Selecting for Innovation*

Gustavo Manso†

October 10, 2026

ABSTRACT

A firm that wants innovation must attract workers whose suitability for exploration it cannot observe. Any contract that attracts explorers is also open to exploiters. In a two-period model with hidden actions, explorers and exploiters generate different performance histories. The firm therefore screens by placing pay on histories relatively more likely under the intended activity than under mimicry, so the timing and path dependence of compensation become instruments of selection. Exploiters are paid for each success; explorers mainly for long-run success and, only when early failure is sufficiently diagnostic, after early failure. Commitment, termination, and feedback screen by the same logic. A common contract can help overcome a lack of commitment; separate contracts earn their keep when termination or feedback must differ. Two predictions follow. Exploration-oriented incentives change who enters, remains in, and chooses exploratory work, not merely innovative output. Competition for talent shifts where pay sits across histories, not merely its level.

Keywords: Innovation, Incentives, Contract Theory, Selection, Exploration, Exploitation.
JEL Codes: D86, J33, M52, O31.

*I thank Pierre Azoulay, Ed Lazear, John Roberts, Carolyn Stein, Scott Stern, Steve Tadelis, and seminar participants at various venues for helpful discussions and comments. All errors are my own.

†Haas School of Business, University of California, Berkeley. Email: manso@berkeley.edu.

I Introduction

Innovation incentives affect not only what workers do, but also who enters, remains in, and chooses particular jobs or contracts. This selection margin is well understood for standard performance pay. When Safelite switched from hourly wages to piece rates, productivity rose by 44%. Lazear (2000) attributes roughly half of that gain to the response of incumbent workers, with the remainder reflecting changes in workforce composition and related effects. Incentive schemes therefore operate through two margins: motivation changes what existing workers do, while selection changes who works under the scheme.

For innovation, however, the relevant dimension of selection is different. Standard sorting under performance pay is largely about productivity, effort cost, or preferences. Here the heterogeneity that matters is comparative advantage in exploration relative to exploitation. Exploratory work uses a novel and uncertain approach, whereas exploitative work uses a known and conventional method. Contracts that attract workers suited to exploration need not be the steepest pay-for-performance contracts. They may instead delay rewards, tolerate early failure, and rely on commitment and feedback. How should a firm design incentives to motivate effort and attract workers suited to exploration or exploitation?

This paper studies that joint screening and incentive problem. Workers privately know their type and choose unobservable actions. In the benchmark, exploiters can use only the conventional method and explorers only the novel one; both can shirk. This restriction isolates selection by fixing each type's productive activity. The firm may serve both types, one, or neither. Each contract must motivate the intended activity and deter imitation by the other type.

The unifying mechanism is that performance histories are endogenous signals of worker type. Explorers and exploiters generate different histories under their respective productive plans and when they mimic the other contract. Subject to incentive compatibility, the firm prefers to place incentives or rents on histories that are relatively more likely under the intended type than under the mimicking type's best response. The timing and path dependence of compensation, not merely the strength of pay for performance, therefore become instruments for selection, as do commitment, termination, and feedback.

That mechanism first shows up in the shape of pay. Under exploitation the contract resembles repeated pay for performance, with a contemporaneous bonus for each success. When the exploiter must be promised more than that baseline, the extra utility goes to first-period success, where it leaks least to an explorer who takes the contract. Under

exploration, baseline compensation is backloaded toward two consecutive successes and recovery after an early setback. Extra explorer utility goes first to long-run success, which a mimicking exploiter who shirks never reaches. Only when promised pay is high and early failure is sufficiently diagnostic does the contract pay after failure.

The same logic extends beyond the wage schedule, beginning with commitment. Its value lies in screening: by protecting rents placed on late, informative histories, it strictly lowers what an exploiter can obtain by taking the explorer contract. With feedback and retention fixed, an organization serving only exploiters would pay nothing for commitment.

Termination draws the contrast more sharply. After the same failure, with the same continuation opportunity, the firm may retain an explorer and dismiss an exploiter. When the explorer contract with retention pays after failure, dismissing an explorer saves no rent and only gives up the value of retaining him; dismissing an exploiter makes the exploitation contract less attractive to explorers.

Feedback cuts both ways. When blind exploration does not cover effort and the two wage-floor payments, every explorer the firm employs receives feedback. Explorers who take the exploitation contract also value it, so the optimal exploitation contract withholds it, lowering what they can obtain there without changing what exploiters do or cost. When the firm can hide the outcome, withholding can replace dismissal of exploiters, keeping them at work.

Because pay and these practices differ across tracks, serving both types might seem to require two contracts. In the wage-only benchmark, it does not: a single contract can give each type the same utility as any menu of two, at the same wage cost. But a single contract is not robust to competition for talent. When a higher outside option raises one type's pay and not the other's, it must be redesigned, while two contracts are only topped up, so I select the two contracts. Without commitment a single contract earns its keep. With explorers and exploiters under the same contract, the firm can sustain a rent for explorers who succeed early, because it cannot cut an explorer's later pay without losing the exploiters' effort. Track-specific contracts earn theirs when a firm hires only one type, or when termination or feedback must differ across tracks.

Two predictions follow. The first concerns which arrangements attract which workers: offered an arrangement built for exploration and one built for exploitation, workers choose the one designed for their comparative advantage. Evidence that exploration-oriented arrangements raise innovative output therefore mixes the behavior of incumbents with changes in composition, and tests should measure composition directly, classifying workers by their comparative advantage in exploration before the incentive change

rather than by generic ability. The second concerns where additional pay sits. When competition for talent raises the pay a type must be promised, the increase lands on different histories on the two tracks: success bonuses for exploiters and, under commitment, long-run rewards for explorers. When promised pay is high and early failure is sufficiently diagnostic, explorers are also paid after failure.

The paper relates to four literatures. The first concerns incentives and worker selection. Salop and Salop (1976) and Lazear (1986) show how compensation can induce self-selection, while Lazear (2000) and Dohmen and Falk (2011) provide evidence that workers sort across incentive schemes, the latter along productivity, risk attitudes, and self-assessment. Oyer and Schaefer (2005) find evidence consistent with sorting and retention explanations for broad-based stock options. Leaver et al. (2021) separate the selection and effort effects of performance pay experimentally. On the theory side, Foarta and Sugaya (2021) jointly choose assignment probabilities and success bonuses when workers privately know their effort costs and act unobservably, and Castro-Pires and Georgiadis (2026) show that steeper pay can worsen the skill composition of the workforce. Bénabou and Tirole (2016) show that competition for talent escalates performance pay, which firms use to attract more productive workers, and shifts effort away from unrewarded tasks. What distinguishes the present paper is the dimension along which workers sort: comparative advantage in exploration relative to exploitation.

The second literature studies incentives and institutions for innovation. Holmstrom (1989) emphasizes the problem created by noisy performance measures in innovative activities, while Aghion and Tirole (1994) study control rights under incomplete contracting. Hellmann (2007) and Hellmann and Thiele (2011) analyze employee innovation, entrepreneurship inside firms, and multitasking. Manso (2011) formulates the explicit exploration–exploitation problem and shows how tolerance for early failure, long-run rewards, commitment, job security, and feedback motivate a given worker to experiment. Subsequent theory studies innovation contests, prizes versus contracts, and experimentation in organizations (Halac et al., 2017; Che et al., 2021; Moroni, 2022), while a large empirical literature examines incentive horizons, failure tolerance, employment arrangements, and innovative activity in field and experimental settings (Azoulay et al., 2011; Ederer and Manso, 2013; Tian and Wang, 2014; Howell, 2021). This paper keeps those instruments and changes what they do. In Manso (2011) the same agent can shirk, exploit, or explore, so the instruments motivate that agent to experiment. Here each worker is suited to one activity, and the instruments determine which workers enter, remain in, and select into exploratory work.

The third literature studies economic experimentation and its delegation. Classical work analyzes learning through experimentation (Weitzman, 1979; Roberts and Weitzman, 1981; Moscarini and Smith, 2001), and strategic models examine informa-

tion spillovers (Bolton and Harris, 1999; Keller et al., 2005). Delegated-experimentation models study contracting with privately informed or privately acting agents (Bergemann and Hege, 2005; Hörner and Samuelson, 2013). Halac et al. (2016) use termination dates and time-dated bonuses for success to screen an experimenting agent's privately known ability under hidden effort. Gomes et al. (2016) use tolerance for failure to screen an entrepreneur's private return from exploitation when experimentation is observable. Khalil et al. (2020, 2026) use the length of experimentation and the timing of payments, including rewards after failure, to screen experimentation efficiency or beliefs about project costs; the 2026 model also includes hidden effort. In this paper, workers differ in their productive capabilities across exploration and exploitation, and whether a worker is exploring, exploiting, or shirking is hidden. A worker who accepts the other type's contract may therefore pursue a different productive activity. The paper characterizes how compensation, commitment, termination, and feedback change the cost and feasibility of selecting workers for exploratory and exploitative work. Two periods are the minimal dynamic setting for this analysis: one early failure has to be interpreted, and how informative it is depends on the type that generated it.

The fourth connection is to organizational ambidexterity, which asks how firms organize exploration and exploitation (March, 1991; Tushman and O'Reilly, 1996; O'Reilly and Tushman, 2013). The selection perspective adds that compensation systems, career paths, failure policies, and employment arrangements can attract different worker types as well as shape behavior. This does not imply that separate units are generally necessary. When commitment is difficult, a common contract can sustain exploration and exploitation within the same unit at lower wage cost than separate contracts. Track-specific contracts become valuable when the firm specializes in one track or when termination or feedback practices need to differ across tracks.

The paper is organized as follows. Section II introduces the model and the solution strategy. Section III characterizes the optimal exploiter and explorer contracts. Section IV studies commitment, Section V studies termination, and Section VI studies feedback. Section VII asks when a common contract should serve both types and when separate contracts are valuable. Section VIII develops empirical implications. Section IX concludes. Proofs are in the Appendix and Online Appendix.

II The Model

This section sets up a two-period principal–agent model in which a firm chooses which types to serve before observing whether the worker is an explorer or an exploiter. A limited-liability contract conditions wages on the history of successes and failures.

The worker's type and actions remain hidden. The firm may serve both types, one type while inducing the other to remain outside, or neither.

A Environment and Screening Problem

Time is two periods. In each period the agent chooses an action and produces either a success (S) or a failure (F), where $S > F \geq 0$. If the agent takes action $i \in \{0, 1, 2\}$, the probability of success is p_i, which may be unknown to both parties. Let $E[p_i]$ denote the prior mean and $E[p_i \mid S, j]$, $E[p_i \mid F, j]$ the expected probabilities that action i succeeds after observing S or F from action j. Learning is action-specific:

$$E[p_j] = E[p_j \mid S, i] = E[p_j \mid F, i] \quad \text{for all } j \neq i.$$

Action 1 represents a conventional method with a known success probability

$$p_1 = E[p_1] = E[p_1 \mid S, 1] = E[p_1 \mid F, 1].$$

Action 2 is a novel method with an uncertain success probability p_2. I assume that success raises beliefs about the effectiveness of the new approach:

$$E[p_2] \;<\; E[p_2 \mid S, 2].$$

After a failure with action 2, the agent can draw on the unsuccessful attempt to pursue a less ambitious version of the exploratory project. For example, work on a radical molecule may reveal a commercially viable backup candidate. This fallback is still action 2, not action 1, but it succeeds with the same probability as the conventional method: $E[p_2 \mid F, 2] = p_1$. Both types therefore face the same continuation opportunity after an early failure, so when the firm keeps an explorer and dismisses an exploiter after the same failure, the difference comes from screening, not technology.

In the baseline, the agent observes the first-period outcome and can adapt to the fallback after failure. Section VI studies the case in which feedback is withheld and the agent cannot make this adjustment.

The *exploratory ordering*

$$E[p_2] \;<\; p_1 \;<\; E[p_2 \mid S, 2]$$

ensures that the novel approach is initially less attractive than the conventional one, but can become strictly superior following a success.

Action 0 corresponds to shirking and yields success with probability $p_0 = 0$. With $p_0 = 0$ a success proves effort, so effort incentives alone require no rent above the wage floors, and neither commitment nor withholding feedback lowers the cost of motivating a given type, as commitment does for exploration and withholding feedback for exploitation in Manso (2011); their value here comes from selection. All other actions considered in the model, including the conventional approach (action 1) and any version of the novel approach (action 2, both after success and after failure), have success probabilities strictly between zero and one.

Taking a productive action in a period entails a private cost $c > 0$, while shirking is costless. Thus $c_0 = 0$ and $c_1 = c_2 = c$. The principal and the agent are risk neutral and share a unit discount factor. Actions are unobservable.

A contract specifies wages $w = \{w_F, w_S, w_{SF}, w_{SS}, w_{FF}, w_{FS}\}$ contingent on the history of successes and failures. Limited liability imposes a minimum wage $w_x \geq \underline{w} > 0$ for all outcomes $x \in \{F, S, SF, SS, FS, FF\}$. It is convenient to write

$$w_x = \underline{w} + \tilde{w}_x, \qquad \tilde{w}_x \geq 0,$$

and to refer to $\tilde{w} = \{\tilde{w}_F, \tilde{w}_S, \tilde{w}_{SF}, \tilde{w}_{SS}, \tilde{w}_{FF}, \tilde{w}_{FS}\}$ as the vector of *bonuses above the floor*. Let $\mathcal{Z}$ be the set of active contractible items. An item $z = (\tilde{w}, \phi, D)$ specifies a bonus vector, whether the firm discloses the first-period outcome to the agent before his second-period action ($\phi \in \{\text{provide}, \text{withhold}\}$), and the first-period outcomes after which employment ends ($D \subseteq \{S, F\}$). The baseline item provides feedback and retains the agent, so that $\phi = \text{provide}$ and $D = \varnothing$. Sections V and VI give the expressions for the other items.

