

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

Felix Z. Feng[†] Brett Green[‡] Curtis R. Taylor[§] Mark Westerfield[¶]

October 1, 2026

Abstract

We study the optimal allocation of tasks over the life cycle of an innovative project. Innovation requires two complementary activities, research and development, both subject to moral hazard. Agents also privately observe the quality of discoveries, so the principal must elicit truthful reports to avoid costly mistakes. Assigning both tasks to a single “solo operator” avoids free-riding but creates a double-deviation problem: the agent can shirk on research and conceal the resulting low-quality discovery. As time passes without a discovery, this deviation becomes increasingly attractive. Separating the tasks across two agents eliminates the double deviation but introduces free-riding. We characterize how this tradeoff shapes organization over time. When low-quality discoveries are readily exposed, the project is run by sequential solo operators; when they are difficult to expose, a two-agent team is optimal from the outset. At intermediate levels of detectability, the project begins with a solo operator and switches endogenously to a team.

JEL Classification: D82, D86, D23, O31

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

*We thank Hengjie Ai, Itay Goldstein, Zhiguo He, Ron Kaniel, Simon Mayer, Marzena Rostek, Wenyu Wang, Lucy White, seminar participants at Berkeley, University of Georgia, University of Illinois Chicago, and Washington University in St. Louis, participants at the 2026 Econometric Society NASM and the 2026 Finance Theory Group (FTG) Member Meeting at Carnegie Mellon University for helpful comments and suggestions.

[†]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

“But of these two thynges he wolde determyne none
Without aide. For two hedds are better than one.”

—John Heywood, 1546

1 Introduction

Innovation requires two broad activities: searching for promising ideas and transforming them into working products or services. However, organizations differ not only in how they divide these activities, but also in when they do so. Some projects begin as a “one-person shop,” with a founder or lone engineer exploring the technological frontier while also building the first working version of the product, and divide responsibilities only as the project evolves (Levy, 2011). Others separate exploratory research from execution-oriented development from the outset, assigning the activities to units with distinct objectives and evaluation criteria (O’Reilly and Tushman, 2004). Still others expand conditionally as progress unfolds, bringing in short-term consultants to diagnose setbacks (Brown, 2008) or activating specialized “tiger teams” when an initial approach disappoints (Pavlak, 2004). These arrangements reflect different choices about whether to concentrate responsibility in a generalist, rely on specialists, or add expertise only after the project reaches a critical stage.

These patterns pose an organizational-design puzzle. Technological comparative advantage can explain why different tasks are assigned to different specialists, but it cannot by itself explain why an organization reallocates unchanged tasks across equally capable agents as a project unfolds. Incentive considerations provide a distinct explanation. Effort in innovation is difficult to observe, and the people producing a discovery often learn its quality before the principal decides whether to accept it; a flaw may become verifiable only later through the realized payoff. Task allocation therefore determines who controls effort, who possesses information, and which deviations are feasible. Bundling both tasks with one agent can economize on incentive payments, but it allows that agent to combine shirking with misreporting. Separating the tasks limits this joint deviation but creates scope for free-riding. We therefore ask how a principal should allocate research and development, when a second agent should be hired, and whether changing incentive conditions alone can generate a transition from a solo operator to a specialized team.

To isolate the organizational tradeoff, we begin with a static benchmark stripped of any technological reason for specialization. A project requires two complementary tasks: *research*, which increases the likelihood that a discovery is high quality, and *development*,

which produces the discovery itself. We call a high-quality discovery a “peach” and a low-quality discovery a “lemon.” A peach surely yields a positive payoff when accepted. A lemon yields the positive payoff with probability $1 - q$, but with probability q acceptance generates a large loss, exposing low quality. Thus, q measures the ex post verifiability of quality: when q is high, lemons are easy to expose; when q is low, they are difficult to distinguish from peaches. The principal can bundle both tasks with a single *solo operator* or separate them between a researcher and a developer. The agents are identical and equally productive at either task, so the two organizational forms are technologically equivalent. Any difference between them arises entirely from incentives.

With a solo operator, the same agent controls research effort, development effort, and the report of quality. This concentration of responsibility avoids free-riding across agents, but it creates a particularly profitable joint deviation. The agent can shirk on research, continue working on development, and, if the resulting discovery is a lemon, report it as a peach. We call this combination of shirking and misreporting the *double deviation*. When q is high, such deception is likely to be exposed, and the relevant constraint is instead *double shirking* – shirking on both tasks. In that region, the principal need only compensate the agent for the private benefits he forgoes by working; i.e., his opportunity cost. As q falls, however, endeavoring to pass off a lemon as a peach becomes increasingly attractive. Deterring the double deviation then requires additional rents, so the incentive cost of a solo operator rises as quality becomes harder to verify.

Separating the tasks removes this information rent altogether. Suppose the researcher controls the effort that affects quality while the developer generates the discovery and reports its quality. Because the developer’s effort affects whether a discovery is produced but not whether it is a peach or a lemon, the principal can give him the same compensation following either report. Truthful reporting therefore requires no additional rent. At the same time, the researcher no longer controls the quality report and hence cannot conceal shirking by misreporting a lemon. In fact, the result is stronger: the minimum cost of employing the two-agent team is exactly the same as it would be if quality were publicly observable. *Separating the tasks makes private information effectively costless to the principal.*

This informational advantage comes at a price, however. Research improves the probability that a discovery is a peach but is not necessary for a discovery to be produced. A shirking researcher can therefore free-ride on the developer: development still produces a discovery, and even without research it may turn out to be a peach. The researcher consequently retains some prospect of being rewarded even when he shirks, so inducing research effort requires compensation above his private benefit from shirking. Unlike the solo operator’s information rent, this free-riding cost does not depend on q . The static organizational choice

therefore reduces to a transparent tradeoff. When lemons are easy to expose, bundling is cheaper because the solo operator avoids free-riding and misreporting is inexpensive to deter. When lemons are difficult to expose, separation is cheaper because the team eliminates the information rent associated with the double deviation. Thus, whether two heads are better than one depends on whether the informational benefit of separating tasks outweighs the free-riding cost it creates.

The static benchmark isolates the two forces that govern organizational design: bundling research and development avoids free-riding but creates an information rent, while separating the tasks eliminates that information rent at the cost of free-riding. Our main model adds the dynamic feature that makes this organizational choice evolve over the life of an innovative project. Discoveries arrive according to a Poisson process as long as development effort continues. Continuous research effort ensures that an arriving discovery is a peach; without research, the discovery may instead be a lemon. The key new element is that the incentive cost of a solo operator now depends on the history of the project. Early on, the prospect of future employment gives the solo operator valuable continuation utility. Because passing off a lemon as a peach risks exposure, termination, and the loss of this future utility, continuation value acts as a *bond* that can deter the double deviation at no additional cost. But this bond is gradually depleted as the project continues without a discovery. As the solo operator's continuation value falls, misreporting becomes harder to deter and eventually requires additional rents. By contrast, a two-agent team continues to eliminate the information problem, while its free-riding cost remains stationary. Dynamics therefore transform the static comparison into an endogenous organizational life cycle: a solo operator can be optimal early in a project, when continuation value provides effective discipline, yet become more expensive than a team after a sufficiently long spell without success.

This history dependence generates three distinct organizational regimes. When lemons are easy to detect ex post, the project is best organized around a succession of solo innovators: one agent performs both research and development until he is replaced by another. When lemons are difficult to detect, the informational cost of bundling the tasks is sufficiently severe that the principal instead begins immediately with a two-agent team, assigning research and development to different agents. The most interesting case arises at intermediate levels of detectability. The project begins with a solo innovator, whose continuation value initially provides an effective bond against misreporting. As that bond erodes, however, the principal brings in a second agent before replacing the first, and the organization transitions endogenously from a solo operation to a specialized team. When the second agent is brought in, he takes over research, resembling an external consultant, while the incumbent agent is relegated to development only. Thus, the model predicts not only whether innovation is

organized around one head or two, but also *when* a second head should enter as the project evolves and *how* responsibilities should be reallocated during the transition.

We make three contributions to the theory of organizational design. First, we identify a new source of history dependence in organizational structure. Even with stationary technology, unchanged tasks, and ex ante identical agents, the optimal organization evolves over the life of a project. The key state variable is the incumbent’s continuation utility: the future stake that can be used to discipline current behavior. As unsuccessful search erodes this stake, a solo operator becomes progressively more expensive to incentivize, eventually making a change in organizational form optimal. Performance history therefore affects organization not through learning about project quality, worker ability, or technology, but through the endogenous depletion of incentive capacity. This mechanism generates transitions from a generalist organization to a specialized team even though the underlying production environment never changes.

Second, we uncover a fundamental incentive tradeoff between bundling and separating tasks when both actions and information are private. Bundling research and development with a single agent avoids free-riding across workers, but it also allows that agent to control both the quality of a discovery and its report. This creates a joint deviation in which he shirks on research and hopes to pass off a low-quality discovery as a high-quality one. Separating the tasks eliminates this deviation: because the reporting agent does not control quality, truthful reporting requires no information rent. Indeed, under a two-agent team, private information about discovery quality becomes effectively public from the principal’s perspective. The cost of separation is instead free-riding in research, since a researcher who shirks may still benefit from discoveries produced by the developer. Organizational design therefore trades off the information rent created by bundling against the free-riding cost created by separation.

