* 88(5), 2149–2178.
- Che, Y.-K. and S.-W. Yoo (2001). Optimal incentives for teams. *American Economic Review* 91(3), 525–541.
- Georgiadis, G. (2015). Projects and team dynamics. *American Economic Review* 81(1), 187–218.
- Gershkov, A. and M. Perry (2012). Dynamic contracts with moral hazard and adverse selection. *Review of Economic Studies* 79(1), 268–306.
- Green, B. and C. R. Taylor (2016). Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *American Economic Review* 106(12), 3660–99.
- Gromb, D. and D. Martimort (2007). Collusion and the organization of delegated expertise. *Journal of Economic Theory* 137(1), 271–299.
- Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *Review of Economic Studies* 83(3), 1040–1091.
- He, Z., B. Wei, J. Yu, and F. Gao (2017). Optimal long-term contracting with learning. *Review of Financial Studies* 30(6), 2006–2065.
- Iossa, E. and D. Martimort (2012). Risk allocation and the costs and benefits of public-private partnerships. *RAND Journal of Economics* 43(3), 442–474.
- Iossa, E. and D. Martimort (2015). The simple microeconomics of public-private partnerships. *Journal of Public Economic Theory* 17(1), 4–48.
- Keller, G., S. Rady, and M. Cripps (2005). Strategic experimentation with exponential bandits. *Econometrica* 73(1), 39–68.
- Laux, C. (2001). Limited-liability and incentive contracting with multiple projects. *RAND Journal of Economics*, 514–526.

- Levy, S. (2011). *In the Plex: How Google Thinks, Works, and Shapes Our Lives*. New York: Simon & Schuster.
- Liu, B. and J. Lu (2018). Pairing provision price and default remedy: optimal two-stage procurement with private r&d efficiency. *RAND Journal of Economics* 49(3), 619–655.
- Martimort, D. and J. Pouyet (2008). To build or not to build: Normative and positive theories of public–private partnerships. *International Journal of Industrial Organization* 26(2), 393–411.
- Marx, L. M. and S. A. Matthews (2000). Dynamic voluntary contribution to a public project. *Review of Economic Studies* 67(2), 327–358.
- O’Reilly, C. A. and M. L. Tushman (2004). The ambidextrous organization. *Harvard Business Review* 82(4), 74–81.
- Pavlak, A. (2004). Project troubleshooting: Tiger teams for reactive risk management. *Project Management Journal*.
- Schmitz, P. W. (2005). Allocating control in agency problems with limited liability and sequential hidden actions. *RAND Journal of Economics*, 318–336.
- Sobek, D. K., A. C. Ward, and J. K. Liker (1999). Toyota’s principles of set-based concurrent engineering. *Sloan Management Review*.
- Varas, F. (2018). Managerial short-termism, turnover policy, and the dynamics of incentives. *Review of Financial Studies* 31(9), 3409–3451.
- Yildirim, H. (2006). Getting the ball rolling: Voluntary contributions to a large-scale public project. *Journal of Public Economic Theory* 8(4), 503–528.
- Yildirim, H. (2023). Who fares better in teamwork? *RAND Journal of Economics* 54(2), 299–324.