The principal can commit to a long-term contract. Section IV removes this commitment: the continuation terms after the first period must then be ones the principal would choose at that point, given what acceptance and the first-period outcome have revealed. Lack of commitment adds this constraint to the problem below and changes nothing else.

An *action plan* is denoted $\langle {}_i{}^j_k \rangle$, where i is the first-period action, j the second-period action after a first-period success, and k the second-period action after a first-period failure. The principal earns revenue $r_1(y_1)$ in the first period and $r_2(y_1, y_2)$ in the second, where $y_t \in \{S, F\}$ is the period-t outcome; in the leading example $r_1(y_1) = y_1$ and $r_2(y_1, y_2) = y_2$. Allowing r_2 to depend on y_1 captures, in reduced form, a firm scaling up or down in response to first-period performance. Let $R(z, a)$ denote the principal's expected total revenue under contract z and action plan a, with no revenue after dismissal. For a fixed action plan, revenue depends on z only through its dismissal

policy D. Let $W(z,a)$ denote the agent's expected total wage payments, including wage floors, under contract z and action plan a. In the baseline, these payments are

$$\begin{aligned} W(z,\langle {}_{i}{}^{j}_{k}\rangle) = {} & \{E[p_i]\tilde{w}_S + (1-E[p_i])\tilde{w}_F\} \\ & + E[p_i]\{E[p_j \mid S,i]\tilde{w}_{SS} + (1-E[p_j \mid S,i])\tilde{w}_{SF}\} \\ & + (1-E[p_i])\{E[p_k \mid F,i]\tilde{w}_{FS} + (1-E[p_k \mid F,i])\tilde{w}_{FF}\} + 2\underline{w}. \end{aligned}$$

Let $C(z,a)$ denote the agent's expected total effort cost under contract z and action plan a, counting only periods in which he remains employed. In the baseline,

$$C(z,\langle {}_{i}{}^{j}_{k}\rangle) = c_i + E[p_i]\, c_j + (1-E[p_i])\, c_k.$$

Before contracting, the agent's type $\theta \in \{\mathrm{E},\mathrm{X}\}$ is privately known. The prior probability of type E is $\Pr(\theta = \mathrm{E}) = \mu \in (0,1)$. Types differ in the productive opportunities they can access. Exploiters can use the conventional action 1 at cost c, but cannot explore. Explorers can use the novel action 2 at cost c, but cannot exploit. This hard-type assumption is a stark benchmark, not a claim that workers literally face infinite costs of other tasks. It shuts down the within-agent choice between exploration and exploitation in order to isolate the selection margin.[1]

The corresponding sets of feasible action plans are

$$\mathcal{A}(\mathrm{E}) = \{\langle {}_{i}{}^{j}_{k}\rangle : i,j,k \in \{0,2\}\}, \qquad \mathcal{A}(\mathrm{X}) = \{\langle {}_{i}{}^{j}_{k}\rangle : i,j,k \in \{0,1\}\}.$$

In the baseline the agent remains employed for two periods. Section V studies termination after a first-period failure.

I focus on the region in which exploration generates more expected successes than exploitation, the *count ordering*,[2]

$$E[p_2] + E[p_2]E[p_2 \mid S,2] + (1-E[p_2])p_1 > 2p_1.$$

[1] The assumption can be relaxed. Online Appendix D.1 lets each type perform the other's activity at an additional cost δ. With both types hired, the separate contracts of Sections III–VI still minimize wage costs whenever δ is at least the larger of the difference between the outside options and a constant, whatever the level of rents. The same bound preserves minimum wage costs when the firm hires one type, so the cost comparisons behind the results on commitment, dismissal, and withholding feedback from exploiters are unchanged. When outside options come from jobs at other firms of this kind, open to both types, the appendix gives conditions under which every $\delta > 0$ suffices for a firm that hires only one type.

[2] Online Appendix D.2 studies environments in which exploration generates fewer expected successes than exploitation. The main results survive: exploitation remains repeated pay for performance, exploration remains backloaded, and commitment, withholding feedback, and the dismissal of exploiters still screen. The firm may even keep explorers when continuing is technologically inefficient.

The screening problem is not whether exploration is socially valuable, but whether the principal can profitably select the exploratory type while preventing exploiters from mimicking the exploration contract.

Each type $\theta \in \{\mathrm{E}, \mathrm{X}\}$ has an exogenous reservation utility $\overline{U}_\theta^{\mathrm{com}} \geq 0$, which captures competition for talent and outside offers. An agent who declines all contracts receives $\overline{U}_\theta^{\mathrm{com}}$.

Given an active contract z and a feasible action plan a for type θ, the agent's utility is

$$U_\theta(z, a) = W(z, a) - C(z, a),$$

and the principal's profit is

$$\Pi(z, a) = R(z, a) - W(z, a).$$

The intended action plan is not part of the contractible item. For each $\theta \in \{\mathrm{E}, \mathrm{X}\}$, the principal chooses $z_\theta \in \mathcal{Z} \cup \{\emptyset\}$ and, when the type is served, an intended plan $a_\theta \in \mathcal{A}(\theta)$. The symbol $\emptyset$ is the outside-option/null employment item.

Every maximization over plans, including within-type and cross-item deviations, is restricted to plans compatible with the feedback and dismissal policies of the item being evaluated.

For a contractible item z, define its value to type θ by

$$V_\theta(z) \equiv \max_{a' \in \mathcal{A}(\theta)} U_\theta(z, a').$$

Extend V_θ and Π to the null item by

$$V_\theta(\emptyset) \equiv \overline{U}_\theta^{\mathrm{com}}, \qquad \Pi(\emptyset, a) \equiv \Pi(\emptyset) \equiv 0.$$

Thus the null item generates no employment, output, wage or wage-floor payment, or effort cost in this relationship. It gives the worker the type-specific competitive outside option.

I maintain the *revenue condition*: with feedback, and whether the worker is retained or dismissed only after a failure, the plan that exerts productive effort in every employed period yields strictly more expected revenue net of effort than any other plan available to the type under the same policy or under a policy that dismisses after additional outcomes; and when such a plan responds to the first-period outcome, acting on it yields strictly more than acting blind. In the leading example, it is implied by the count

ordering and the *adaptive-exploration condition*

$$E[p_2](S-F) \;<\; c \;<\; p_1(S-F):$$

repeating the novel method without acting on the first-period outcome does not pay for its effort, while the conventional method and the fallback do.

The exploratory ordering, the count ordering, and the revenue condition are maintained throughout. Section V adds the low-floor condition, and any further condition is stated in the result that uses it.

The offered menu is

$$\mathcal{M}(z_{\mathrm{E}}, z_{\mathrm{X}}) = \{\emptyset, z_{\mathrm{E}}, z_{\mathrm{X}}\}.$$

If $z_{\mathrm{E}} = z_{\mathrm{X}} \neq \emptyset$, the menu contains one active common contract. Both types may accept it and choose different optimal plans because their feasible action sets differ. An *allocation* specifies a designated item and an intended action plan for each type. The principal solves the allocation problem

$$\begin{aligned} \max_{\substack{z_{\mathrm{E}}, z_{\mathrm{X}} \in \mathcal{Z} \cup \{\emptyset\} \\ a_{\mathrm{E}} \in \mathcal{A}(\mathrm{E}),\, a_{\mathrm{X}} \in \mathcal{A}(\mathrm{X})}} \quad & \mu\, \Pi(z_{\mathrm{E}}, a_{\mathrm{E}}) + (1-\mu)\, \Pi(z_{\mathrm{X}}, a_{\mathrm{X}}) \\ \text{s.t.} \quad & U_\theta(z_\theta, a_\theta) \geq U_\theta(z_\theta, a), \quad \forall a \in \mathcal{A}(\theta), \quad \text{if } z_\theta \neq \emptyset, \\ & z_\theta \in \arg \max_{z \in \mathcal{M}(z_{\mathrm{E}}, z_{\mathrm{X}})} V_\theta(z), \qquad \theta \in \{\mathrm{E}, \mathrm{X}\}. \end{aligned}$$

The first constraint ensures that each served type chooses its intended action plan. The second combines participation, self-selection, and exclusion: each type must weakly prefer its designated item to every alternative, including its outside option. Limited liability is built into $\mathcal{Z}$.

Omitting a contract for one type does not mean that the firm observes that type and refuses to make an offer. It is feasible only when that type prefers its outside option to every offered contract. Thus explorer-only hiring requires

$$V_{\mathrm{X}}(z_{\mathrm{E}}) \leq V_{\mathrm{X}}(\emptyset) = \overline{U}_{\mathrm{X}}^{\mathrm{com}}.$$

Likewise, exploiter-only hiring requires $V_{\mathrm{E}}(z_{\mathrm{X}}) \leq \overline{U}_{\mathrm{E}}^{\mathrm{com}}$.

B Solution Strategy

I solve the screening problem by first characterizing the least-cost contract for each type, conditional on its intended action plan and a pair of utility requirements. The

explorer contract must satisfy $V_{\mathrm{E}}(z_{\mathrm{E}}) \geq \overline{U}_{\mathrm{E}}$ and $V_{\mathrm{X}}(z_{\mathrm{E}}) \leq \overline{U}_{\mathrm{X}}$, and the exploiter contract $V_{\mathrm{X}}(z_{\mathrm{X}}) \geq \overline{U}_{\mathrm{X}}$ and $V_{\mathrm{E}}(z_{\mathrm{X}}) \leq \overline{U}_{\mathrm{E}}$: each contract has its own type's requirement as a floor and the other type's as a ceiling. Each type then gets at least its requirement from its own contract and at most it from the other, so it weakly prefers its own. Given the pair, the two contracts can be designed separately. Appendix A shows that this decomposition loses nothing, including when both types accept a common contract and choose different plans.

When both types are served, the firm chooses the two requirements subject to participation. For selective hiring, the requirements are the two outside options: the hired type must receive at least its outside option, while the excluded type must obtain no more than its outside option from the offered contract. Incentives or limited liability may require paying the hired type more. If exclusion cannot be achieved, selective hiring is infeasible.

A requirement pair is feasible when a contract can satisfy the incentive and utility bounds. Serving both types requires feasibility for both contracts. Unless noted otherwise, the contract characterizations below apply to feasible requirement pairs. Appendix B.1 gives the baseline conditions. The extension conditions are in Online Appendix A.3.

Each conditional contract problem is a finite linear program. I first find the least-cost contract that motivates the intended activity, setting aside participation and cross-type imitation. I then ask whether pay can be reallocated to deter imitation while preserving incentives. Finally, when the intended type's utility requirement binds, I determine how to deliver the additional surplus.

Likelihood-ratio comparisons guide these choices. A history is more diagnostic when it is relatively more likely under the intended plan than under a particular deviation or mimicking plan. Paying on such histories can reward the intended type while limiting what the other type obtains. Preserving incentives may require changing several wages together. Appendix B proves the baseline results. Appendix C gives the central organizational arguments, with the remaining proofs in Online Appendix B and the arguments for Section VII in Online Appendix C.

When several contracts implement the same action plan at the same cost, I select one in three steps: the first two by what the other type would obtain, the third by timing. First, I minimize the other type's payoff from accepting the contract. This is the screening objective itself, applied after cost: the firm does not observe who applies, and the least-cost contract that is least attractive to the other type is the one that relaxes self-selection and exclusion, so it remains optimal whenever those constraints

bind. Second, among contracts with the same cost and the same mimicking payoff, I keep those that would remain optimal if competition for talent raised the utility the firm must promise: outside options are not known exactly, and a contract that would have to be redesigned, rather than topped up, as the requirement rises is not one the firm would post. If the other type could not accept the contract, every least-cost contract would pass both steps, since an unconditional raise delivers additional utility at its minimum cost. Third, among the remaining contracts, I choose the one that pays earliest, that is, the one with the largest expected payment at the interim date. Every contract this step sets aside costs the same and gives each type the same payoff, but pays less at the interim date: the firm would be holding part of the worker's pay, acting as his bank. Online Appendix D.3 obtains the last two steps as the limit of two small perturbations: a chance that the required utility rises after the contract is posted, and a smaller chance that the worker needs cash at the interim date.[3]

Minimum wage costs, the least a contract can leave the other type, the service sets that are feasible, and the retention thresholds of Section V do not depend on which contract is selected. Where mimicking raises the other type's pay above what he would otherwise earn, the first step is not a choice: only contracts that leave him the least are optimal. The last two steps matter only with commitment and feedback: without commitment the first step already leaves a single contract, and the results without feedback hold for every optimum, display one optimal contract, or use only the first step. Throughout, "the optimal contract" and "the optimal allocation" refer to one these steps select.