Finally, the model yields a set of practical predictions about when organizations should rely on generalists, specialists, and transitional arrangements between the two. When low-quality discoveries are likely to be exposed ex post, the threat of losing future continuation value provides effective discipline, allowing the organization to rely longer on solo innovators. When flawed output is difficult to detect, this disciplinary force is weaker, specialization occurs earlier, and a two-agent team may be optimal from the outset. The model also predicts how an additional worker should be deployed during the transition: as a research specialist because research is more costly to incentivize. These predictions distinguish our mechanism from explanations based on comparative advantage, learning-by-doing, or changes in the underlying production technology.

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.

[Fulghieri, Hu, and Varas \(2025\)](#) study, in the context of venture capital stage financing, whether it is optimal to assign monitoring and screening to one investor or separate them across two investors. Similar to the aforementioned studies, their model is also static with three dates, and the mechanism favoring separation differs materially from ours, arising from the differences in discounting between active and passive investors. Separation can be

more efficient than integration because the former allows part of the compensation to the impatient investors to be paid earlier at an interim date. If active and passive investors are equally patient, the entrepreneur is indifferent between the two arrangements whenever both are feasible. In comparison, we consider both a static benchmark and a dynamic environment, where the latter allows us to demonstrate how the optimal organizational design can endogenously vary over time, transitioning from one organization to another even when the underlying production technology and the agents’ expertise remain stationary. All parties in our model are equally patient, so the comparison between a solo operator and a two-agent team stems from the total incentives required for each arrangement. Finally, our study focuses on the design of internal organizations while [Fulghieri, Hu, and Varas \(2025\)](#) focuses on the design of securities in VC financing.

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 (e.g., 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, e.g., [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 aforementioned 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

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.

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 \(2016\)](#), [Liu and Lu \(2018\)](#), [Varas \(2018\)](#), [Che, Iossa, and Rey \(2021\)](#), [Feng, Taylor, Westerfield, and Zhang \(2024\)](#), 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 (a 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 highlights the static tradeoff between employing a single agent and an agent team. 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. Additional derivations and proofs are in the Appendix.

2 Static Benchmark

2.1 Basic Environment

There are two dates, $t \in \{0, 1\}$, a principal, and two ex ante identical agents. The principal undertakes a project requiring two complementary tasks, which we call *research* and *development*. Either agent is equally capable of performing either task, or both. Thus, there is no technological gain from specialization; any advantage of assigning the tasks to different agents arises entirely from incentives. 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 either or both tasks, each at cost c . Development requires the assigned agent to choose unobservable effort $e \in \{s, w\}$. If he works, the project produces a discovery at $t = 1$; if he shirks, he obtains private benefit $b < c$ and no discovery is generated.

Research affects the quality rather than the arrival of the discovery. The agent assigned to research chooses unobservable effort $a \in \{s, w\}$. If he shirks, he obtains private benefit b , and a discovery generated by development is a high-quality “peach” with probability $\lambda_0 \in (0, 1)$. If he works, he forgoes this private benefit and raises the probability of a peach to $\lambda_1 \in (\lambda_0, 1)$. Let

$$\theta \equiv \lambda_1 - \lambda_0 > 0 \tag{1}$$

denote the improvement in discovery quality generated by research.

The principal can assign both tasks to a single agent, whom we call a *solo operator*, or assign research and development to different agents, whom we call the *researcher* and *developer*.

At $t = 1$, if a discovery is generated, the agent or agents observe its quality and report it to the principal. The principal then decides whether to accept it. An accepted peach yields $R > 0$. An accepted low-quality “lemon” yields R with probability $1 - q$ but generates a loss $-L < 0$ with probability q , where $q \in (0, 1]$. Thus, q measures the ex post detectability of a lemon: a larger q makes it more difficult to pass off a lemon as a successful discovery.

We impose the following assumption.

Assumption 1 *Lemon discoveries should not be accepted: $(1 - q)R < qL$. Research is socially optimal: $\theta R > c$.*

To isolate the organizational tradeoff, we restrict attention to incentive-compatible contracts under which the principal invests in both tasks, the assigned agents exert effort, and quality is reported truthfully.² On the equilibrium path, a discovery is therefore generated with probability one and accepted if and only if it is a peach. When indifferent, an agent exerts effort and reports truthfully.

Before solving for optimal contracts, it is useful to preview the organizational tradeoff. Bundling the tasks avoids a free-riding problem because one agent bears responsibility for both research and development. But it also gives that agent control over both the quality of the discovery and the report of its quality. He can therefore combine two deviations: shirk on research and, if a lemon results, conceal his shirking by reporting the lemon as a peach. Separating the tasks eliminates this *double deviation*, because the agent whose effort determines quality need not control the quality report. The price of separation is free-riding:

²Development is essential, so research is the only optional task. This static comparison is conditional on implementing both tasks; Appendix B.4 studies the research decision in the dynamic model when only one agent remains.

a researcher who shirks can still benefit from a discovery produced by the developer. The optimal organizational form is determined by which of these two incentive costs is larger.

2.2 Analysis

A. Solo Operator

Suppose first that a single agent performs both tasks. A contract is a pair (β, γ) , where the solo operator receives β if he reports a peach and the accepted discovery yields R , and receives γ if he reports a lemon.

Following the prescribed actions, his expected compensation is

$$U(\beta, \gamma) \equiv \lambda_1 \beta + (1 - \lambda_1) \gamma. \quad (2)$$

There are two economically relevant deviations. First, the agent can shirk on both tasks, obtaining private benefit $2b$ and producing no discovery. Deterring this double shirking requires

$$U(\beta, \gamma) \geq 2b. \quad (\text{Double-Shirking IC})$$

Second, the agent can shirk on research, obtain private benefit b , continue to work on development, and then report the resulting discovery strategically. Deterring this deviation requires

$$U(\beta, \gamma) \geq b + \lambda_0 \beta + (1 - \lambda_0) \max\{\gamma, (1 - q)\beta\}. \quad (3)$$

Truthful reporting of a lemon separately requires

$$\gamma \geq (1 - q)\beta. \quad (\text{Truth-telling IC})$$

At the optimum, the truth-telling constraint binds. Substituting $\gamma = (1 - q)\beta$ into (3) yields

$$U(\beta, \gamma) \geq b + \lambda_0 \beta + (1 - \lambda_0)(1 - q)\beta. \quad (\text{Double-Deviation IC})$$

The deviation on the right-hand side combines shirking on research with misreporting: the agent works on development but reports every discovery as a peach. We therefore call it the *double deviation*. Bundling the two tasks is valuable because it avoids free-riding, but costly because it makes precisely this joint deviation feasible.

The principal's problem is consequently to minimize compensation while deterring both double shirking and the double deviation.

Lemma 1 (Cost of a solo operator) *Let*

$$D(q) \equiv \lambda_1 + (1 - \lambda_1)(1 - q). \quad (4)$$

For a solo operator, the minimum expected compensation for effort on both tasks and truthful reporting is

$$C^{solo}(q) = \max \left\{ 2b, \frac{bD(q)}{q\theta} \right\}. \quad (5)$$

This cost is implementable by

$$\beta^* = \max \left\{ \frac{2b}{D(q)}, \frac{b}{q\theta} \right\}, \quad \gamma^* = (1 - q)\beta^*. \quad (6)$$

The two terms in $C^{solo}(q)$ correspond directly to the two deviations above. They are equal at

$$\bar{q} \equiv \frac{1}{1 + \lambda_1 - 2\lambda_0}. \quad (7)$$

Whenever $\bar{q} \leq 1$,

$$C^{solo}(q) = \begin{cases} \frac{bD(q)}{q\theta}, & q < \bar{q}, \\ 2b, & q \geq \bar{q}. \end{cases} \quad (8)$$

When q is high, an accepted lemon is likely to be exposed. Misreporting is therefore unattractive, and double shirking is the relevant deviation. The principal can then induce both tasks at expected cost $2b$, exactly equal to the agent's total private benefit from shirking. In this region, bundling achieves the lowest possible incentive cost.

When q is low, by contrast, a lemon can more easily masquerade as a peach. Shirking on research while continuing development and concealing a bad discovery becomes the more attractive deviation. The principal must then leave the solo operator an information rent to induce both research effort and truthful reporting. Accordingly, $C^{solo}(q)$ is strictly decreasing in q on this branch: making quality harder to verify raises the cost of bundling the tasks.

B. Two-Agent Team

Now suppose that research and development are assigned to different agents. Both agents observe quality, and the principal may solicit the quality report from either or both of them. The minimum cost is the same regardless of how the reporting responsibility is assigned.

Lemma 2 (Cost of a two-agent team) *The minimum expected compensation required to induce effort on both tasks and truthful reporting from a two-agent team is*

$$C^{team} = \left(\frac{2\lambda_1 - \lambda_0}{\lambda_1 - \lambda_0} \right) b > 2b. \quad (9)$$

No alternative allocation of reporting responsibility or report-contingent compensation can lower this cost.