Given the plan, the practices, and the utility it must deliver, the least-cost contract does not depend on the firm's revenue. Revenue, the type distribution, and outside options determine which types the firm serves and whether it retains an explorer after an early failure: at the full plans, a premium on first-period success favors exploiters, since $p_1 > E[p_2]$, whereas a premium on two successes favors explorers, since $E[p_2]E[p_2 \mid S, 2] > p_1^2$.

The analysis below focuses on the maintained productive plans $\langle {}_2{}^2_2 \rangle$ and $\langle {}_1{}^1_1 \rangle$ and, with dismissal, on their versions that end employment after a first-period failure. Proposition 1 gives conditions under which every optimal allocation takes this form. One condition concerns exploration without feedback. The *blind-loss condition* holds if every explorer plan that explores without feedback generates expected revenue net of effort below $2\underline{w}$, so that blind exploration does not cover the two wage-floor payments.

[3]Together with the cash need, a small chance that effort turns out to be cheap after hiring selects the same explorer contracts at every required utility (Online Appendix D.4).

PROPOSITION 1 (Productive employment)*: Suppose, in addition to the maintained assumptions, that* $\underline{w} \leq c$*, the blind-loss condition holds, and* $\overline{U}_{\mathrm{E}}^{\mathrm{com}}, \overline{U}_{\mathrm{X}}^{\mathrm{com}} \geq 2\underline{w}$*. For every* $\mu \in (0,1)$ *the allocation problem has a solution, and every optimal allocation induces productive effort from each worker it serves in every period of employment and ends employment, if at all, only after a first-period failure.*

Under these conditions, the proof eliminates shirking during employment and dismissal after success while preserving participation, self-selection, and exclusion. The only productive continuation screening can sacrifice is the one after a first-period failure, the trade-off of Section V. The proof is in Appendix C.4. The contract characterizations below do not require these additional conditions.

III Optimal Exploiter and Explorer Contracts

This section characterizes the least-cost contracts for exploitation and exploration.

A Exploiter Contract

I first characterize the least-cost contract that implements exploitation, $a_{\mathrm{X}} = \langle {}_{1}{}^{1}_{1} \rangle$. This contract provides the benchmark against which the explorer contract's distinctive features will be identified.

Because $p_0 = 0$, success perfectly distinguishes productive effort from shirking on each effort margin. The least-cost way to induce second-period effort is therefore the familiar one-period bonus c/p_1 following success. The same logic applies to first-period effort, so the baseline contract pays a constant bonus for each success, and the exploiter earns baseline utility $2\underline{w}$.

When the exploiter's utility requirement exceeds that baseline level, the principal places the additional surplus where it leaks least to a misreporting explorer. First-period success is the least leaky of the histories that end in success: S is more likely under exploitation relative to exploration than either SS or the failure-success branch. A success followed by a failure is more diagnostic still, but pay there would have to be matched after a second success to keep the exploiter working, and so amounts to pay for first-period success.

PROPOSITION 2 (Exploiter Contract)*: If* $\overline{U}_{\mathrm{X}} \leq 2\underline{w}$*, the optimal exploiter contract implements* $a_{\mathrm{X}} = \langle {}_{1}{}^{1}_{1} \rangle$ *with*

$$\tilde{w}_S = \tilde{w}_{SS} = \tilde{w}_{FS} = \frac{c}{p_1}, \qquad \tilde{w}_F = \tilde{w}_{SF} = \tilde{w}_{FF} = 0,$$

and the exploiter receives $2\underline{w}$*.*

If $\overline{U}_{\mathrm{X}} > 2\underline{w}$*, the principal starts from this baseline contract and delivers the required additional surplus* $\Delta = \overline{U}_{\mathrm{X}} - 2\underline{w} > 0$ *by increasing* $\tilde{w}_S$ *only, while keeping* $\tilde{w}_{SS}$ *and* $\tilde{w}_{FS}$ *at their incentive-minimal levels.*

B Explorer Contract

I now characterize the least-cost contract that implements exploration, $a_{\mathrm{E}} = \langle {}_{2}{}^{2}_{2} \rangle$. Its distinctive form reflects how informative double success and early failure are about exploration relative to exploitation.

The baseline contract rewards two consecutive successes rather than first-period success alone. After an exploratory success, learning makes a second success more likely than under exploitation, so delaying the reward better distinguishes the two activities. A separate bonus for recovery maintains effort after failure, and the explorer receives baseline utility $2\underline{w}$.

When additional surplus is required, the principal first raises the double-success reward without increasing the exploiter's best-response payoff. Once this is no longer possible, the principal compares the exploration-to-exploitation likelihood ratios of double success and early failure. A payment after early failure must be accompanied by higher double-success pay, so that an exploiter cannot gain by switching to first-period shirking.

PROPOSITION 3 (Explorer Contract)*: If* $\overline{U}_{\mathrm{E}} \leq 2\underline{w}$*, the optimal explorer contract implements* $a_{\mathrm{E}} = \langle {}_{2}{}^{2}_{2} \rangle$ *with*

$$\tilde{w}_S = \tilde{w}_F = \tilde{w}_{SF} = \tilde{w}_{FF} = 0, \qquad \tilde{w}_{FS} = \frac{c}{p_1}, \qquad \tilde{w}_{SS} = \frac{(1 + E[p_2])c}{E[p_2]E[p_2 \mid S, 2]},$$

and the explorer receives $2\underline{w}$*.*

If $\overline{U}_{\mathrm{E}} > 2\underline{w}$*, the principal starts from that baseline contract and delivers the additional surplus* $\Delta = \overline{U}_{\mathrm{E}} - 2\underline{w} > 0$ *by raising* $\tilde{w}_{SS}$ *alone, stopping as soon as either the surplus is*

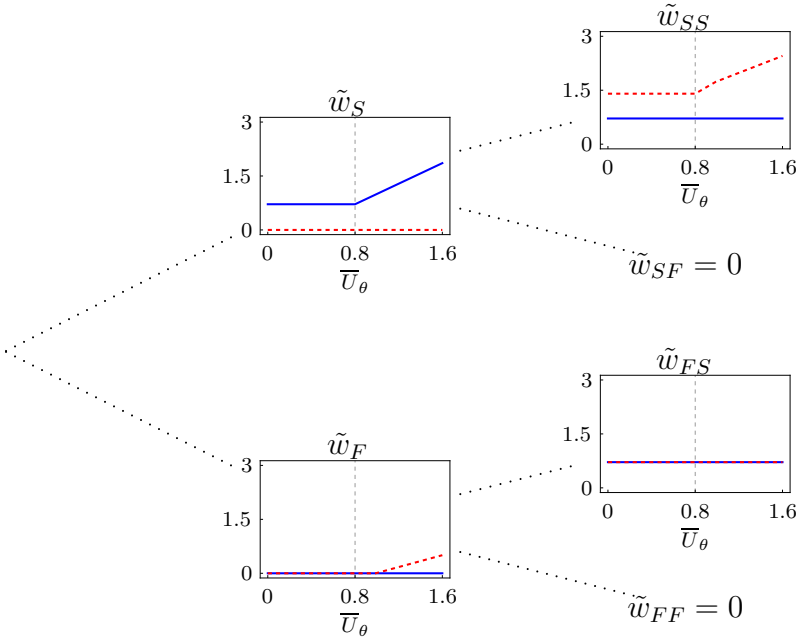


Figure 1. Optimal exploiter contract (solid blue curves) and optimal explorer contract (dashed red curves) under the benchmark parameters, each as a function of the utility requirement $\overline{U}_\theta$ of the type it serves. Exploitation is paid as repeated pay for performance, and extra surplus raises only $\tilde{w}_S$. Exploration is paid on double success and on recovery: extra surplus first raises $\tilde{w}_{SS}$ and, once the misreporting exploiter's first-period IC binds, $\tilde{w}_F$ and $\tilde{w}_{SS}$ together.

exhausted or the misreporting exploiter is just indifferent between first-period work and shirking. Any surplus still owed beyond that point is delivered on the more diagnostic history. If double success is more diagnostic of exploration than first-period failure, it is optimal to continue raising $\tilde{w}_{SS}$ alone. If first-period failure is at least as diagnostic, it is optimal to raise $\tilde{w}_F$ and $\tilde{w}_{SS}$ together, in the fixed proportion that holds the misreporting exploiter's first-period incentive constraint at equality.

Figure 1 plots both contracts at the benchmark parameters,[4] for which first-period failure is the more diagnostic history, so past the threshold the explorer contract raises $\tilde{w}_F$ and $\tilde{w}_{SS}$ together.

The next three sections relax the baseline institutional assumptions. Commitment, termination, and feedback are not separate departures from the selection logic. They determine where rents can be promised, whether early failure is insured or punished,

[4]The benchmark parameters are $p_0 = 0$, $E[p_2] = 0.60$, $E[p_2 \mid S, 2] = 0.95$, $p_1 = 0.70$, $c = 0.50$, and a wage floor $\underline{w} = 0.40$. The plotted wage schedules do not depend on the firm's revenue function, holding the intended plans and utility requirements fixed.

and whether exploration preserves its option value. Section VII then asks whether a common contract can serve both types.

IV Lack of Commitment

This section asks how selection changes when the principal cannot commit to a long-term contract and must offer a sequence of one-period contracts instead, with the continuation contract chosen after the first-period outcome is observed. No payment can then be promised merely to influence an action already taken. Under the revenue condition, productive continuation is sequentially optimal after failure and after exploratory success. Contracts are compared with their committed counterparts at a fixed utility requirement for the type they serve.

A Exploitation without commitment

The exploiter contract of Proposition 2 pays only for current success, so each payment falls in the period whose effort it motivates and no part of it is a promise about the future. The baseline repeats the one-period incentive contract, so it remains optimal once the first-period action is sunk. Additional surplus is delivered through the first-period success payment $\tilde{w}_S$, which the first-period contract itself covers.

PROPOSITION 4 (Exploitation without commitment)*: For a given exploiter utility requirement, the optimal exploiter contract of Proposition 2 is implemented by a sequence of one-period contracts: the baseline pays only current success bonuses, and additional surplus is delivered by raising $\tilde{w}_S$ alone. Its cost, and the payoff an explorer obtains by accepting it, are unchanged by commitment.*

B Exploration without commitment

I now turn to the contract that implements the explorer's plan $\langle {}_2{}^2_2 \rangle$. Under commitment the explorer's rent sits on later histories. Once the first-period action is sunk, $\tilde{w}_{SS}$ and $\tilde{w}_{FS}$ sit at their one-period incentive minima, and exploration must be induced by the spread between $\tilde{w}_S$ and $\tilde{w}_F$. The failure-tolerance payment itself is not the obstacle: under the committed contract, it is combined with rewards for long-term success to prevent the explorer from shirking. Those above-minimum continuation rewards, however, are not sequentially optimal.

Additional utility must therefore be delivered by raising $\tilde{w}_S$ and $\tilde{w}_F$ equally, holding their spread and with it the explorer's incentive to explore fixed. A common first-period transfer says nothing about which activity the worker chose, and an exploiter accepting the explorer contract collects it in full. Because early success is relatively more likely under exploitation, that exploiter works in the first period from the outset, so each extra dollar owed to the explorer raises the mimicking payoff one for one. Unlike in Manso (2011), commitment therefore does not lower the direct expected wage cost of serving the explorer at a fixed requirement at which both contracts are feasible. What it lowers, at every such requirement, is the payoff a mimicking exploiter can obtain from the explorer contract, and that is the margin that can relax self-selection or exclusion and change which allocations are feasible.

PROPOSITION 5 (Exploration without commitment)*: If* $\overline{U}_{\mathrm{E}} \leq 2\underline{w}$*, the optimal contract implementing* $\langle {}_{2}{}^{2}_{2} \rangle$ *is the baseline contract*

$$\tilde{w}_F = \tilde{w}_{SF} = \tilde{w}_{FF} = 0, \qquad \tilde{w}_{FS} = \frac{c}{p_1},$$
$$\tilde{w}_{SS} = \frac{c}{E[p_2 \mid S, 2]}, \quad \tilde{w}_S = \frac{c}{E[p_2]},$$

and the explorer receives $2\underline{w}$*. If* $\overline{U}_{\mathrm{E}} > 2\underline{w}$*, the principal starts from this baseline and delivers the required additional surplus by raising* $\tilde{w}_S$ *and* $\tilde{w}_F$ *equally, holding* $\tilde{w}_{SS}$ *and* $\tilde{w}_{FS}$ *at their baseline levels.*

Figure 2 plots the mimicking exploiter's payoff, or *ride-along* payoff: the utility an exploiter earns by accepting the explorer contract and best-responding. Without commitment the mimicking payoff, already above $2\underline{w}$ at the baseline, rises one for one once the requirement exceeds $2\underline{w}$. With commitment the mimicking payoff starts at $2\underline{w}$, is initially flat, and thereafter rises more slowly. The two payoff schedules never meet. Online Appendix Figure OA.1 plots the two contracts.

COROLLARY 1 (Value of commitment)*: Conditional on the exploiter's utility requirement, commitment leaves the optimal exploiter contract and its cost unchanged. For every explorer utility requirement, it strictly reduces the minimum ride-along utility of a misreporting exploiter under an explorer-implementing contract. Commitment therefore weakly raises optimal profit and weakly lowers the minimum cost of screening.*

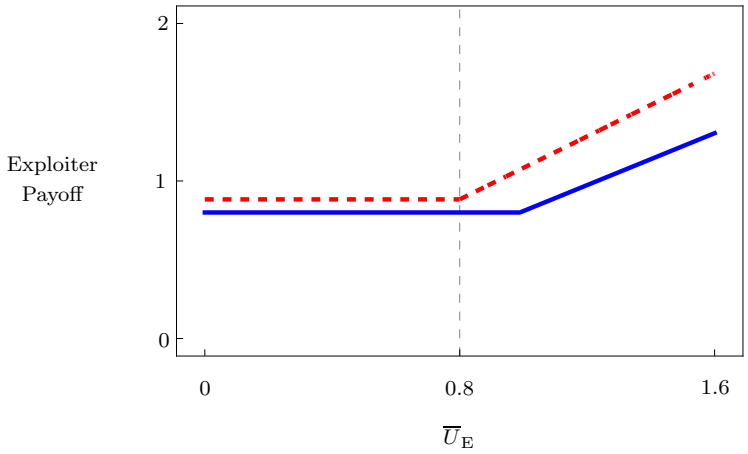


Figure 2. Exploiter ride-along payoff under the explorer contract with commitment (solid blue curve) and without commitment (dashed red curve). Commitment moves rents to informative histories and keeps the mimicking payoff below its no-commitment counterpart.

Read as a willingness to pay for institutions that make long-term promises credible, commitment is worth nothing directly to an organization serving only exploiters and weakly more to one serving explorers. With separate contracts, that value is strict when mimicking raises the exploiter's requirement above his outside option or when commitment makes a more profitable explorer-serving allocation feasible. When it is, there is a nonempty interval of common positive commitment costs at which an explorer-serving organization invests in commitment and an exploiter-only organization does not. Section VII shows when a common contract can overcome a lack of commitment instead.

V Termination

This section studies termination after first-period failure: firing a manager, or discontinuing funding for a project, after poor interim performance. Dismissal disciplines, because it withdraws the income that continued employment would have paid, and it screens, because it withdraws that income at a history an explorer reaches more often than an exploiter. Both effects are bought with production forgone. Feedback is provided throughout this section; Section VI shows what changes when the firm can withhold the outcome.

Throughout this section, assume the *low-floor condition*

$$\underline{w} < \frac{c}{p_1}.$$

The continuation floor a first-period success secures in expectation, $p_1\underline{w}$, falls short of the effort cost c. Dismissal therefore assists effort incentives without inducing effort on its own. Because $E[p_2] < p_1$, the condition also implies $E[p_2]\underline{w} < c$.

I write $\langle {}_{i}{}^{j}_{t} \rangle$ for the plan with action i in period 1, action j after first-period success, and termination after first-period failure. The exploiter and explorer plans with termination are $\langle {}_{1}{}^{1}_{t} \rangle$ and $\langle {}_{2}{}^{2}_{t} \rangle$.

A Exploiter Contract with Termination

On the exploitation track dismissal substitutes for money. Failure costs the worker the floor otherwise collected in period 2, so first-period effort can be induced with a smaller success payment. Additional surplus is placed on first-period success, a history the target exploiter reaches relatively more often than a mimicking explorer.

PROPOSITION 6 (Exploiter contract with termination)*: If* $\overline{U}_{\mathrm{X}} \leq \underline{w}$*, the optimal contract implementing* $a_{\mathrm{X}} = \langle {}_{1}{}^{1}_{t} \rangle$ *is the baseline contract*

$$\tilde{w}_{SS} = \frac{c}{p_1}, \qquad \tilde{w}_{S} = \frac{c}{p_1} - \underline{w}, \qquad \tilde{w}_{F} = \tilde{w}_{SF} = 0.$$

If $\overline{U}_{\mathrm{X}} > \underline{w}$*, the principal starts from this baseline and delivers the required additional surplus by increasing* $\tilde{w}_S$ *only, leaving* $\tilde{w}_{SS} = c/p_1$ *and the remaining payments unchanged.*

Online Appendix Figure OA.2 displays the exploiter contract with and without termination. The terminated contract is strictly less attractive to a misreporting explorer at every utility requirement, as Lemma 2 in the Appendix records. When only exploiters are served, that same reduction relaxes exclusion instead. Whether the principal fires is a separate comparison of the cost of inducing effort, the explorer's payoff from the item, and the production forgone. Corollary 2 makes that comparison on both tracks.

B Explorer Contract with Termination

On the exploration track dismissal is blunter. First-period failure is not evidence against exploration, since it arrives more often under exploration than under exploitation. Shutting down continuation therefore lowers the intended explorer's payoff along with the mimic's, and whatever part of the lost value the explorer is owed must be returned on the surviving histories.

The baseline pays on double success alone, and additional surplus is first delivered there, where a misreporting exploiter still prefers to shirk and collect the floor. Severance never rises alone, because raising pay after failure would otherwise induce a mimicking exploiter to shirk in the first period.

PROPOSITION 7 (Optimal explorer contract with termination)*: If* $\overline{U}_{\mathrm{E}} \leq \underline{w}$*, the optimal contract implementing* $\langle {}_{2}{}^{2}_{t} \rangle$ *is the baseline contract*

$$\tilde{w}_S = \tilde{w}_F = \tilde{w}_{SF} = 0, \qquad \tilde{w}_{SS} = \frac{(1 + E[p_2])c - E[p_2]\underline{w}}{E[p_2]E[p_2 \mid S, 2]}.$$

If $\overline{U}_{\mathrm{E}} > \underline{w}$*, the principal starts from this baseline and raises* $\tilde{w}_{SS}$ *alone, stopping when either the required surplus is delivered or the misreporting exploiter is just indifferent between working and shirking in the first period. Any surplus still owed is delivered on the more diagnostic history. If double success is more diagnostic of exploration than first-period failure,* $\tilde{w}_{SS}$ *continues to rise alone. If first-period failure is at least as diagnostic,* $\tilde{w}_F$ *and* $\tilde{w}_{SS}$ *rise together, in the fixed proportion that holds the misreporting exploiter's first-period incentive constraint at equality.*

Online Appendix Figure OA.3 displays the explorer contract under the benchmark parameters, for which first-period failure is the more diagnostic history. Termination initially lowers $\tilde{w}_{SS}$, because lost continuation income does part of the work of first-period incentives. At a fixed explorer utility requirement, termination only weakly lowers the payoff a mimicking exploiter can obtain from this contract; Lemma 3 in the Appendix identifies when this reduction is strict.

Let g denote the operating value of retaining the worker after first-period failure, before the continuation wage floor. The same opportunity is available on both tracks: idle retention yields $r_2(F, F)$, while productive continuation yields $p_1 r_2(F, S) + (1 - p_1) r_2(F, F) - c$. Thus

$$g = r_2(F, F) + \big[p_1[r_2(F, S) - r_2(F, F)] - c\big]^+,$$

which is $F + [p_1(S - F) - c]^+$ in the leading example. Because g is net of the fallback's effort cost, wage costs below are counted net of the pay that reimburses that effort.

The *severance region* consists of environments in which the optimal exploration contract with retention reaches or exceeds the threshold at which severance begins, whether the exploitation track retains or terminates after failure. If first-period failure is at least

as diagnostic of exploration as double success and both types are served, the firm enters this region once competition for talent is strong enough: the explorer's own outside option reaches that threshold, or the exploiter's is high enough that the exploitation contract, even one that terminates after failure, would give an explorer at least that much.

COROLLARY 2 (Termination across exploration and exploitation)*: Fix a feasible set of served types and, when comparing termination policies on one track, hold the other active track's policy fixed. For each active track θ there is a threshold τ_θ such that the optimal contract on that track terminates after first-period failure when $g < \tau_\theta$ and retains when $g > \tau_\theta$.*

(i) *On an active exploitation track on which retention is feasible, τ_{X} is the sum of two nonnegative terms, each per unit of forgone continuation: the incentive cost that termination saves directly on the exploiter's contract, and the screening cost it saves by making that contract less attractive to explorers. If only termination is feasible, $\tau_{\mathrm{X}} = +\infty$.*

(ii) *On an active exploration track in the severance region, $\tau_{\mathrm{E}} = 0$. Termination saves no wage cost, so the principal retains after exploratory failure whenever $g > 0$.*

(iii) *Suppose both tracks are active and the environment lies in the severance region. If $0 < g < \tau_{\mathrm{X}}$, the principal retains after exploratory failure and terminates after exploitative failure.*

The corollary ties job security to the pay that follows early failure. In the severance region, the retention thresholds satisfy $0 = \tau_{\mathrm{E}} \leq \tau_{\mathrm{X}}$. The firm needs weakly less operating value to keep an explorer after failure than to keep an exploiter. It needs strictly less exactly when dismissing exploiters lowers the optimized cost of the menu.

The leniency toward explorers comes from a swap. In the severance region, the firm can replace idle retention with dismissal and add the forgone continuation floor to severance. Total pay after failure, the explorer's incentives, a mimicking exploiter's payoff, and total wage cost are unchanged. Dismissal saves nothing on wages and gives up the operating value of retention. Even idle retention is preferable when $r_2(F, F) > 0$. Job security here reflects the absence of a saving from dismissal, with exploration incentives unchanged.

Exploitation allows no such swap. With exploration retained, an exploiter could take the exploration contract, shirk, and still collect the wage floor in both periods. The

exploiter's utility requirement is therefore at least $2\underline{w}$. Dismissal still disciplines effort, but the income it takes away must be paid back to meet that requirement, so it saves nothing on the exploiter's own wages at the fixed utility requirement. Its value comes entirely from screening savings. A contract that ends after failure is less attractive to explorers, which lowers the pay needed to keep them on the exploration track when the exploitation contract raises that pay above their outside option. When $0 < g < \tau_{\mathrm{X}}$, the firm therefore keeps the explorer and dismisses the exploiter, even though both have the same productive opportunity after failing.

The tolerance of exploratory failure also goes beyond what the principal would choose after selection. Once incentive and screening consequences are sunk, the principal would retain a worker only when $g \geq \underline{w}$. Because $\tau_{\mathrm{E}} = 0$, exploration shows excessive continuation whenever $0 < g < \underline{w}$. Exploitation exhibits excessive termination when $\underline{w} < g < \tau_{\mathrm{X}}$, and excessive continuation when $\tau_{\mathrm{X}} < g < \underline{w}$. These comparisons measure departures from the post-selection retention decision, not social or technological inefficiency.

VI Feedback

The principal often learns how first-period work turned out before the worker can: it sees how customers respond, what a test shows, or how outside reviewers assess the output. This section asks whether the principal should disclose that information. The principal observes first-period performance and commits ex ante in each menu item to providing or withholding feedback, the disclosure of the first-period outcome. Without feedback, the worker chooses his second-period action without knowing the first-period outcome, and $w_S = w_F$ prevents interim pay from revealing it. The complete history SS, SF, FS, or FF remains verifiable and contractible at the end of period 2, so terminal wages may differ across histories.

Let 2^{NF} denote action 2 chosen without observing the first-period outcome. It is the same action and technology as action 2 and has the same effort cost c. The superscript records only the information available when the action is chosen. Because the fallback can be accessed only after a recognized failure, the hidden-failure continuation probability is pinned down by the law of iterated expectations rather than separately assumed, and satisfies $E[p_{2^{\mathrm{NF}}} \mid F, 2] < E[p_2] < p_1$. With feedback, by contrast, the worker recognizes a first-period failure and can access the fallback, whose continuation success probability is $E[p_2 \mid F, 2] = p_1$.

A Exploiter contract without feedback

In Manso (2011), withholding interim outcomes economizes on incentives for exploitation: when effort is costly and $p_0 > 0$, it rules out deviations that condition second-period behavior on the first-period outcome.[5] Here $p_0 = 0$, so shirking never produces success and repeated effort can already be implemented at baseline utility $2\underline{w}$. Withholding feedback therefore delivers no motivational saving: intended output and the minimum expected wage cost are the same either way.

The value of withholding feedback comes from selection. A misreporting explorer who is kept uninformed cannot recognize failure and adapt by accessing the fallback. He must choose one continuation action across both hidden histories, including the low hidden-failure continuation, so the exploiter contract is worth less to him while the intended exploiter is unaffected.

PROPOSITION 8 (Exploiter contract without feedback)*: If* $\overline{U}_{\mathrm{X}} \leq 2\underline{w}$*, an optimal baseline contract is*

$$\tilde{w}_S = \tilde{w}_F = \tilde{w}_{FF} = 0, \qquad \tilde{w}_{SF} = \tilde{w}_{FS} = \frac{c}{p_1}, \qquad \tilde{w}_{SS} = \frac{2c}{p_1}.$$

If $\overline{U}_{\mathrm{X}} > 2\underline{w}$*, an optimal contract delivers the required additional surplus through terminal bonuses for one and two successes, with equal bonuses for the two single-success histories. The split is chosen to minimize the misreporting explorer's best-response utility subject to the exploiter's incentive constraints. At every feasible exploiter utility requirement, withholding feedback leaves the intended exploiter's output and minimum expected wage cost unchanged.*

Online Appendix Figure OA.4 plots the exploiter contract with and without feedback as a function of the utility requirement $\overline{U}_{\mathrm{X}}$. At the baseline, compensation is standard pay for performance: the bonus depends only on the number of successes. Above the baseline, the split of additional rent between single-success histories and double success balances the explorer's temptation to work for one period against the temptation to work in both, while preserving the exploiter's incentives. Figure 3 illustrates the selection effect of withholding feedback.