To see the source of this cost, suppose the developer reports quality. Because his effort determines whether a discovery is produced but has no effect on its quality, the cheapest way to induce both development effort and truthful reporting is

$$\beta_D = \gamma_D = b. \quad (10)$$

The developer is compensated for his foregone shirking benefit but is given no reason to prefer reporting a peach to reporting a lemon.

For the researcher, it is optimal to set $\gamma_R = 0$. His incentive constraint is then

$$\lambda_1 \beta_R \geq b + \lambda_0 \beta_R. \quad (11)$$

At the optimum it binds, giving

$$\beta_R = \frac{b}{\theta} > b. \quad (12)$$

Expected total compensation is therefore

$$C^{team} = \lambda_1(\beta_R + b) + (1 - \lambda_1)b = \left(\frac{2\lambda_1 - \lambda_0}{\lambda_1 - \lambda_0} \right) b. \quad (13)$$

Separation thus has two sharply contrasting effects. Its benefit is informational. The developer controls the report but not the probability that the discovery is a lemon, so the principal can make him indifferent across reports. The researcher controls quality but not the report. Neither agent can combine shirking on research with concealing the resulting

lemon. Indeed, C^{team} is exactly the cost the principal would incur if quality were publicly observable: private information creates no additional incentive cost under separation.

Its cost is free-riding. Research is valuable because it raises the probability that a discovery produced by the developer is a peach, but research is not itself necessary for a discovery to occur. A shirking researcher therefore still benefits from the developer's effort and receives $\lambda_0\beta_R$ with positive probability. Inducing research must compensate him not only for the private benefit b from shirking, but also for this foregone chance of being rewarded while free-riding. Consequently, $C^{team} > 2b$.

Crucially, this free-riding cost is independent of q . Separation removes the connection between the researcher's hidden action and the agent responsible for reporting quality, so making lemons harder to detect does not make the team more expensive to incentivize.

2.3 Cost Comparison

Because both organizational forms implement the same equilibrium actions, the principal's investment costs and expected project revenue are identical under bundling and separation. Organizational design therefore reduces to a comparison of $C^{solo}(q)$ and C^{team} .

Proposition 1 *Define*

$$q^* = \frac{1}{1 + \theta}. \tag{14}$$

The principal strictly prefers a two-agent team if $q < q^$ and strictly prefers a solo operator if $q > q^*$.*

Figure 1 illustrates the comparison, which has a simple interpretation. A solo operator avoids free-riding, but when q is low the principal must pay an information rent to prevent him from shirking on research and misreporting. A team eliminates this information rent, but splitting the tasks creates a fixed free-riding cost. The optimal organization trades off these two distortions. When lemons are easy to expose, the information problem is small and bundling dominates. When lemons are difficult to expose, the information rent becomes sufficiently large that separation is cheaper.

The cutoff in Proposition 1 has two useful features. First,

$$q^* < \bar{q}. \tag{15}$$

Thus, the organizational switch occurs while the solo operator's double-deviation constraint




Figure 1: Expected compensation as q varies. The illustration sets $b = 1$, $\lambda_0 = 0.2$, and $\lambda_1 = 0.6$, so $q^* = 5/7$ and $\bar{q} = 5/6$. The team is cheaper to the left of q^* , while the solo operator is cheaper to its right.

is binding. The team becomes attractive precisely because separation eliminates the information rent generated by that deviation, rather than because of the double-shirking constraint.

Second, the cutoff depends on $\theta = \lambda_1 - \lambda_0$, but not separately on λ_0 , b , or the scale of compensation. A larger θ lowers q^* , shrinking the range of q for which separation is cheaper: when research has a larger effect on quality, the principal is willing to tolerate a less informative quality-verification technology before abandoning the solo operator.

The static benchmark thus isolates the central organizational tension. Bundling economizes on moral-hazard rents but creates an information rent; separation eliminates that information rent but creates free-riding. Whether two heads are better than one depends on which distortion is more costly.

Private information is essential to this comparison. If quality were publicly observable, the solo operator could be induced at cost $\max\{2b, \lambda_1 b/\theta\}$, while the team's cost would remain $b + \lambda_1 b/\theta$, so bundling would always dominate. What the static model cannot explain, however, is why the preferred organization should change over the life of a project. In the dynamic model, the cost of controlling the solo operator's double deviation becomes history dependent, while the team's free-riding cost remains stationary. This history dependence generates endogenous transitions between organizational forms.

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. All players are risk neutral and do not discount future payoffs. 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 a flow 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.³ 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

³We discuss our motivation for allowing two forms of shirking in Section 6.

the payoff $R > 0$. If a lemon is built and accepted by the principal, a flaw materializes with probability q and pays $-L < 0$. With probability $(1 - q)$, no flaw develops and the lemon also yields R . Agency frictions arise because efforts and the quality of the discovery are both private to the agent. 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.

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 $\theta > 1/2$ and $R > 2c$, then the principal finances research and development until a peach is discovered. Her first-best expected payoff is given by*

$$\pi = R - 2c. \tag{16}$$

Proof. With full information, the principal's options are: fund both tasks (peaches arrive at rate 1, yielding $R - 2c$ in expectation), fund development only and discard lemons (peaches arrive at rate $1 - \theta$, yielding $R - c/(1 - \theta)$), or terminate. Accepting lemons is never optimal since $(1 - q)R - qL < 0$. Funding both tasks dominates development alone if and only if $2c \leq c/(1 - \theta)$, i.e., $\theta \geq 1/2$. ■

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 the agent must receive beyond those required to incentivize development alone. Because the question of whether one or two agents should be employed is only relevant when both tasks are desired, throughout the rest of the 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 2 *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 sufficient condition for research after one agent has been terminated and only one agent remains. 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 the arrival of discoveries, reported quality and realized payoffs, and has access to a public randomization device; her 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 non-decreasing 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. Each task is performed by at most one agent at any date. 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) .

An agent is employed if and only if he is assigned at least one task. 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 $\tau^i = \inf\{t > t_i : \chi_{it} = (0, 0)\}$. Termination is permanent, so each agent can be hired at most once, i.e., $\chi_{it} = (0, 0)$ for all $t \geq \tau^i$.

Let $\tau = \max\{\tau^1, \tau^2\}$ denote the project termination date. Given a contract, agent i 's continuation utility at time t is

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

The principal funds a task at flow cost c whenever it is assigned. Let n_t denote the number of tasks funded at date t . The principal's continuation value is

$$F_t = \mathbb{E} \left[Y_\tau - c \int_t^\tau n_s ds - \int_t^\tau (dC_s^1 + dC_s^2) \middle| \hat{\mathcal{F}}_t \right], \quad (17)$$

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. Given a contract, each agent chooses an action process adapted to his own filtration and a reporting strategy. A contract and a profile of agent strategies are incentive compatible if each agent's strategy maximizes his expected utility given the contract and the other agent's strategy. A contract is optimal if, together with

an incentive-compatible profile, it maximizes the principal's date-0 payoff over all incentive-compatible pairs that deliver each agent non-negative initial utility: $W_{t_i}^i \geq 0$ for $i \in \{1, 2\}$.

4 The Agents' 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 rate (zero) so deferring payments to when the contract ends is costless. Furthermore, because the agent is risk-neutral, any payment that the agent would receive when the contract ends without a product can 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 one.

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

$$dW_t^i = \beta_t^i(dN_t - dt), \quad (18)$$

The evolution of the agent's continuation utility given in (18) 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 incentivize 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 the agent exerts effort on both tasks conditional on truthful reporting. The second set of constraints ensures that the 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

$$\gamma_t^S + W_t^S \geq (1 - q)(\beta_t^S + W_t^S) + q \cdot 0 \quad (19)$$

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:

$$\beta_t^S = \beta_S(W_t^S) = \max \left\{ 2b, \frac{b}{\theta q} - W_t^S \right\}. \quad (20)$$

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^S 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.⁴ Whether this temptation exceeds the benefit of double shirking depends on model parameters. Define

$$K \equiv \frac{b}{\theta q} \quad \text{and} \quad \underline{W} \equiv K - 2b = \frac{b(1 - 2\theta q)}{\theta q}, \quad (21)$$

so that $\beta_S(W) = K - W$ for $W < \underline{W}$ and $\beta_S(W) = 2b$ for $W \geq \underline{W}$. If $\theta q \geq 1/2$, then $\underline{W} \leq 0$ and $\beta_S(W) = 2b$ 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 ($2b$).

⁴Producing lemons requires the agent’s effort in development, so the agent cannot reap unlimited utility by continuously fabricating lemons.

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 principal solicits a report of its quality. As in the static benchmark, the principal achieves the minimum cost if the report comes from the developer. If the developer reports a peach, the principal accepts the discovery and realizes R or $-L$. If he 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:

$$\beta_t^D \geq b. \quad (22)$$

To ensure the developer truthfully reports peaches and lemons requires

$$\beta_t^D \geq \gamma_t^D, \quad (23)$$

$$\gamma_t^D \geq (1 - q)\beta_t^D - qW_t^D. \quad (24)$$

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, as in the static benchmark.