[5] Similar results appear in Lizzeri et al. (2002) and Fuchs (2007), who show that suppressing performance evaluations can reduce incentive costs in repeated moral-hazard settings.

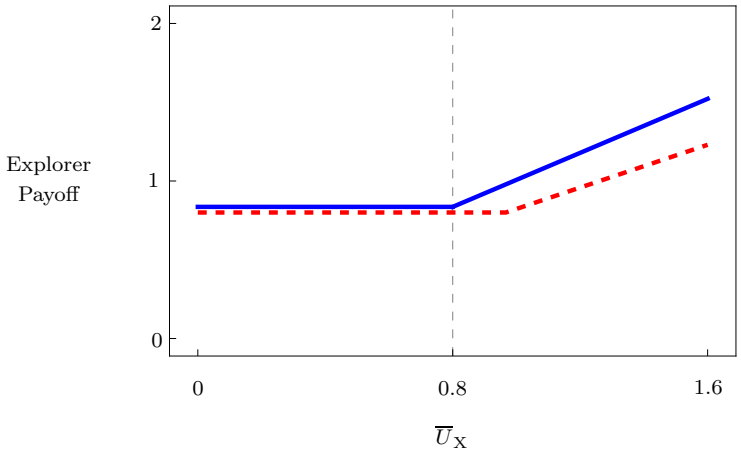


Figure 3. Explorer ride-along payoff under the exploiter contract with feedback (solid blue curve) and without feedback (dashed red curve). At a fixed $\overline{U}_{\mathrm{X}}$, withholding feedback leaves intended output and minimum wage cost unchanged and lowers what an explorer can obtain by mimicking.

B Exploration without feedback

Feedback allows the explorer to continue the original variant after success and access the fallback after failure. An explorer kept uninformed can only repeat the novel method blindly, which is worth less than exploring with feedback.

PROPOSITION 9 (Feedback and productive exploration)*: Suppose the blind-loss condition holds. Every optimal allocation that serves explorers provides them feedback and implements* $\langle {}_{2}{}_{2}^{2} \rangle$*. Blind exploration loses money, and no screening gain from withholding feedback can justify it.*

The screening clause is the substantive one. A blind exploration item loses money, so it could earn its place only by lowering what a mimicking exploiter obtains from it. It cannot. The exploiter's outside option and the explorer's own value on the exploitation item give lower bounds on the utilities the firm must deliver. If the explorer's outside option does not exceed his bound, adaptive exploration with feedback delivering that bound gives the exploiter no more than his own, and the firm replaces the blind item. If it does, the exploitation item already keeps explorers out, and the firm closes the track.

Feedback plays different roles on the two tracks.

COROLLARY 3 (Feedback and selection)*:*

(i) *Under the blind-loss condition, the optimal allocation provides feedback on the exploration track and withholds it on the exploitation track.*

(ii) *When dismissal is also contractible, the wage floor does not exceed the effort cost, and each type's outside option is at least* $2\underline{w}$*, every optimal allocation retains a served exploiter: withholding feedback replaces dismissal on the exploitation track.*

The screening gain from withholding feedback on the exploitation track can lower selection costs or make exploiter-only hiring feasible. When dismissal is also available, the firm has two ways to make the exploitation contract unattractive to explorers: end employment after failure, or hide the outcome. Under the conditions of part (ii), hiding it keeps the productive continuation and never raises the utilities the firm must deliver, so the firm hides it. The dismissal of exploiters in Section V is then the choice of a firm that cannot hide the outcome, as when the worker observes his own output or interim results are public. An exact example in Online Appendix B.4 shows both choices as optima at the same parameters: dismissal when the outcome cannot be hidden, withholding when it can.

VII One Contract or Two?

The analysis so far has characterized the contract offered to each track. A firm that serves both types must also decide how many contracts to post: a separate item for each type, or a single contract that both accept and that elicits type-specific behavior. The natural starting point is a wage-only benchmark, in which nothing but the schedule of bonuses distinguishes the two arrangements.

PROPOSITION 10 (One contract in the wage-only benchmark)*: Suppose both types are served, the firm induces* $\langle {}_{2}{}^{2}_{2} \rangle$ *from explorers and* $\langle {}_{1}{}^{1}_{1} \rangle$ *from exploiters, and the baseline commitment and feedback environment applies without termination.*

(i) *Every feasible two-item menu can be replaced by a single common contract that leaves each type the same utility at the same prior-weighted expected wage cost.*

(ii) *If explorers earn what the exploiter contract would give them, as they do when both outside options are slack, the optimal common contract is the exploiter contract of Proposition 2. If exploiters earn what the explorer contract would give them for exploiting, it is the explorer contract of Proposition 3. Otherwise neither type's*

pay is set by the other's contract, and neither of the two is an optimal common contract.

(iii) *In each case, a small rise in one type's outside option above what he already earns raises only his pay: the explorer's in the first case, the exploiter's in the second, either in the third. A common contract cannot be topped up to raise only one type's pay and remain least-cost, because every bonus it pays reaches both types. The contracts of Propositions 2 and 3 are only ever topped up, one of them or both, since no bonus of either falls as the utility it must deliver rises. The second step of Section II.B therefore selects the two contracts.*

In the first two cases the firm posts one type's contract and the other type takes it as well. Explorers take the exploiter contract, an ordinary piece rate when outside options are slack, or exploiters take the explorer contract. Gottlieb and Moreira (2022) show that, with risk-neutral agents and limited liability, a single contract is optimal under adverse selection and moral hazard when all types rank contracts by incentive power in the same way. Here that ranking fails, yet one schedule elicits two activities because the two types cannot do the same things.

Track-specific contracts earn their keep when a firm hires only one type. When both outside options cover the two wage-floor payments and serving both types leaves one type more than his outside option, a firm can avoid paying that excess by hiring only that type, under his own track's contract. It does so when the total rent saved exceeds the total profit forgone on the other type and the type it keeps is profitable on his own (Online Appendix C.1).

When competition raises only one type's pay, separate contracts can adjust through top-ups, whereas a common contract must be redesigned to remain least-cost. If such changes are possible and redesign is costly, separate contracts are strictly cheaper in expectation (Online Appendix D.3). I next examine the differences that arise without commitment or through different termination and feedback policies.

For these comparisons, fix an implementable allocation serving both types and compare only $\langle {}_{2}{}^{2}_{2} \rangle$ and $\langle {}_{1}{}^{1}_{1} \rangle$, or their terminated versions $\langle {}_{2}{}^{2}_{t} \rangle$ and $\langle {}_{1}{}^{1}_{t} \rangle$. Each departure changes one feature of the benchmark, keeping the others fixed.

PROPOSITION 11 (When separate contracts are valuable)*:*

(i) *Without commitment, a common contract does not reveal the worker's type, so the firm cannot revise continuation terms type by type. If the pooled continuation,*

which keeps both types on their plans, is sequentially optimal, a common contract never costs more than separate contracts without commitment; with slack outside options it costs strictly less, and the same as separate contracts with commitment. Otherwise only separate contracts can implement the prescribed plans.

(ii) *Under the conditions of Corollary 2(iii), separate contracts that retain explorers and terminate exploiters after failure yield strictly higher profit than any common contract at the maintained plans.*

(iii) *Providing feedback on the exploration contract and withholding it on the exploitation contract weakly lowers minimum wage cost relative to a common contract. The reduction is strict if and only if the least-cost common contract leaves explorers strictly more than their participation requirement.*

Pooling itself supplies a form of commitment. Acceptance of a common contract reveals nothing about the worker's type, so the principal cannot cut an explorer's later pay without losing the exploiters' effort. The logic is that of the ratchet effect, in which a principal who cannot commit uses what the agent reveals and pooling results (Laffont and Tirole, 1988); here the firm pools by design, with one contract that elicits two activities. When the pooled continuation is sequentially optimal, a common contract can therefore sustain a continuation rent for successful explorers, which Section IV shows a separate exploration contract cannot do without commitment. Pooling has a limit. When explorers are sufficiently prevalent after a first-period success, the principal would rather induce only exploration than keep the incentives for repeated exploitation; then only separate contracts, whose acceptance identifies the track, sustain the fixed plans.

Separate contracts allow the firm to treat failure differently across tracks. A common contract must apply the same termination rule to both types after failure. Under the conditions of Corollary 2(iii), retaining both forgoes the selection gain on the exploitation track; terminating both sacrifices valuable continuation on the exploration track. Separate contracts avoid both losses.

A common contract that implements both productive plans must provide feedback to both types. Separate contracts can withhold it from exploiters and so reduce explorer rents when screening forces their utility above the participation requirement. When both outside options are slack, withholding feedback from exploiters eliminates the explorer's rent altogether: both types receive $2\underline{w}$.

VIII Empirical implications

Innovation incentives can raise measured innovative output in two ways. They can change what incumbent workers do, or they can change which workers do exploratory work. The second margin is an outcome of the model and not merely a nuisance for identification, and testing it requires data on who works under an arrangement rather than only on what is produced under it. The two predictions below concern that margin and the compensation and practices that support it.

Online Appendix E establishes the conditions underlying these predictions. Competition for talent refers to changes in explorers' or exploiters' outside opportunities.

Prediction 1. *Long-run rewards and commitment help organizations select explorers, while feedback enables adaptive exploration; when promised pay is high and early failure is sufficiently diagnostic, explorers receive pay after failure and can be retained when exploiters would be dismissed. Organizations seeking exploiters reward success through a contract that needs no commitment with feedback and retention fixed; withholding feedback or dismissal after failure can discourage explorers from taking those jobs. In a separate menu serving both types, workers choose the arrangement designed for their comparative advantage.*

The prediction applies wherever exploration- and exploitation-oriented arrangements coexist: career tracks, funding programs, or units within one organization, or employers that specialize in one kind of work. Around an exploration-oriented change, the observables are applications and acceptances by arrangement, moves between arrangements by workers already inside, entry from outside, and retention.

Tests should distinguish sorting between arrangements from selective hiring. When both types are served under separate arrangements, examine which workers choose each. When an employer hires only one type, examine who enters, remains, or stays outside. The central question is whether workers with stronger pre-treatment exploratory comparative advantage become disproportionately more likely to enter and remain in exploratory work.

Everything therefore depends on classifying workers before the change. The relevant type is comparative advantage, so the informative measure is the same worker's relative performance in exploratory and exploitative tasks: how a scientist's novel lines of work compare with routine ones, how an engineer's speculative projects compare with incremental ones. Generic productivity is not the type, since a worker can be better at both activities without being relatively better at exploration. Nor are risk preferences,

wealth, or confidence, although they may affect which arrangement a worker chooses. Post-treatment outcomes cannot serve as measures of type, because they are the response the design is meant to identify.

Separating composition from behavior then requires following people rather than output alone. Worker fixed effects do not identify changes in workforce composition. A design should also follow entrants, leavers, and switchers using their pre-treatment characteristics and compare them with incumbents as hiring and turnover accumulate.

Dismissal after early failure provides a further test. With both tracks active in the severance region, explorers require weakly less continuation value to be retained after failure than exploiters, because dismissing an exploiter saves screening cost and dismissing an explorer does not. Where firms can hide outcomes, withholding feedback can replace dismissal of exploiters. Tests should compare dismissal conditional on early failure, accounting for continuation opportunities and promised compensation, and examine worker composition and arrangement choices alongside dismissal decisions.

Prediction 2. *When competition for talent raises required pay, it changes the shape of compensation, not merely its level: rather than a uniform raise across all histories, the additional pay lands on different histories on the two tracks, success bonuses for exploiters and, under commitment, long-run rewards for explorers. At sufficiently high pay levels, and when early failure is sufficiently diagnostic of exploration, explorer compensation expands to include pay after early failure.*

Motivation alone does not determine where additional pay goes: with a single worker type, any placement that preserves incentives costs the same. Screening does, because how much pay leaks to the other type depends on the history. The measurable counterparts are contractual: vesting and payout horizons, the share of pay contingent on long-run rather than immediate results, severance and continuation payments after early failure, and the size of contemporaneous success bonuses. A design that measures only the level of compensation, or a single pay-for-performance slope, will not detect the prediction.

Changes in outside opportunities provide variation for testing this prediction. The entry or departure of an employer competing for exploratory talent and the opening or closing of a research campus can raise required pay on one track and, through self-selection, on the other. Type-specific shocks are informative but not necessary. The prediction concerns the track-specific contracts characterized above and applies where the shock raises the relevant utility requirement, holding intended activities, the types served, and organizational practices fixed. An observed outside option maps one-for-one into that requirement only when participation binds for the type in question. When

both types are served, self-selection can instead set the requirement. The comparison to run is across tracks in the same organization or market, and it must distinguish changes in outside options from changes in project technology or the returns to exploration.

Existing evidence and empirical implementation.

Existing research studies investment, innovation, and entrepreneurship in relation to incentive horizons and governance (Laux, 2012; Bernstein, 2015; He and Tian, 2013; Ladika and Sautner, 2020; Chemmanur and Tian, 2018; Balsmeier et al., 2017; Lerner and Wulf, 2007), the consequences of failure (Acharya and Subramanian, 2009; Mann, 2018; Acharya et al., 2014; Hombert et al., 2020; Tian and Wang, 2014; Gottlieb et al., 2022), and evaluation and feedback (Gross, 2017; Chu et al., 2019; Howell, 2021). These studies examine outcomes: innovative output, investment, entry, or the abandonment of ventures. Where they follow people, as Bernstein (2015) does for inventors around an IPO, they do not observe comparative advantage in exploration before treatment.

The HHMI program studied by Azoulay et al. (2011) bundles a long horizon, tolerance of early failure, freedom to redirect research, feedback, and evaluation. Its effects on scientific output leave open whether it also attracts and retains scientists with stronger pre-appointment exploratory comparative advantage. Data on applicants and nominees, take-up where observable, and retention, linked to pre-appointment measures of comparative advantage, would make the selection margin testable.

An experiment could distinguish selection from motivation. First measure subjects' comparative advantage in exploratory and exploitative tasks before treatment. In a choice arm, offer exploration- and exploitation-oriented schemes that differ in reward timing and treatment of early failure. Test whether pre-treatment exploratory advantage predicts choosing the exploration-oriented scheme. In a random-assignment arm, assign the same schemes and observe subsequent behavior to identify their average effect on behavior, the margin studied by Manso (2011).

IX Conclusion

Long-run rewards, failure tolerance, commitment, job security, and feedback motivate a given worker to experiment. Here they also determine who experiments. With type and effort hidden, explorers and exploiters generate different performance histories under their own plans and each other's contracts. A contract that must motivate its intended activity can screen by where its rents sit: immediate success for exploiters, long-run success for explorers, and early failure only when it is sufficiently diagnostic.

Commitment limits an exploiter's gains from imitating an explorer, while feedback preserves exploration's option value. Withholding feedback from exploiters improves selection without changing their output or minimum wage cost. A firm that cannot hide outcomes may retain an explorer and dismiss an exploiter after the same failure, despite the same continuation opportunity. A common contract can help overcome a lack of commitment: acceptance does not reveal the worker's type, limiting the firm's ability to revise continuation terms separately for each type.

Evidence that exploration-oriented institutions affect innovative output does not by itself identify selection. Separating motivation from selection requires following workers as well as output: who enters, who stays, and who chooses exploratory work. And because competition for talent places additional pay on different histories on the two tracks, its effect on compensation shows up in where pay sits, not only in its level.

References

Acharya, Viral and Krishnamurthy Subramanian, "Bankruptcy Codes and Innovation," *Review of Financial Studies*, 2009, *22* (12), 4949–4988.

_ , Ramin Baghai, and Krishnamurthy Subramanian, "Wrongful Discharge Laws and Innovation," *Review of Financial Studies*, 2014, *27* (1), 301–346.

Aghion, Philippe and Jean Tirole, "The Management of Innovation," *Quarterly Journal of Economics*, 1994, *109* (4), 1185–1209.

Azoulay, Pierre, Joshua Graff Zivin, and Gustavo Manso, "Incentives and Creativity: Evidence from the Academic Life Sciences," *RAND Journal of Economics*, 2011, *42* (3), 527–554.

Balsmeier, Benjamin, Lee Fleming, and Gustavo Manso, "Independent Boards and Innovation," *Journal of Financial Economics*, 2017, *123* (3), 536–557.

Bénabou, Roland and Jean Tirole, "Bonus Culture: Competitive Pay, Screening, and Multitasking," *Journal of Political Economy*, 2016, *124* (2), 305–370.

Bergemann, Dirk and Ulrich Hege, "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, 2005, *36* (4), 719–752.

Bernstein, Shai, "Does Going Public Affect Innovation?," *Journal of Finance*, 2015, *70* (4), 1365–1403.

Bolton, Patrick and Christopher Harris, "Strategic Experimentation," *Econometrica*, 1999, *67* (2), 349–374.

Castro-Pires, Henrique and George Georgiadis, "Incentives and Selection," Working Paper, Northwestern University 2026.

Che, Yeon-Koo, Elisabetta Iossa, and Patrick Rey, "Prizes versus Contracts as Incentives for Innovation," *Review of Economic Studies*, 2021, *88* (5), 2149–2178.

Chemmanur, Thomas J. and Xuan Tian, "Do Antitakeover Provisions Spur Corporate Innovation? A Regression Discontinuity Analysis," *Journal of Financial and Quantitative Analysis*, 2018, *53* (3), 1163–1194.

Chu, Yongqiang, Xuan Tian, and Wenyu Wang, "Corporate Innovation Along the Supply Chain," *Management Science*, 2019, *65* (6), 2445–2466.

Dohmen, Thomas and Armin Falk, "Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender," *American Economic Review*, 2011, *101* (2), 556–590.

Ederer, Florian and Gustavo Manso, "Is Pay for Performance Detrimental to Innovation?," *Management Science*, 2013, *59* (7), 1496–1513.

Foarta, Dana and Takuo Sugaya, "The Management of Talent: Optimal Contracting for Selection and Incentives," *RAND Journal of Economics*, 2021, *52* (1), 49–77.

Fuchs, William, "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review*, 2007, *97* (4), 1432–1448.

Gomes, Renato, Daniel Gottlieb, and Lucas Maestri, "Experimentation and Project Selection: Screening and Learning," *Games and Economic Behavior*, 2016, *96*, 145–169.

Gottlieb, Daniel and Humberto Moreira, "Simple Contracts with Adverse Selection and Moral Hazard," *Theoretical Economics*, 2022, *17* (3), 1357–1401.

Gottlieb, Joshua D., Richard R. Townsend, and Ting Xu, "Does Career Risk Deter Potential Entrepreneurs?," *Review of Financial Studies*, 2022, *35* (9), 3973–4015.

Gross, Daniel, "Performance Feedback in Competitive Product Development," *RAND Journal of Economics*, 2017, *48* (2), 438–466.

Halac, Marina, Navin Kartik, and Qingmin Liu, "Optimal Contracts for Experimentation," *Review of Economic Studies*, 2016, *83* (3), 1040–1091.

_ , _ , and _ , "Contests for Experimentation," *Journal of Political Economy*, 2017, *125* (5), 1523–1569.

He, Jie and Xuan Tian, "The Dark Side of Analyst Coverage: The Case of Innovation," *Journal of Financial Economics*, 2013, *109* (3), 856–878.

Hellmann, Thomas, "When Do Employees Become Entrepreneurs?," *Management Science*, 2007, *53* (6), 919–933.

_ and Veikko Thiele, "Incentives and Innovation: A Multi-tasking Approach," *American Economic Journal: Microeconomics*, 2011, *3* (1), 78–128.

Holmstrom, Bengt, "Agency Costs and Innovation," *Journal of Economic Behavior & Organization*, 1989, *12* (3), 305–327.

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar, "Can Unemployment Insurance Spur Entrepreneurial Activity? Evidence from France," *Journal of Finance*, 2020, *75* (3), 1247–1285.

Hörner, Johannes and Larry Samuelson, "Incentives for Experimenting Agents," *RAND Journal of Economics*, 2013, *44* (4), 632–663.

Howell, Sabrina, "Learning from Feedback: Evidence from New Ventures," *Review of Finance*, 2021, *25* (3), 595–627.

Keller, Godfrey, Sven Rady, and Martin Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 2005, *73* (1), 39–68.

Khalil, Fahad, Jacques Lawarrée, and Alexander Rodivilov, "Learning from Failures: Optimal Contracts for Experimentation and Production," *Journal of Economic Theory*, 2020, *190*, 105107.

_ , _ , and _ , "Rewarding Success and Failure: Moral Hazard and Adverse Selection in Strategic Experimentation," *Games and Economic Behavior*, 2026, *159*, 241–267.

Ladika, Tomislav and Zacharias Sautner, "Managerial Short-Termism and Investment: Evidence from Accelerated Option Vesting," *Review of Finance*, 2020, *24* (2), 305–344.

Laffont, Jean-Jacques and Jean Tirole, "The Dynamics of Incentive Contracts," *Econometrica*, 1988, *56* (5), 1153–1175.

Laux, Volker, "Stock Option Vesting Conditions, CEO Turnover, and Myopic Investment," *Journal of Financial Economics*, 2012, *106* (3), 513–526.

Lazear, Edward P., "Salaries and Piece Rates," *Journal of Business*, 1986, *59* (3), 405–431.

_ , "Performance Pay and Productivity," *American Economic Review*, 2000, *90* (5), 1346–1361.

Leaver, Clare, Owen Ozier, Pieter Serneels, and Andrew Zeitlin, "Recruitment, Effort, and Retention Effects of Performance Contracts for Civil Servants: Experimental Evidence from Rwandan Primary Schools," *American Economic Review*, 2021, *111* (7), 2213–2246.

Lerner, Josh and Julie Wulf, "Innovation and Incentives: Evidence from Corporate R&D," *Review of Economics and Statistics*, 2007, *89* (4), 634–644.

Lizzeri, Alessandro, Margaret Meyer, and Nicola Persico, "The Incentive Effects of Interim Performance Evaluations," CARESS Working Paper 02-09, University of Pennsylvania 2002.

Mann, William, "Creditor Rights and Innovation: Evidence from Patent Collateral," *Journal of Financial Economics*, 2018, *130* (1), 25–47.

Manso, Gustavo, "Motivating Innovation," *Journal of Finance*, 2011, *66* (5), 1823–1860.

March, James G., "Exploration and Exploitation in Organizational Learning," *Organization Science*, 1991, *2* (1), 71–87.

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

Moscarini, Giuseppe and Lones Smith, "The Optimal Level of Experimentation," *Econometrica*, 2001, *69* (6), 1629–1644.

O'Reilly, Charles A. and Michael L. Tushman, "Organizational Ambidexterity: Past, Present, and Future," *Academy of Management Perspectives*, 2013, *27* (4), 324–338.

Oyer, Paul and Scott Schaefer, "Why Do Some Firms Give Stock Options to All Employees? An Empirical Examination of Alternative Theories," *Journal of Financial Economics*, 2005, *76* (1), 99–133.

Roberts, Kevin and Martin L. Weitzman, "Funding Criteria for Research, Development, and Exploration Projects," *Econometrica*, 1981, *49* (5), 1261–1288.

Salop, Joanne and Steven C. Salop, "Self-Selection and Turnover in the Labor Market," *Quarterly Journal of Economics*, 1976, *90* (4), 619–627.

Tian, Xuan and Tracy Yue Wang, "Tolerance for Failure and Corporate Innovation," *Review of Financial Studies*, 2014, *27* (1), 211–255.

Tushman, Michael L. and Charles A. O'Reilly III, "Ambidextrous Organizations: Managing Evolutionary and Revolutionary Change," *California Management Review*, 1996, *38* (4), 8–30.

Weitzman, Martin L., "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

A Service Sets and Utility Requirements

For a contractible item z, let $V_\theta(z) = \max_{a \in \mathcal{A}(\theta)} U_\theta(z, a)$. The null item gives $V_\theta(\emptyset) = \overline{U}_\theta^{\text{com}}$ and zero profit. For $\theta \in \{\text{E}, \text{X}\}$, fix the intended plan a_θ and policies (ϕ_θ, D_θ), and write $z_\theta = (\tilde{w}, \phi_\theta, D_\theta)$. The minimization is over the bonus vector. With $-\theta$ denoting the other type, define the conditional wage cost

$$\begin{aligned}
&\mathcal{K}_\theta(\overline{U}_\theta, \overline{U}_{-\theta}; a_\theta, \phi_\theta, D_\theta) = \min_{\substack{\tilde{w} \geq 0 \\ z_\theta \in \mathcal{Z}}} W(z_\theta, a_\theta) \\
&\text{s.t.} \quad U_\theta(z_\theta, a_\theta) \geq U_\theta(z_\theta, a) \quad \text{for all } a \in \mathcal{A}(\theta), \\
&\qquad U_\theta(z_\theta, a_\theta) \geq \overline{U}_\theta, \qquad V_{-\theta}(z_\theta) \leq \overline{U}_{-\theta}.
\end{aligned}$$

An infeasible inner problem has value $+\infty$. The outer problem is

$$\begin{aligned}
&\mathcal{K}(a_\text{E}, a_\text{X}; \phi_\text{E}, D_\text{E}, \phi_\text{X}, D_\text{X}) \\
&= \min_{\substack{\overline{U}_\theta \geq \overline{U}_\theta^{\text{com}} \\ \theta \in \{\text{E},\text{X}\}}} \Big\{ \mu \mathcal{K}_\text{E}(\overline{U}_\text{E}, \overline{U}_\text{X}; a_\text{E}, \phi_\text{E}, D_\text{E}) \\
&\qquad\qquad + (1-\mu) \mathcal{K}_\text{X}(\overline{U}_\text{X}, \overline{U}_\text{E}; a_\text{X}, \phi_\text{X}, D_\text{X}) \Big\}.
\end{aligned}$$

Each requirement bounds the same type's utility from below on its own item and from above on the other item. It need not equal delivered utility.

LEMMA 1 (Exact Two-Utility-Requirement Decomposition)*: Fix target plans* (a_E, a_X) *and each item's feedback and dismissal policies. The minimum prior-weighted expected wage cost of serving both types and implementing these plans equals*

$$\mathcal{K}(a_\text{E}, a_\text{X}; \phi_\text{E}, D_\text{E}, \phi_\text{X}, D_\text{X}).$$

A menu attains this minimum if and only if each designated item minimizes its inner problem at an outer-optimal requirement pair. The two items may coincide at a common contract.