The IC condition for the researcher is

$$\beta_t^R \geq \max\{\underbrace{(1 - \theta)\beta_t^R + b}_{a\text{-shirking}}, \underbrace{(1 - \theta)\beta_t^R + b + \theta\gamma_t^R}_{p\text{-shirking}}\} \quad (25)$$

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 $\beta_t^R = \beta_R = b/\theta > b$. The last inequality is due to the same free-riding problem in the static benchmark, which reflects the additional cost of the two-agent team compared to the minimum possible cost of incentivizing a solo operator.

4.3 Cost Comparison

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

Proposition 3 *The principal’s flow cost of incentivizing two separate agents is lower than the flow cost of a solo innovator if both q and the continuation utility of the solo innovator are sufficiently low, i.e.,*

$$\beta_D + \beta_R < \beta_S(W_t^S) \text{ if } q < \frac{1}{1+\theta} \text{ and } W_t^S < K - \frac{b(1+\theta)}{\theta} \quad (26)$$

Notice that the condition $q < \frac{1}{1+\theta}$ is identical to the condition for two agents to be cheaper in Proposition 1 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. The two-agent team creates a free-riding problem but saves the cost of inducing truthful reporting, where the latter is higher when q is lower.

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 the double deviation, 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 3

$$\pi > \frac{(1 + \theta)b}{\theta} = \beta_D + \beta_R \quad (27)$$

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 structure with a two-agent team. We do not make the analogous assumption for a structure with a solo innovator 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. For each $\vec{W} = (W_1, W_2) \in \mathbb{R}_{++}^2$, the principal chooses between two assignments: agent 1 researches and agent 2 develops, or the reverse.

Because the principal can randomize over continuation values, V is concave, and it satisfies the HJB equation as an inequality,

$$V(W_1, W_2) \geq \pi - (W_1 + W_2) - \min \left\{ \begin{array}{l} \beta_D(1 + V_1) + \beta_R(1 + V_2), \\ \beta_R(1 + V_1) + \beta_D(1 + V_2) \end{array} \right\}, \quad (28)$$

where $V_i = \frac{\partial V}{\partial W_i}$, with equality at every state at which the principal does not randomize. Along the optimal path characterized below the principal does not randomize. The boundary conditions 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 agents are interchangeable, relabeling the agents in any incentive-compatible contract yields an incentive-compatible contract with the same cost. Thus, the problem is symmetric, and 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). In this region, it is better to assign the agent with the higher continuation value the task that requires stronger incentives. Thus, we have:

Proposition 4 *For $W_2 \geq W_1 > 0$, it is optimal to assign agent 2 to research and assign agent 1 to development. The HJB for the principal's problem can be simplified to*

$$V(\vec{W}) = \pi - (W_1 + W_2) - \beta_R(1 + V_2(\vec{W})) - \beta_D(1 + V_1(\vec{W})) \quad (29)$$

Intuitively, because V is symmetric and is jointly concave in its arguments, it is never strictly optimal to assign the agent with less continuation value the more expensive task. Compared to development, research is the costlier task requiring a higher flow cost of incentives. Thus, whenever $W_2 \geq W_1$ it is without loss to assign research to Agent 2 and development to Agent 1 until either $W_2 = 0$ or $W_2 < W_1$.⁵

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.

If the principal hires the second agent when the state reaches $W_t = w > 0$, she enters the two-agent regime characterized by Proposition 4, with each agent immediately assigned one task. Hiring at $w = 0$ instead replaces the terminated incumbent with a new solo innovator. Conditional on hiring at that point, she chooses the initial promise to the new agent to maximize her continuation payoff:

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

Define $W_2^*(w) \in \arg \max_{y \geq 0} V(w, y)$. The boundary choice $y = 0$ represents continuing with the incumbent alone, with value $V(w, 0) = F_1(w)$; agent 2 is not hired in this case. 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

⁵The optimality is only weak. Wherever both promises can be exhausted at the same date, V depends on them only through their sum, and any assignment path that preserves joint termination is optimal.

agent 1 at hiring, 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 the incumbent's continuation values at the time of hire.⁶ 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:

$$\Pi = \sup_{(\tau, W_0)} \int_0^\tau e^{-t} (\pi - W_t - \beta_S(W_t)) dt + e^{-\tau} \bar{G}(W_\tau), \quad dW_t = -\beta_S(W_t) dt. \quad (30)$$

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. Hiring must occur no later than agent 1's termination date $\int_0^{W_0} du/\beta_S(u)$. Let (τ^*, W_0^*) denote the solution to the problem and let $\hat{W}_1 = W_{\tau^*}$. We refer to $\hat{W}_1$ as the *hiring state*: agent 1's continuation utility when agent 2 is hired.

For a fixed W_0 , differentiating the objective in (30) with respect to τ gives the marginal value of waiting an additional instant before hiring the second agent:

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

The function $\Phi(W)$ therefore captures the *benefit of delay*: if $\Phi(W) > 0$, a brief delay increases the payoff, while if $\Phi(W) < 0$, a brief delay reduces it. At a differentiable interior solution, $\hat{W}_1$ must satisfy the first-order condition

$$\Phi(\hat{W}_1) = 0, \quad (31)$$

which corresponds to a smooth-pasting condition between the single-agent and two-agent value functions. Observe that for $K > 2b$ and $W \in [0, \underline{W}]$,

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

Thus, Φ satisfies the (weak) single-crossing property on this branch. For $W > \max\{0, \underline{W}\}$, the derivative is $-(1 + \bar{G}'(W)) - 2b\bar{G}''(W)$, whose sign need not be positive. At kinks, the corresponding one-sided conditions apply.

⁶For 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 $\bar{G}(W_1)$.

Connection to the QVI formulation. The stopping problem in (30) is equivalent to a one-obstacle quasi-variational inequality (QVI) for the principal's value function $F(W)$ when she has the option to hire a second agent. Let the solo-innovator HJB operator $\mathcal{A}$ 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

$$\max \{ \mathcal{A}F(W), \bar{G}(W) - F(W) \} = 0, \quad F(0) = \bar{G}(0). \quad (32)$$

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

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

are value matching, smooth pasting, and the continuation HJB at the boundary, respectively. Substituting value matching and smooth pasting into the continuation HJB yields 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 (32) 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.

Let $\mathcal{W}_1 = \mathbb{R}_+$ denote the possible continuation values for agent 1. 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.

$$\begin{aligned} \mathcal{C}(K) &\equiv \{W \in \mathcal{W}_1 : F(W; K) > \bar{G}(W; K)\}, \\ \mathcal{S}(K) &\equiv \{W \in \mathcal{W}_1 : F(W; K) = \bar{G}(W; K)\}. \end{aligned}$$

In words, $\mathcal{C}(K)$ is the region where it is optimal for the principal to delay hiring agent 2, whereas $\mathcal{S}(K)$ is the region where immediately hiring agent 2 is optimal. On a range where Φ is continuous and weakly increasing, the hiring threshold is the largest state at which $\Phi \leq 0$, with hiring delayed until $W = 0$ if Φ is strictly positive. The following lemma characterizes the hiring state after optimizing the initial promise.

Lemma 5 *There exist 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 $K \in (K_{\text{seq}}, K_{\text{sim}})$, $\hat{W}_1(K) \in (0, W_0^*)$.
3. For all $K \geq K_{\text{sim}}$, $\hat{W}_1(K) = W_0^*$.

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 deter the double deviation. 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 deterring the double deviation 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, the double deviation is very costly to deter, making early hiring of the second agent optimal despite the fixed cost.

5.3 Organizational Regimes

Lemma 5 characterizes the optimal policy for when to hire the second agent $\tau = \inf\{t : W_t \leq \hat{W}_1(K)\}$. We now relate the hiring state to 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 optimal contract is not unique: any assignment path along which the two promises reach zero at the same date delivers the same payoff. However, if there were an arbitrarily small cost of reallocating tasks across agents, as when human capital is task specific (Gibbons and Waldman, 2004) or coordinating work across agents is costly (Becker and Murphy, 1992), then fixed specialist roles would be uniquely optimal. 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 $\bar{G}$. When $K > 2b$ and $\hat{W}_1 \leq \underline{W}$, 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

$$A(\hat{W}_1) = \frac{\pi - K - \bar{G}(\hat{W}_1)}{K - \hat{W}_1} \quad (33)$$

For intuition, we can note that $A(\hat{W}_1)$ corresponds to $d\Pi/dW_0$ evaluated at $W_0 = \hat{W}_1$. Thus, a positive $A(\hat{W}_1)$ implies $W_0^* > \hat{W}_1$, while a negative slope makes increasing the initial promise above the hiring state unprofitable on this branch. Note also that at an interior threshold, the first-order condition $\Phi(\hat{W}_1) = 0$ rearranges to exactly $A(\hat{W}_1) = \bar{G}'(\hat{W}_1)$: the slope of the attached single-agent value function equals the slope of the obstacle at the contact point (smooth pasting).

Lemma 6 *Agent 1 is hired strictly before agent 2 if and only if $K < K_{\text{sim}}$.*

From Lemmas 5 and 6, we can conclude that the optimal organizational design can take one of three possible forms as summarized in the following proposition.

Proposition 5 (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. **External Consultant:** 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 and immediately takes over research, while agent 1 continues development. Both agents are terminated at the same time if no success occurs.