Proof. For a feasible menu (z_E, z_X), set $\overline{U}_\theta = \max\{\overline{U}_\theta^{\text{com}}, V_\theta(z_{-\theta})\}$. Participation and self-selection give $V_\theta(z_{-\theta}) \leq \overline{U}_\theta \leq V_\theta(z_\theta)$. Each item is therefore feasible for its inner problem, so the menu's cost is at least the outer minimum. The minima are attained whenever the problems are feasible: each inner problem, and the joint problem in the two bonus vectors and the two requirements, is a linear program whose objective is bounded below. Conversely, inner minimizers at an outer-optimal pair satisfy within-

type incentives, participation, and these same self-selection inequalities. If the items coincide at z, the lower bound from type θ's inner problem and the upper bound from the other problem imply $V_\theta(z) = \overline{U}_\theta$ for both types. The common item implements both target plans and has exactly the prior-weighted inner cost. Finally, for a cost-minimizing menu, the first construction must attain equality in both the inner and outer cost bounds. Since $\mu \in (0,1)$, each item must minimize its inner cost. This also proves the converse optimizer correspondence. □

Without commitment, continuation terms are chosen after acceptance against the belief it induces, so the decomposition applies to items that identify the track. A common item's pooled continuation is treated directly in Online Appendix C.2.1.

Requirements can be chosen minimally as in the proof. Conditional on serving only type θ, participation and exclusion instead fix them at $(\overline{U}_\theta^{\text{com}}, \overline{U}_{-\theta}^{\text{com}})$. Writing $m_{\text{E}} = \mu$ and $m_{\text{X}} = 1 - \mu$, the service-set values below maximize over feasible plans and admissible feedback and dismissal policies. In the revenue terms, z_θ is any item with policies (ϕ_θ, D_θ).

$$\Pi^*_{\{\theta\}} = \max_{a_\theta, \phi_\theta, D_\theta} m_\theta \Big[R(z_\theta, a_\theta) - \mathcal{K}_\theta(\overline{U}_\theta^{\text{com}}, \overline{U}_{-\theta}^{\text{com}}; a_\theta, \phi_\theta, D_\theta) \Big].$$

$$\begin{aligned} &\Pi^*_{\{\text{E,X}\}} \\ &= \max_{\substack{a_{\text{E}}, a_{\text{X}} \\ \phi_{\text{E}}, D_{\text{E}}, \phi_{\text{X}}, D_{\text{X}}}} \Big\{ \mu R(z_{\text{E}}, a_{\text{E}}) + (1-\mu) R(z_{\text{X}}, a_{\text{X}}) \\ &\qquad - \mathcal{K}(a_{\text{E}}, a_{\text{X}}; \phi_{\text{E}}, D_{\text{E}}, \phi_{\text{X}}, D_{\text{X}}) \Big\}, \\ &\Pi^* = \max\{0, \Pi^*_{\{\text{E}\}}, \Pi^*_{\{\text{X}\}}, \Pi^*_{\{\text{E,X}\}}\}. \end{aligned}$$

An infeasible plan contributes $-\infty$ to profit; the null allocation ensures $\Pi^* \geq 0$. These formulas follow by applying the lemma to both-type service and the outside-option bounds to singleton service. Lemma 4 in Appendix C.4 gives conditions under which every allocation that attains Π^* uses the maintained plans. The conditional exclusion result and screening-value decomposition are in Online Appendix A.

B Proofs of Optimal Screening Contracts

This appendix states when the utility requirements are feasible and proves Propositions 2 and 3.

B.1 Feasibility of Utility Requirements

Write $x^+ = \max\{x, 0\}$. At the maintained target plans, the minimum explorer best-response value on an exploitation item has baseline

$$V_{\mathrm{E}}^{\mathrm{X},0} = 2\underline{w} + \frac{c}{p_1}\left[E[p_2](1 + E[p_2 \mid S, 2] - p_1) - p_1\right].$$

The conditional problem in Proposition 2 is feasible exactly when

$$V_{\mathrm{E}}^{\mathrm{X},0} + \frac{E[p_2]}{p_1}(\overline{U}_{\mathrm{X}} - 2\underline{w})^+ \leq \overline{U}_{\mathrm{E}}. \tag{B.1}$$

For exploration, define the requirement at which the mimicking exploiter's first-period incentive constraint begins to bind:

$$\overline{U}_N^* = 2\underline{w} + c\left[\frac{(1 + p_1)E[p_2]E[p_2 \mid S, 2]}{p_1^2} - (1 + E[p_2])\right].$$

The conditional problem in Proposition 3 is feasible exactly when

$$\begin{aligned} 2\underline{w} + \min\left\{\frac{p_1^2}{E[p_2]E[p_2 \mid S, 2]}, \frac{p_1}{E[p_2]E[p_2 \mid S, 2] + (1 - E[p_2])p_1}\right\} \\ \times (\overline{U}_{\mathrm{E}} - \overline{U}_N^*)^+ \leq \overline{U}_{\mathrm{X}}. \end{aligned} \tag{B.2}$$

The following proofs establish both conditions globally over implementing contracts and the other type's full action-plan set. Both-type service needs both conditions. For selective hiring, substitute the intended and excluded types' outside options for the two requirements; feasibility is not a claim that singleton service is optimal. The corresponding conditions for the organizational extensions are collected in Online Appendix A.3.

B.2 Proof of the Exploiter Contract

Proof of Proposition 2. Every contract implementing $\langle {}_1{}^1_1 \rangle$ can be written using continuation slacks $x, y \geq 0$ and first-period slack $z \geq 0$ as

$$\begin{aligned} \tilde{w}_{SS} &= \tilde{w}_{SF} + c/p_1 + x, \qquad \tilde{w}_{FS} = \tilde{w}_{FF} + c/p_1 + y, \\ \tilde{w}_S &= c/p_1 + \tilde{w}_F - \tilde{w}_{SF} + \tilde{w}_{FF} - p_1 x + p_1 y + z, \end{aligned}$$

subject to nonnegative bonuses. The exploiter's expected bonus is

$$B_{\mathrm{X}} = 2c + \tilde{w}_F + \tilde{w}_{FF} + p_1 y + p_1 z.$$

Hence $B_{\mathrm{X}} \geq 2c$, with equality requiring $\tilde{w}_F = \tilde{w}_{FF} = y = z = 0$. Relative to the baseline $\tilde{w}_S = \tilde{w}_{SS} = \tilde{w}_{FS} = c/p_1$ and all other bonuses zero, the change in an explorer's payoff from $\left\langle {}_2{}^2_2 \right\rangle$ is

$$\tilde{w}_F + \tilde{w}_{FF} + E[p_2](E[p_2 \mid S, 2] - p_1)x + p_1 y + E[p_2]z.$$

All coefficients are nonnegative, and the coefficient of x is positive.

At this baseline the explorer works after success because $E[p_2 \mid S, 2] > p_1$; after a genuine exploratory failure, the fallback and shirking give the same continuation utility. First-period shirking creates neither information nor the fallback, so subsequent exploration pays $E[p_2]c/p_1 - c < 0$ and the explorer shirks again. Thus the relevant first-period comparison is between exploration and full shirking. Its gain is

$$\frac{c}{p_1}\left[E[p_2](1 + E[p_2 \mid S, 2] - p_1) - p_1\right] > 0$$

by the count ordering. The target exploratory plan therefore attains the baseline best-response value $V_{\mathrm{E}}^{\mathrm{X},0}$. Since its payoff cannot be lower under any exploiter-implementing contract, this is a global lower bound on explorer utility, including under contracts with higher intended cost.

For surplus delivery, exploiter utility above $2\underline{w}$ is $\Delta_{\mathrm{X}} = \tilde{w}_F + \tilde{w}_{FF} + p_1 y + p_1 z$. Transfers through $\tilde{w}_F$, $\tilde{w}_{FF}$, and y raise the explorer's payoff one-for-one with exploiter utility; x raises only the explorer's payoff. The z direction instead has rate $E[p_2]/p_1 < 1$. Thus every implementing contract meeting the exploiter requirement gives the explorer at least

$$V_{\mathrm{E}}^{\mathrm{X},0} + \frac{E[p_2]}{p_1}(\overline{U}_{\mathrm{X}} - 2\underline{w})^+.$$

This bound is attained by setting $z = (\overline{U}_{\mathrm{X}} - 2\underline{w})^+/p_1$ and all other slacks to zero. Only $\tilde{w}_S$ increases. The explorer still explores in period 1, since this increase raises his exploration payoff and leaves his shirking payoff unchanged. The cost- and value-neutral $\tilde{w}_{SF}$ direction delays payment; paying earliest selects $\tilde{w}_{SF} = 0$.

The constructed contract has the minimum intended cost and attains the global minimum explorer best-response value. It is therefore conditionally optimal exactly under (B.1); otherwise the requirement pair is infeasible. Evaluating the same attained minimum at $\overline{U}_{\mathrm{X}}^{\mathrm{com}}$ and comparing it with $\overline{U}_{\mathrm{E}}^{\mathrm{com}}$ establishes the condition for hiring only exploiters. The exclusion inequality is weak under the paper's menu-choice convention. □

B.3 Proof of the Explorer Contract

Proof of Proposition 3. Set $\tilde{w}_S = \tilde{w}_F = \tilde{w}_{SF} = \tilde{w}_{FF} = 0$ and $\tilde{w}_{FS} = c/p_1$. After genuine exploratory failure, the fallback success probability is p_1, so continued effort and shirking are tied. After first-period shirking, however, the success probability remains $E[p_2] < p_1$; continued exploration is strictly unprofitable. Full shirking is consequently the relevant first-period deviation. The exploratory target is optimal when

$$E[p_2]E[p_2 \mid S, 2]\tilde{w}_{SS} + (1 - E[p_2])p_1\tilde{w}_{FS} \geq 2c.$$

At equality this gives

$$b_N = \frac{(1 + E[p_2])c}{E[p_2]E[p_2 \mid S, 2]}, \qquad \tilde{w}_{SS} = b_N.$$

Continuation effort after success is strictly optimal, after genuine failure it is weakly optimal, and after shirking it is strictly dominated. These comparisons cover all eight explorer plans. Expected bonuses equal $2c$. No implementing contract can pay less, since full shirking guarantees at least $2\underline{w}$ under every wage-floor contract. Implementation is weak; no additional action-choice tie-break is needed.

Hold explorer utility fixed and move pay from SS to S. With $d\tilde{w}_{SS} = -d\tilde{w}_S/E[p_2 \mid S, 2]$, the working exploiter's payoff rises by

$$p_1(1 - p_1/E[p_2 \mid S, 2])d\tilde{w}_S > 0.$$

Backloading therefore weakly improves screening. Because pay at S raises the working exploiter's payoff, any contract with $\tilde{w}_S > 0$ stops being optimal before the explorer's requirement reaches $\overline{U}^*_N$; the second step of Section II.B therefore selects $\tilde{w}_S = 0$. At the selected baseline, the exploiter works after success because $p_1 b_N > c$, and is indifferent after failure. His first-period gain from effort rather than shirking is $p_1^2 b_N - (1+p_1)c < 0$. To verify the sign,

$$\begin{aligned}&(1 + p_1)E[p_2]E[p_2 \mid S, 2] - p_1^2(1 + E[p_2]) \\ &\quad = p_1 - E[p_2] + (1 + p_1)\left[E[p_2](1 + E[p_2 \mid S, 2] - p_1) - p_1\right] > 0.\end{aligned}$$

Thus his best-response value is $2\underline{w}$. This is globally minimal because full shirking guarantees that floor under every contract.

Additional explorer utility is first delivered through $\tilde{w}_{SS}$ alone. The exploiter continues to receive $2\underline{w}$ until

$$\tilde{w}_{SS} = b_N^* = \frac{(1+p_1)c}{p_1^2}, \qquad \overline{U}_{\mathrm{E}} = \overline{U}_N^* = 2\underline{w} + E[p_2]E[p_2 \mid S,2](b_N^* - b_N).$$

Beyond this point, compare an increase in $\tilde{w}_{SS}$ alone with the joint adjustment $d\tilde{w}_F = p_1 d\tilde{w}_{SS}$, which preserves the exploiter's first-period indifference. Their respective rates of increase in exploiter utility per unit of explorer utility are

$$\frac{p_1^2}{E[p_2]E[p_2 \mid S,2]} \quad \text{and} \quad \frac{p_1}{E[p_2]E[p_2 \mid S,2] + (1 - E[p_2])p_1}.$$

The joint direction is better exactly when

$$\frac{E[p_2]E[p_2 \mid S,2]}{p_1^2} < \frac{1 - E[p_2]}{1 - p_1}.$$

It raises the explorer's first-period incentive gain by

$$E[p_2](E[p_2 \mid S,2] - p_1) d\tilde{w}_{SS} > 0.$$

When double success and first-period failure are equally diagnostic, the two directions raise the exploiter's utility at the same rate, so every mix of them is optimal and remains optimal as the requirement rises; paying earliest selects the joint adjustment in the proposition.