In the low- K , 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 high- K , 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.⁷

⁷Absent reallocation costs, the specialist arrangement is one of a continuum of optimal contracts. For example, the principal could give the agents equal initial utilities and exchange their roles at some point before termination. However, any other contract would incur an arbitrarily small reallocation costs. Therefore, under our selection criterion, the optimal contract is to grant the researcher higher initial utility and have him specialize in research for the duration of the project.

In the medium- K region, 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 his continuation value reaches $\hat{W}_1$, then 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. The research task is more costly to incentivize, so $W_2^*(\hat{W}_1) = \hat{W}_1/\theta$ and the agents’ continuation values reach zero at the same time.

Why is the new agent assigned to the more expensive task? The reason is that promises to agent 2 are deferred, and are therefore paid only if the project does not succeed before agent 2 is hired, whereas promises to agent 1 raise the cost of compensation for earlier successes as well. To see this, fix the lengths of the solo and team phases, so that the probability of success is fixed, and consider reallocating the marginal unit of promised utility the moment agent 2 is hired. Promised to agent 2, it is paid only if the project reaches that date. Promised to agent 1, it is paid in that event as well, but also if he succeeds before his continuation value reaches $\underline{W}$, since until then his reward for a discovery is $2b$ on top of his continuation value.⁸ Because $W_0 > \underline{W}$, agent 1’s promise is strictly the more expensive one, and the principal optimally gives him the developer’s share, the smallest share consistent with joint termination.

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 arrive at rate $\epsilon \in (0, 1 - \theta)$, both when the agents exert effort and when the developer exerts effort and the researcher actively shirks. The arrival rates of peaches and lemons in this extension are:

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.

⁸Below $\underline{W}$, agent 1 collects K upon a discovery regardless of his continuation value, so if success occurs over that stretch the unit costs no more than it would if promised to agent 2.

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.

Proposition 6 *Suppose lemons arise on path as in Table 2 and fix $\epsilon \in (0, 1 - \theta)$. 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 the double-shirking constraint 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 Proposition 3) 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 1 *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.,

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

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 his continuation value is sufficiently high.⁹ 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

⁹This 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.

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 2 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 truthful reporting.

Lemma 7 *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, assigning both tasks to the sole remaining agent dominates allowing development-only intervals for all continuation values. 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

In this paper we study 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 the double deviation in which the agent both reduces effort and misreports quality. Separating tasks across two agents eliminates this double deviation 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, in which the second hire serves as an external consultant, beginning research immediately while the incumbent continues development. 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 core insights developed here and to engage with additional institutional features of innovation-focused organizations.

References

- Admati, A. R. and M. Perry (1991). Joint projects without commitment. *Review of Economic Studies* 58(2), 259–276.
- Becker, G. S. and K. M. Murphy (1992). The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics* 107(4), 1137–1160.
- 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.
- Feng, F. Z., C. R. Taylor, M. M. Westerfield, and F. Zhang (2024). Setbacks, shutdowns, and overruns. *Econometrica* 92(3), 815–847.
- Fulghieri, P., Y. Hu, and F. Varas (2025). Multilateral contracting in stage financing. Working paper, Available at SSRN 5753183.
- Georgiadis, G. (2015). Projects and team dynamics. *Review of Economic Studies* 82(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.
- Gibbons, R. and M. Waldman (2004). Task-specific human capital. *American Economic Review* 94(2), 203–207.
- 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.
- 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.
- 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.

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. Recall from Section 4 that $K \equiv \frac{b}{\theta q}$.

Region $0 \leq W \leq \underline{W}$ when $K > 2b$. 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 $W > \max\{0, \underline{W}\}$. 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): $2b F'_1(W) + F_1(W) + W = \pi - 2b$. Using $x_1 = 2b$ and $x_0 = \pi - 2b$ (so that $x_0 + x_1 = \pi$) in (A.2):

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

When $K > 2b$, continuity at $\underline{W}$ determines C_1 . When $K \leq 2b$, $\beta_S = 2b$ for all $W \geq 0$, and the boundary condition is $F_1(0) = 0$. Thus,

$$C_1 = \begin{cases} \exp(\rho_1 \underline{W})(1 + A)(\underline{W} - K), & K > 2b, \\ -\pi, & K \leq 2b. \end{cases} \quad (\text{A.6})$$

Optimal starting utility. Under Assumption 3, $\pi > 2b$. If $K < \pi$, the optimal initial utility is interior and satisfies $F_1'(W_1^*) = 0$, yielding

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

If $K > \pi$, the maximum is attained at $W_1^* = 0$, with value zero. If $K = \pi$, every $W \in [0, \underline{W}]$ maximizes F_1 , again with value zero.

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 2. (The derivation follows steps analogous to those above and is omitted.)*

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 $x_0 = \pi - x_1$ (so that $x_0 + x_1 = \pi$). With boundary condition $F_2(0) = 0$, substituting into (A.2) gives $C = -(x_0 + x_1) = -\pi$:

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

Under Assumption 3, $F_2'(0) = \pi\rho_2 - 1 > 0$, so the optimal starting utility is interior:

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

B Proofs

B.1 Proofs in Section 2

Proof of Lemma 1 Let $x \equiv \beta - \gamma$. Every feasible contract satisfies

$$x \geq \frac{b}{\theta}. \quad (\text{B.1})$$

Because $\beta = \gamma + x$, the lower truthtelling constraint implies

$$\gamma \geq (1 - q)(\gamma + x) \iff \gamma \geq \frac{1 - q}{q}x. \quad (\text{B.2})$$

It follows that

$$U(\beta, \gamma) = \gamma + \lambda_1 x \geq \left(\frac{1 - q}{q} + \lambda_1 \right) x = \frac{D(q)}{q} x \geq \frac{bD(q)}{q\theta}. \quad (\text{B.3})$$

Together with (Double-Shirking IC), this establishes the lower bound

$$U(\beta, \gamma) \geq \max \left\{ 2b, \frac{bD(q)}{q\theta} \right\}. \quad (\text{B.4})$$

Now consider the contract in (6). It satisfies truthtelling by construction, and

$$\theta(\beta^* - \gamma^*) = q\theta\beta^* \geq b, \quad (\text{B.5})$$

so the research-shirking constraint is satisfied. Its expected compensation is

$$U(\beta^*, \gamma^*) = D(q)\beta^* = \max \left\{ 2b, \frac{bD(q)}{q\theta} \right\}. \quad (\text{B.6})$$

Thus, the contract attains the lower bound and is optimal.

Proof of Proposition 1 First suppose that $\lambda_1 \geq 2\lambda_0$. Then the deviation cutoff

$$\bar{q} = \frac{1}{1 + \lambda_1 - 2\lambda_0} \quad (\text{B.7})$$

lies in $(0, 1]$. Let $D(q) \equiv \lambda_1 + (1 - \lambda_1)(1 - q)$. From the analysis of the solo operator, expected compensation is

$$C^{solo}(q) = D(q)\beta^* = \begin{cases} \frac{bD(q)}{q\theta}, & q < \bar{q}, \\ 2b, & q \geq \bar{q}. \end{cases} \quad (\text{B.8})$$

For $q < \bar{q}$, the double deviation binds. For $q \geq \bar{q}$, the double-shirking constraint binds.

With two agents, the developer's expected compensation is b and the researcher's ex-

pected compensation is $\lambda_1 b / \theta$. Thus,

$$C^{team} = b + \frac{\lambda_1 b}{\theta} = \left(\frac{2\lambda_1 - \lambda_0}{\lambda_1 - \lambda_0} \right) b. \quad (\text{B.9})$$

The team cost is independent of q because the researcher need not report quality, while the developer's effort is already sunk when he observes it.

Define the organizational cutoff

$$q^* \equiv \frac{1}{1 + \theta}. \quad (\text{B.10})$$

In the region where the double deviation binds,

$$C^{solo}(q) - C^{team} = \frac{b}{\theta} \left(\frac{1}{q} - (1 + \theta) \right). \quad (\text{B.11})$$

Moreover, $q^* < \bar{q}$ so the organizational cutoff lies in the region where the double deviation binds. For $q \geq \bar{q}$,

$$C^{team} - C^{solo} = \frac{\lambda_0}{\theta} b > 0. \quad (\text{B.12})$$

If $\lambda_1 < 2\lambda_0$, then $\bar{q} > 1$, so research shirking determines $C^{solo}(q)$ for every feasible $q \in (0, 1]$ and the same q^* applies.

B.2 Proofs for Section 5

Lemma B.1 *Fix $K > 2b$ so that $\beta_S(W; K) = K - W$ on $W \in [0, \underline{W}]$. Wherever $\bar{G}$ is twice differentiable on this interval, $\Phi(W; K)$ is weakly increasing in W , and strictly increasing whenever $\bar{G}_{WW} < 0$.*

Proof of Lemma B.1. 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 $G(w) = \max_{y \geq 0} V(w, y)$. Differentiate w.r.t. W wherever $\bar{G}$ is C^2 :

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

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. Where $\beta_S(W) = 2b$, direct differentiation instead gives $\Phi_W = -(1 + \bar{G}_W) - 2b\bar{G}_{WW}$. Thus the single-crossing argument applies only for $W \leq \underline{W}$. For this reason, the proof of Lemma 5 below instead compares the principal's total payoffs, with the initial promise also chosen optimally. ■

Lemma B.2 *Let $F_2^{\max} \equiv F_2(W_{sim}^*) > 0$ and*

$$K_{sim} = \pi - F_2^{\max} = (\beta_D + \beta_R) \left(1 + \ln \frac{\pi}{\beta_D + \beta_R} \right). \quad (\text{B.13})$$

Then:

1. $\max_{W \in \mathcal{W}_1} \bar{G}(W; K) < \pi - K$ for all $K < K_{sim}$;
2. $\max_{W \in \mathcal{W}_1} \bar{G}(W; K) = F_2^{\max}$ for all $K \geq K_{sim}$, attained by a team whose promises reach zero simultaneously.

Proof. Since the concave envelope preserves the maximum, $\max_{\mathcal{W}_1} \bar{G} = \max_{\mathcal{W}_1} G$. While both agents are employed, each is assigned one task, and the sum of their promises decreases at rate $\beta_D + \beta_R$. If the first agent is terminated after t units of time and the other agent has remaining continuation utility z , the initial total promise is $(\beta_D + \beta_R)t + z$. The principal's payoff is then

$$\pi - (\beta_D + \beta_R)t - z - e^{-t} [\pi - z - F_1(z)]. \quad (\text{B.14})$$

This expression also holds if tasks are reassigned before either agent is terminated, because the sum of the agents' rewards, $\beta_D + \beta_R$, is unchanged.

The one-agent solution in Appendix A gives

$$\pi - z - F_1(z) = \begin{cases} \pi(1 - z/K), & K > 2b, \quad 0 \leq z \leq \underline{W}, \\ \frac{2b\pi}{\max\{2b, K\}} \exp\left(\frac{\max\{0, \underline{W}\} - z}{2b}\right), & z \geq \max\{0, \underline{W}\}. \end{cases} \quad (\text{B.15})$$

If $\max\{2b, K\} \geq \pi$, the value of z that maximizes (B.14) is zero. Otherwise, for $0 \leq t < \ln[\pi/\max\{2b, K\}]$, a maximizing z is

$$z = \max\{0, \underline{W}\} + 2b \left(\ln \frac{\pi}{\max\{2b, K\}} - t \right).$$

The resulting payoff is $F_1(W_1^*) - (\beta_D + \beta_R - 2b)t$, which is decreasing in t because $\theta < 1$. For $t > \ln[\pi/\max\{2b, K\}]$, the maximizing z is zero and the payoff is $F_2((\beta_D + \beta_R)t) \leq F_2^{\max}$. At $t = \ln[\pi/\max\{2b, K\}]$, the two expressions coincide. Thus

$$\max_{W \in \mathcal{W}_1} \bar{G}(W; K) = \max \{F_1(W_1^*), F_2^{\max}\}. \quad (\text{B.16})$$

The bound also holds when the principal randomizes, since it holds for every draw.

For part 1, take $K < K_{sim}$: the principal's payoff from a team is at most $F_2^{\max} = \pi - K_{sim} < \pi - K$, and the single-agent value function in Appendix A gives $F_1(W_1^*) < \pi - K$. Hence $\max \bar{G} < \pi - K$.

For part 2, take $K \geq K_{sim}$. The same single-agent value function gives $F_1(W_1^*) \leq F_2^{\max}$, so $\max \bar{G} = F_2^{\max}$. Assigning development to agent 1 and research to agent 2, with initial

promises $\theta W_{sim}^*/(1 + \theta)$ and $W_{sim}^*/(1 + \theta)$, attains this value. Both promises reach zero at the same time. ■

Proof of Lemma 5. Part (3) follows from Lemma 6. To prove parts (1) and (2), we first show that, when $K < K_{sim}$, the principal need only compare the sequential-generalist contract with hiring an external consultant. We then show that the comparison has a unique cutoff in K .

Consider an initial solo phase of length $\tau > 0$, followed by a team phase of length t . If agent 1 is terminated first and agent 2 continues alone with continuation utility z , the continuation utilities when agent 2 was hired satisfy $W_1 \in [bt, bt/\theta]$ and $W_2 = (\beta_D + \beta_R)t - W_1 + z$. If agent 2 is terminated first, then $W_1 - z \in [bt, bt/\theta]$ and $W_2 = (\beta_D + \beta_R)t - W_1 + z$. The principal's initial payoff is

$$\pi - W_0 - e^{-\tau} (W_2 + e^{-t} [\pi - z - F_1(z)]), \quad (\text{B.17})$$

where W_0 is agent 1's initial continuation utility, which falls to W_1 after τ units of time without a discovery. We call $W_0 + e^{-\tau} (W_2 + e^{-t} [\pi - z - F_1(z)])$, the amount by which this payoff falls short of π , the principal's *cost*. If the optimal τ is positive,

$$\beta_S(W_0) = e^{-\tau} (W_2 + e^{-t} [\pi - z - F_1(z)]).$$

Moreover, Lemma B.2 implies $W_2 + e^{-t} [\pi - z - F_1(z)] > K - W_1$.

Holding t and z fixed, increasing agent 1's share of the total promise raises the cost. Its derivative, after optimizing τ , is

$$e^{-\tau} \left(\frac{W_2 + e^{-t} [\pi - z - F_1(z)]}{\beta_S(W_1)} - 1 \right) > 0.$$

When $W_1 < \underline{W}$, the inequality follows from the preceding bound; in the region where $\beta_S = 2b$, it follows from $\tau > 0$ and the first-order condition. Thus agent 1 performs development, so $W_1 = bt$ if agent 1 is terminated first and $W_1 = bt + z$ if agent 2 is terminated first.

If agent 2 is terminated first, the derivative of the cost with respect to z is

$$e^{-\tau} \left(\frac{W_2 + e^{-t} [\pi - z - F_1(z)]}{\beta_S(W_1)} - \frac{e^{-t} [\pi - z - F_1(z)]}{\beta_S(z)} \right) > 0 \quad (t > 0),$$

since $W_1 \geq z$, $\beta_S(W_1) \leq \beta_S(z)$, and $W_2 > 0$. Hence it is never optimal for agent 2 to be terminated first and agent 1 to continue alone.

If agent 1 is terminated first, optimize z as in the proof of Lemma B.2. When the optimal z is positive,

$$W_2 + e^{-t} [\pi - z - F_1(z)] = \max\{2b, K\} + 2b \ln \frac{\pi}{\max\{2b, K\}} + b \left(\frac{1}{\theta} - 2 \right) t.$$

The derivative of the cost with respect to t , after choosing τ optimally, is

$$be^{-\tau} \left(\frac{W_2 + e^{-t}[\pi - z - F_1(z)]}{\beta_S(bt)} + \frac{1}{\theta} - 2 \right) > 0.$$

Indeed, the fraction exceeds one: when $bt < \underline{W}$, its numerator exceeds $K - bt$ by $2b \ln(\pi/K) + b(1/\theta - 1)t > 0$; in the region where $\beta_S = 2b$, this follows from the optimal τ being positive. Consequently, a team phase after which agent 2 continues alone is dominated either by the sequential-generalist contract ($t = 0$) or by terminating both agents at the same time ($z = 0$).

A potentially optimal external-consultant contract therefore has continuation utilities $(bt, bt/\theta)$ at the hiring date. Hiring agent 2 in the region where $\beta_S = 2b$, i.e., $bt \geq \max\{0, \underline{W}\}$, cannot improve on the optimal sequential-generalist contract. To see this, choose $\tau > 0$ optimally. At an interior optimum in t in this region, the first-order condition is $\pi e^{-t} = b(t+2)/\theta$. Compare with the sequential-generalist contract at $K = b(t+2)$, which is at least the original K . The consultant's total cost is

$$b(t+2) + 2b \ln \frac{t+1}{\theta},$$

whereas the cost of the sequential-generalist contract at $K = b(t+2)$ is

$$b(t+2) + 2b \ln \left(\frac{3t+2+2 \ln(1/\theta)}{t+2} \right).$$

The latter is strictly smaller, since $(t+1)(t+2)/\theta > 3t+2+2 \ln(1/\theta)$ for $\theta < 1$ and $t \geq 0$. The principal's payoff under the sequential-generalist contract is weakly decreasing in K , so it also dominates at the original K . Corner solutions in t either lie in the region $W_1 \leq \underline{W}$, reduce to immediate hiring, or are covered by the preceding comparison for $z > 0$. Thus hiring an external consultant can be optimal only if $K > 2b$, in which case agent 2 is hired at some $0 < W_1 < \underline{W}$.

For a hiring state $W_1 < \underline{W}$ with payoff g from hiring agent 2, choosing the length of the solo phase optimally gives

$$\begin{aligned} W_0 &= K - 2b + 2b \ln \psi(W_1, g), \\ \Pi &= \pi - K - 2b \ln \psi(W_1, g), \end{aligned} \tag{B.18}$$

where

$$\psi(W_1, g) \equiv \frac{\pi - W_1 - g}{K - W_1}.$$

The numerator is agent 2's continuation utility plus the expected surplus lost if the project ends without a discovery. The denominator is agent 1's reward for a discovery at the hiring state, $\beta_S(W_1) = K - W_1$. Because Π is decreasing in ψ , the principal prefers whichever hiring arrangement yields the smaller ψ , and the rest of the proof compares values of ψ . By Lemma B.2, $\psi > 1$. The sequential-generalist contract has $W_1 = 0$ and $g = F_1(W_1^*)$, so

$\psi = 1 + (2b/K) \ln(\pi/K)$. For the external-consultant contract,

$$\psi = \frac{W_1/\theta + \pi e^{-W_1/b}}{K - W_1}. \quad (\text{B.19})$$

Thus the principal compares the sequential-generalist contract with the minimum of (B.19).

The same comparison applies to public lotteries. For any team phase with $W_1 < K$, increasing agent 1's share of a fixed total promise raises ψ , since its derivative is

$$\frac{W_2 + e^{-t}[\pi - z - F_1(z)] - (K - W_1)}{(K - W_1)^2} > 0.$$

Leaving agent 1 with positive continuation utility after agent 2 is terminated also raises ψ , since $\beta_S(z) \geq K - z > K - W_1$ and $W_2 > 0$. If agent 1 is terminated first and the optimal z is positive, then ψ , evaluated at the optimal z , increases with $W_1 = bt$:

$$\frac{d\psi}{dW_1} = \frac{(1/\theta - 1)K + 2b \ln(\pi/K)}{(K - W_1)^2} > 0.$$

Hence the smallest ψ is attained by the sequential-generalist contract or by an external-consultant contract in which both agents are terminated at the same time. The comparison below shows that the best external-consultant contract never hires agent 2 at $W_1 \geq \underline{W}$.

Suppose that, when agent 2 is hired, the principal uses a lottery over agent 1's continuation utility: with probability μ_j , agent 1's continuation utility is $w_j < K$ and the principal's payoff from hiring is g_j . Then ψ for the lottery is a weighted average of the values $\psi(w_j, g_j)$ it would take if each draw occurred for sure:

$$\psi\left(\sum_j \mu_j w_j, \sum_j \mu_j g_j\right) = \frac{\sum_j \mu_j (\pi - w_j - g_j)}{K - \sum_j \mu_j w_j} = \sum_j \frac{\mu_j (K - w_j)}{K - \sum_\ell \mu_\ell w_\ell} \psi(w_j, g_j).$$

Draws with $w_j \geq K$ cannot improve the comparison, because $g_j \leq \max \bar{G} < \pi - K$. Equivalently, since $A(\hat{W}_1) = \psi(\hat{W}_1, \bar{G}(\hat{W}_1)) - 1$ by (33), every pair (w, g) satisfies $g \leq \pi - K + A(\hat{W}_1)(w - K)$: it lies on or below the line through the optimal pair with slope $A(\hat{W}_1) > 0$. In the region where $\beta_S = 2b$, it cannot be strictly optimal to hire agent 2 at a state where $\bar{G}$ is linear, which would require a lottery over agent 1's continuation utility, where $\Phi'(W) = -(1 + \bar{G}'(W)) \leq 0$. Moving the hiring state to an endpoint of the linear segment weakly improves the payoff. If, instead, $\bar{G} = G$ at the hiring state, no lottery is needed. Agent 2 is hired at a single, deterministic value of W_1 , which is the case analyzed above. Hiring at $\underline{W}$ is covered by (B.18). If the principal randomizes during the team phase, the expected continuation utilities and payoffs are averages of those in (B.14), so the bound applies to these lotteries as well. Hence the comparison also holds with G replaced by its concave envelope $\bar{G}$.

Hiring an external consultant improves on the sequential-generalist contract precisely

when

$$K + 2b \ln \frac{\pi}{K} - \pi e^{-W_1/b} - \left(1 + \frac{1}{\theta} + \frac{2b}{K} \ln \frac{\pi}{K}\right) W_1 > 0. \quad (\text{B.20})$$

For $\beta_D + \beta_R \leq K \leq K_{\text{sim}}$, the left-hand side is strictly concave in W_1 , with an unconstrained maximizer

$$W_1 = b \ln \frac{\pi}{\beta_D + \beta_R + (2b^2/K) \ln(\pi/K)} > 0.$$

The left-hand side of (B.20) is negative for every $W_1 \geq \underline{W}$. Hence, whenever its maximum is nonnegative, the maximizer above satisfies $W_1 < \underline{W}$. To see this for $\underline{W} \leq W_1 < K$, maximize the left-hand side of (B.20) over $\pi > 0$. Using $\beta_D + \beta_R > 2b$, the maximum is bounded above by

$$\frac{2b(K - W_1)}{K} \left[-\left(\frac{K}{2b}\right)^2 + \frac{2K}{b} - 3 - 2 \ln \frac{K}{2b} \right] < 0.$$

The bracket is zero at $K = 2b$ and strictly decreasing in K , so it is negative for every $K > 2b$. Hiring at $W_1 \geq K$ cannot make the left-hand side of (B.20) positive either.

K_{seq} therefore solves

$$K + 2b \ln \frac{\pi}{K} = \left(\beta_D + \beta_R + \frac{2b^2}{K} \ln \frac{\pi}{K} \right) \times \left[1 + \ln \frac{\pi}{\beta_D + \beta_R + (2b^2/K) \ln(\pi/K)} \right]. \quad (\text{B.21})$$

We now show that the root is unique and satisfies $\beta_D + \beta_R < K_{\text{seq}} < K_{\text{sim}}$. The difference between the left- and right-hand sides is strictly increasing in K on $[\beta_D + \beta_R, K_{\text{sim}}]$, since its derivative is

$$1 - \frac{2b}{K} + \frac{2b^2[1 + \ln(\pi/K)]}{K^2} \ln \frac{\pi}{\beta_D + \beta_R + (2b^2/K) \ln(\pi/K)} > 0.$$

At $K = \beta_D + \beta_R$, the difference is zero when $\pi = \beta_D + \beta_R$ and strictly decreasing in π thereafter, by direct differentiation. It is therefore negative under Assumption 3. At $K = K_{\text{sim}}$, evaluate (B.20) at $W_1 = \theta W_{\text{sim}}^*/(1 + \theta)$. Its left-hand side is

$$2b \ln \frac{\pi}{K_{\text{sim}}} \left(1 - \frac{\theta W_{\text{sim}}^*}{(1 + \theta) K_{\text{sim}}} \right) > 0.$$

Since the difference in (B.21) is continuous in K , the intermediate value theorem implies that there exists a unique $K_{\text{seq}} \in (\beta_D + \beta_R, K_{\text{sim}})$ at which it is zero. Moreover, $K_{\text{sim}} < \pi$ follows from (B.13).

For $2b < K < \beta_D + \beta_R$, the left-hand side of (B.20) is strictly increasing in K at each

fixed $W_1 < K$, since

$$1 - \frac{2b}{K} + \frac{2bW_1[1 + \ln(\pi/K)]}{K^2} > 0.$$

Thus no external consultant improves on the sequential-generalist contract there. For $K \leq 2b$, $\beta_S = 2b$ for all W , and the earlier comparison for the region where $\beta_S = 2b$ gives the same conclusion.

For (1), the sequential-generalist contract is therefore optimal for all $K \leq K_{\text{seq}}$ (at $K = K_{\text{seq}}$ the principal is indifferent, and we select the sequential-generalist contract), so $\hat{W}_1(K) = 0$. For (2), an external consultant strictly improves on the sequential-generalist contract when $K \in (K_{\text{seq}}, K_{\text{sim}})$, while immediate hiring is strictly suboptimal by Lemma 6. Hence $0 < \hat{W}_1(K) < W_0^*$ on this interval. This completes the proof. ■

Proof of Lemma 6. Define $M(K) \equiv \max_{W \in W_1} \bar{G}(W; K)$ and recall $K_{\text{sim}} = \pi - F_2^{\max}$ (Lemma B.2).

Case 1: $K \leq \beta_D + \beta_R$. Here $M(K) = F_1(W_1^*)$. A short solo phase, with agent 2 hired once agent 1's continuation utility reaches zero, improves on immediate hiring, since the marginal gain from delay at $\tau = 0$ is

$$\pi - \max\{2b, K\} - M(K) = 2b \ln \frac{\pi}{\max\{2b, K\}} > 0.$$

The strict inequality follows from Assumption 3, since $\max\{2b, K\} \leq \beta_D + \beta_R < \pi$. Hence agent 1 is hired strictly before agent 2.

Case 2: $\beta_D + \beta_R < K < K_{\text{sim}}$. By Lemma B.2 (part 1), $\pi - K > M(K)$. For any initial promise W_0 and hiring date τ , since $\beta_S(W) \geq K - W$, the flow payoff in (30) satisfies $\pi - W_t - \beta_S(W_t) \leq \pi - K$, with equality if and only if $W_t \leq \underline{W}$. Hiring immediately therefore yields at most $M(K)$ from any state.

If $M(K) = F_1(W_1^*)$, the construction in Case 1 again gives a strict improvement. Otherwise, choose the team's initial continuation utilities $W_1 = \theta W_{\text{sim}}^*/(1 + \theta)$ and $W_2 = W_{\text{sim}}^*/(1 + \theta)$. If $W_1 < \underline{W}$, starting at $W_0 = W_1 + \delta$ for small $\delta > 0$ and hiring agent 2 when W_t reaches W_1 , at date $\tau_\delta > 0$, yields

$$\Pi(\tau_\delta, W_1 + \delta) = (\pi - K)(1 - e^{-\tau_\delta}) + e^{-\tau_\delta} M(K) > M(K),$$

because W_t stays below $\underline{W}$, so the flow payoff is exactly $\pi - K$. If $W_1 \geq \underline{W}$, the marginal payoff from a short initial solo phase is instead

$$\pi - W_1 - 2b - F_2^{\max} = \beta_R - \beta_D + \frac{W_{\text{sim}}^*}{1 + \theta} > 0.$$

Hence $\tau^* > 0$: agent 1 is hired strictly before agent 2.

Case 3: $K \geq K_{\text{sim}}$. By Lemma B.2 (part 2), $M(K) = F_2^{\max}$, and $\pi - K \leq \pi - K_{\text{sim}} = F_2^{\max}$. The same flow bound gives, for every (τ, W_0) ,

$$\Pi(\tau, W_0) \leq (\pi - K)(1 - e^{-\tau}) + e^{-\tau} M(K) \leq M(K),$$

with the second inequality strict whenever $\tau > 0$ and $K > K_{\text{sim}}$. Hiring agent 2 immediately at a maximizer of $\bar{G}$, with total promised utility W_{sim}^* , attains $M(K)$. At $K = K_{\text{sim}}$ we select immediate hiring. Hence $W_0^* = \hat{W}_1$ and the two agents are hired simultaneously. ■

Proof of Proposition 5. From Lemmas 5 and 6, the optimal organizational design takes the three forms stated in the proposition. In the sequential-generalist regime each agent performs both tasks. In the simultaneous-specialist regime, Lemma B.2 gives the initial promises and common termination date. On $(K_{\text{seq}}, K_{\text{sim}})$, the proof of Lemma 5 shows that agent 2 begins research immediately upon hiring, while agent 1 continues development, and both promises reach zero simultaneously. An arbitrarily small reallocation cost selects fixed specialist roles.

For completeness, minimizing (B.19) gives the hiring threshold directly:

$$\frac{bK}{\theta} = \pi e^{-\hat{W}_1/b} (K - \hat{W}_1 - b). \quad (\text{B.22})$$

The left-hand side minus the right-hand side is strictly increasing in $\hat{W}_1 \in (0, K)$, with derivative $(\pi/b)e^{-\hat{W}_1/b}(K - \hat{W}_1) > 0$. It is negative at zero and positive at $K - b$, so the solution is unique. The preceding comparison places it in $(0, \underline{W})$. Agent 2's continuation utility and agent 1's optimal initial continuation utility are

$$W_2^*(\hat{W}_1) = \frac{\hat{W}_1}{\theta}, \quad W_0^* = K - 2b + 2b \ln(1 + A(\hat{W}_1)), \quad (\text{B.23})$$

where $A(\hat{W}_1)$ is given by (33).

We next establish that $\hat{W}_1(K)$ is strictly increasing in K on $(K_{\text{seq}}, K_{\text{sim}})$. Implicit differentiation of (B.22) yields

$$\hat{W}_1'(K) = \frac{\pi e^{-\hat{W}_1/b} - b/\theta}{(\pi/b)e^{-\hat{W}_1/b}(K - \hat{W}_1)} > 0.$$

At $K = K_{\text{sim}}$, the solution is $\hat{W}_1 = \theta W_{\text{sim}}^*/(1 + \theta)$ and $A(\hat{W}_1) = 0$, so immediate hiring is optimal under the selection in Lemma 5. At $K = K_{\text{seq}}$, the sequential-generalist and external-consultant contracts give the same payoff, so $\hat{W}_1(K)$ can jump from zero to a positive value at this cutoff. ■

B.3 Proofs for Section 6

Proof of Proposition 6. Fix $\epsilon \in (0, 1 - \theta)$. 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:

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

Note that double shirking and truthful peaches jointly imply $\beta \geq 2b$, so active shirking is implied whenever $\theta \geq 1/2$ (Assumption 2) and we may drop it in what follows. Moreover, once truthful lemons holds, any deviation that combines shirking with misreporting is weakly dominated by the corresponding shirk-and-report-truthfully deviation, so the constraints above are sufficient for full effort and truthful reporting. Define $K_\epsilon := \frac{b}{q(\theta - \epsilon)}$ and $(\bar{\beta}_\epsilon, \bar{\gamma}_\epsilon)$ as in the statement, and set

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

(The threshold $\widehat{W}_\epsilon$, at which the truthful-lemons constraint starts to bind, is distinct from the cost-comparison threshold $\bar{W}_\epsilon$ defined in Corollary 1.) 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 keep 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 ($W^S < \widehat{W}_\epsilon$). Suppose truthful lemons were slack at the optimum. Then the logic above pins down the contract from the double-shirking and passive-shirking constraints alone, yielding the constant pair $(\bar{\beta}_\epsilon, \bar{\gamma}_\epsilon)$ derived in the high- W^S case below; but that pair violates truthful lemons precisely when $W^S < \widehat{W}_\epsilon$ —a contradiction. Hence truthful lemons binds. 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 ($W^S \geq \widehat{W}_\epsilon$). 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 $W^S \geq \widehat{W}_\epsilon$, i.e., $\bar{\gamma}_\epsilon \geq (1-q)\bar{\beta}_\epsilon - qW^S$. Combining the two cases gives the expressions 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 for implementing research after one agent has been terminated and only one agent remains. We begin with the proof of Lemma 7, which characterizes the optimal contract that implements effort on development and truthful reporting.

Proof of Lemma 7. 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 $\delta = (1 - \theta)\beta + \theta\gamma$ and rewrite the objective as $-\delta F'(W) - (\delta - \theta\gamma) + \theta F(W + \gamma)$. Since $F'(W) \geq -1$, the objective is decreasing in δ (its derivative in δ , holding γ fixed, is $-(1 + F'(W)) \leq 0$), so the constraint $\delta \geq b$ binds. Holding $\delta = b$ fixed, the objective is increasing in γ (its derivative is $\theta(1 + F'(W + \gamma)) \geq 0$), so γ increases until $\beta \geq \gamma$ binds, giving $\beta = \gamma = \delta = b$. The remaining constraint $\gamma \geq (1 - q)\beta - qW$ is then satisfied. ■

Research with one remaining agent. We compare a solo innovator with a solo developer only after one agent has been terminated and no further hire is possible. While both agents are employed, each performs one task, so dropping research would mean terminating the researcher.

Let $v(w) = F_1(w)$ be the value obtained by assigning both tasks to the sole remaining agent. We compare this policy with allowing the same agent to switch between both tasks and development alone. Either assignment keeps him employed, so these switches are allowed.

Lemma B.3 (Research with one remaining agent) *Suppose F_1 is concave. Adding the option to assign only development to the sole remaining agent does not increase the principal's*

value if

$$(1 - \theta) \left[R - \min \left\{ \frac{1}{2}, \theta q \right\} (R - 2c) \right] \leq c.$$

In particular, Assumption 2 implies this sufficient condition.

Proof. Let $\beta_0 = \beta_S(0) = \max\{2b, b/(\theta q)\}$. The solo-innovator value satisfies

$$v(w) = R - 2c - w - \beta_S(w)(1 + v'(w)), \quad v(0) = 0. \quad (\text{B.24})$$

Evaluating at zero gives $1 + v'_+(0) = (R - 2c)/\beta_0$. Concavity and the one-agent value function then give

$$0 \leq 1 + v'(w) \leq \frac{R - 2c}{\beta_0}, \quad \beta_S(w) \leq \beta_0.$$

Under development alone, Lemma 7 gives $\beta = \gamma = b$. The corresponding HJB expression evaluated at the candidate v is

$$\mathcal{D}v(w) = (1 - \theta)(R - w - b) - c - bv'(w) + \theta v(w + b).$$

Concavity implies $v(w + b) \leq v(w) + bv'(w)$, so

$$\begin{aligned} \mathcal{D}v(w) - v(w) &\leq (1 - \theta) [R - w - b - v(w) - bv'(w)] - c \\ &= (1 - \theta) [2c + (\beta_S(w) - b)(1 + v'(w))] - c \\ &\leq (1 - \theta) \left[2c + (\beta_0 - b) \frac{R - 2c}{\beta_0} \right] - c \\ &= (1 - \theta) \left[R - \min \left\{ \frac{1}{2}, \theta q \right\} (R - 2c) \right] - c. \end{aligned}$$

The second line uses (B.24). Under the stated condition, $\mathcal{D}v \leq v$, whereas the solo-innovator HJB holds with equality. Hence v also satisfies the HJB inequality for development-only periods. A standard verification argument with $v(0) = 0$ then shows that allowing the agent to switch to development alone cannot raise the principal's value.

Finally, Assumption 2 gives $(1 - \theta)R \leq c$. Since $R - 2c > 0$, the bracket in the sufficient condition is at most R , so that assumption implies the condition. ■