To establish global optimality, write every contract implementing $\langle {}_2{}^2_2 \rangle$ using continuation slacks $x, y \geq 0$ as

$$\tilde{w}_{SS} = \tilde{w}_{SF} + c/E[p_2 \mid S,2] + x, \qquad \tilde{w}_{FS} = \tilde{w}_{FF} + c/p_1 + y,$$

subject to nonnegative bonuses, and put $k = \tilde{w}_S + \tilde{w}_{SF}$ and $f = \tilde{w}_F + \tilde{w}_{FF} + p_1 y$. The explorer's utility is $2\underline{w} + u$, where

$$u = E[p_2]k + E[p_2]E[p_2 \mid S,2]x + (1 - E[p_2])f - c.$$

An exploiter who shirks in period 1 and works in period 2 obtains $2\underline{w} + f$, and one who works in both periods obtains $2\underline{w} + p_1 k + p_1^2(x + c/E[p_2 \mid S,2]) + (1-p_1)f - (1+p_1)c$. Write $\ell_{SS} = p_1^2/(E[p_2]E[p_2 \mid S,2])$ and $u^* = \overline{U}_N^* - 2\underline{w}$. Eliminating x, the exploiter's best-response value under any implementing contract is at least

$$2\underline{w} + \max\Big\{ f,\ \ell_{SS}(u - u^*) + \big[1 - p_1 - \ell_{SS}(1 - E[p_2])\big] f + p_1\big(1 - p_1/E[p_2 \mid S,2]\big) k \Big\}.$$

The coefficient of k is positive, and the coefficient of f in the second term is nonnegative exactly when $E[p_2]E[p_2 \mid S,2]/p_1^2 \geq (1-E[p_2])/(1-p_1)$. In that case the bound is smallest at $f=k=0$, where it equals $2\underline{w}+\ell_{SS}(u-u^*)^+$. Otherwise the second term falls in f: the bound is smallest at $k=0$ and, for $u \leq u^*$, at $f=0$, where it equals $2\underline{w}$, and for $u > u^*$ where the two terms are equal,

$$f = \frac{p_1(u-u^*)}{E[p_2]E[p_2 \mid S,2] + (1-E[p_2])p_1}.$$

In both cases the smallest value is the left side of (B.2) at $\overline{U}_{\mathrm{E}} = 2\underline{w}+u$.

The contracts constructed above attain it. They have $k=y=0$, with $f=0$ when $\tilde{w}_{SS}$ rises alone and $f=\tilde{w}_F$ on the joint path, and under them the exploiter's best-response value is exactly

$$2\underline{w} + \max\left\{\tilde{w}_F, p_1^2\tilde{w}_{SS} + (1-p_1)\tilde{w}_F - (1+p_1)c\right\},$$

the two terms corresponding to shirking and working in period 1 with continuation chosen optimally. The slope comparison therefore gives the global lower envelope, over all implementing contracts and all exploiter plans, rather than only a local adjustment.

The bound also gives the selections of Section II.B. A contract with $k>0$ leaves the exploiter more than the minimum at every $u \geq u^*$, and topping it up keeps $k>0$, so it stops being optimal before the explorer's requirement reaches $\overline{U}_N^*$: the second step selects $\tilde{w}_S = \tilde{w}_{SF} = 0$. Where the minimum is $2\underline{w}$, attaining it requires $f=0$, that is, $\tilde{w}_F = \tilde{w}_{FF} = 0$ and $\tilde{w}_{FS} = c/p_1$. Past $\overline{U}_N^*$, if double success is more diagnostic the minimum is attained only at $f=k=0$. If first-period failure is at least as diagnostic, the contracts that attain it have $k=0$ and f no larger than the value displayed above, and their expected interim payment $(1-E[p_2])\tilde{w}_F$ is at most $(1-E[p_2])f$; it is largest on the joint path, where $\tilde{w}_F = f$ and $\tilde{w}_{FF} = y = 0$, which the third step selects.

The minimum exploiter utility is exactly the left side of (B.2). The construction delivers the required explorer utility at the minimum expected wage cost, so that inequality is both necessary and sufficient for conditional feasibility. At the two outside options the same attained envelope establishes the condition for hiring only explorers, again with weak exclusion. □

C Arguments for the Organizational Results

This appendix gives the central arguments for the organizational results. Complete optimality proofs and supporting cases are in Online Appendix B, and the feasibility bounds in Online Appendix A.3.

C.1 Lack of Commitment

For exploitation, the sequentially optimal continuation bonus is c/p_1 after either first-period history. Repeating this bonus in period 1 implements the committed baseline. Additional exploiter surplus is paid through $\tilde{w}_S$ and is already settled before the continuation terms are set. Thus Proposition 4 leaves the contract, its cost, and its explorer best-response value unchanged.

For exploration, the least-cost continuation bonuses are instead $\tilde{w}_{SS} = c/E[p_2 \mid S, 2]$ and $\tilde{w}_{FS} = c/p_1$, with $\tilde{w}_{SF} = \tilde{w}_{FF} = 0$. Neither genuine first-period history leaves an explorer continuation rent above the floor. A first-period shirker generates no fallback and then shirks again, since $E[p_2]c/p_1 < c$. First-period incentives therefore require

$$E[p_2](\tilde{w}_S - \tilde{w}_F) \geq c.$$

The baseline sets $\tilde{w}_S = c/E[p_2]$ and $\tilde{w}_F = 0$ and costs $2c$ in bonuses. A mimicking exploiter shirks after success because $p_1 c/E[p_2 \mid S, 2] < c$, and is indifferent after failure. His first-period effort gain is at least $p_1 c/E[p_2] - c > 0$, so he strictly works in period 1 under every feasible exploration contract without commitment.

Let $\Delta = (\overline{U}_{\mathrm{E}} - 2\underline{w})^+$ and $d = \tilde{w}_S - \tilde{w}_F$. Required explorer surplus satisfies $\tilde{w}_F + E[p_2]d - c \geq \Delta$, while the exploiter obtains $\tilde{w}_F + p_1 d - c$ above the floor. Wage cost is least when the surplus requirement binds, and the exploiter then obtains $\Delta + (p_1 - E[p_2])d$, so the least-cost contract that is least attractive to him keeps $d = c/E[p_2]$ and raises both first-period bonuses by Δ. Because $d \geq c/E[p_2]$ under every feasible contract, none leaves him less. This yields the exact minimum

$$V_{\mathrm{X}}^{\mathrm{E,NC}}(\overline{U}_{\mathrm{E}}) = 2\underline{w} + \frac{p_1 - E[p_2]}{E[p_2]} c + (\overline{U}_{\mathrm{E}} - 2\underline{w})^+.$$

The exploration problem is feasible exactly when this value is at most $\overline{U}_{\mathrm{X}}$. This establishes the contract and feasibility argument for Proposition 5.

Under commitment, the corresponding minimum starts at $2\underline{w}$, remains flat until $\overline{U}^*_N$, and then increases at one of the two rates in (B.2). Both are below one: the count ordering gives $E[p_2]E[p_2 \mid S, 2] > p_1 - E[p_2](1 - p_1) > p_1^2$, and $E[p_2 \mid S, 2] > p_1$ gives $E[p_2]E[p_2 \mid S, 2] + (1 - E[p_2])p_1 > p_1$. Since $E[p_2]E[p_2 \mid S, 2] > p_1^2$, also $\overline{U}^*_N > 2\underline{w}$. The committed value is consequently strictly lower at every common explorer requirement, while the direct cost of delivering that requirement to the explorer is unchanged.

For Corollary 1, the exploiter item is unchanged and the lower mimicking value expands the feasible set of requirement pairs. Lemma 1 therefore gives a weak reduction in both-type wage cost. More generally, commitment can replicate any no-commitment allocation, including its continuation policy and null items, so optimized profit weakly rises. Strictness needs a screening constraint that affects cost or feasibility. In a minimally normalized both-type optimum, suppose

$$\overline{U}^{\mathrm{NC}}_{\mathrm{X}} = V^{\mathrm{E,NC}}_{\mathrm{X}}(\overline{U}^{\mathrm{NC}}_{\mathrm{E}}) > \max\{\overline{U}^{\mathrm{com}}_{\mathrm{X}}, 2\underline{w}\}.$$

Replacing exploration by its committed contract permits a strictly lower exploiter requirement. Lowering $\tilde{w}_S$ on the exploitation item then strictly lowers its cost and weakly lowers the explorer's mimicking payoff. Explorer cost is unchanged, so menu cost falls strictly. If the outside option or baseline pins the exploiter requirement, this strict implication need not follow.

C.2 Termination

Termination after failure removes period 2 employment and its floor. For exploitation, continuation and first-period incentive slacks give

$$\begin{aligned} \tilde{w}_{SS} &= \tilde{w}_{SF} + c/p_1 + x, \\ \tilde{w}_S &= c/p_1 - \underline{w} + \tilde{w}_F - \tilde{w}_{SF} - p_1 x + z, \qquad x, z \geq 0. \end{aligned}$$

The intended utility is $\underline{w} + \tilde{w}_F + p_1 z$, while an explorer's best-response value is

$$\begin{aligned} V^{\mathrm{X},T,0}_{\mathrm{E}} + \tilde{w}_F + E[p_2](E[p_2 \mid S, 2] - p_1)x + E[p_2]z, \\ V^{\mathrm{X},T,0}_{\mathrm{E}} = \underline{w} + \frac{c}{p_1}\left[E[p_2](1 + E[p_2 \mid S, 2]) - p_1(1 + E[p_2])\right]. \end{aligned}$$

The count ordering makes exploratory effort preferable to shirking for the mimic. Thus $x = 0$ minimizes his value, and z delivers exploiter surplus at rate $E[p_2]/p_1 < 1$, compared with one through $\tilde{w}_F$. Paying earliest sets $\tilde{w}_{SF} = 0$. This gives Proposition 6 and

the envelope

$$V_{\mathrm{E}}^{\mathrm{X},T}(\overline{U}_{\mathrm{X}}) = V_{\mathrm{E}}^{\mathrm{X},T,0} + \frac{E[p_2]}{p_1}(\overline{U}_{\mathrm{X}} - \underline{w})^+.$$

LEMMA 2 (Termination and selection on the exploiter contract)*: For every exploiter utility requirement, the optimal terminated exploiter contract gives a misreporting explorer strictly lower utility than the optimal exploiter contract without termination.*

Proof. Subtracting the terminated envelope from the retained envelope gives

$$\underline{w} + \frac{E[p_2]}{p_1}\left[(\overline{U}_{\mathrm{X}} - 2\underline{w})^+ - (\overline{U}_{\mathrm{X}} - \underline{w})^+\right] \geq \underline{w}(1 - E[p_2]/p_1) > 0.$$

□

For exploration, the terminated baseline backloads to SS:

$$b_T = \frac{(1 + E[p_2])c - E[p_2]\underline{w}}{E[p_2]E[p_2 \mid S, 2]}, \qquad \tilde{w}_{SS} = b_T, \qquad \tilde{w}_S = \tilde{w}_F = \tilde{w}_{SF} = 0.$$

The intended utility is $\underline{w}$. A shirking exploiter receives this floor; effort becomes profitable at

$$b_T^* = \frac{(1 + p_1)c - p_1\underline{w}}{p_1^2}, \qquad \overline{U}_T^* = \underline{w} + E[p_2]E[p_2 \mid S, 2](b_T^* - b_T).$$

The maintained count ordering and low-floor condition imply $b_T < b_T^*$, so the baseline attains the global lower bound $\underline{w}$. Increasing SS first delivers explorer surplus without imitation rents; thereafter the same two rates as in the retained problem select between SS alone and $d\tilde{w}_F = p_1 d\tilde{w}_{SS}$. The full optimality proof is in Online Appendix B.1. Writing $J = E[p_2]E[p_2 \mid S, 2] + (1 - E[p_2])p_1$ and $\ell = \min\{p_1^2/(E[p_2]E[p_2 \mid S, 2]), p_1/J\}$, its envelope is

$$V_{\mathrm{X}}^{\mathrm{E},T}(\overline{U}_{\mathrm{E}}) = \underline{w} + \ell(\overline{U}_{\mathrm{E}} - \overline{U}_T^*)^+, \qquad \overline{U}_N^* - \overline{U}_T^* = \frac{\underline{w}J}{p_1}.$$

LEMMA 3 (Termination and selection on the explorer contract)*: For every explorer utility requirement, termination weakly lowers the misreporting exploiter's utility. The reduction is strict at every requirement if double success is strictly more diagnostic of exploration than first-period failure. If failure is at least as diagnostic, the reduction is strict below $\overline{U}_N^*$ and the two values coincide at and above that threshold.*

Proof. The retained envelope is $2\underline{w} + \ell(\overline{U}_{\mathrm{E}} - \overline{U}^*_N)^+$. The gap between it and the terminated envelope begins at $\underline{w}$, falls between the thresholds, and is constant thereafter at $\underline{w}(1 - \ell J/p_1)$. This is positive when SS is strictly more diagnostic and zero otherwise, including when the two histories are equally diagnostic. The positive initial gap and these slopes prove the claim. □

C.2.1 Firing thresholds and the severance region

Fix an implementable service set H, an active type θ, and the other active track's policy j when present. Put $m_{\mathrm{E}} = \mu$, $m_{\mathrm{X}} = 1 - \mu$, $q_{\mathrm{E}} = 1 - E[p_2]$, and $q_{\mathrm{X}} = 1 - p_1$. Let $\mathcal{C}_H^{I,\theta;j}$ and $\mathcal{C}_H^{T,\theta;j}$ be the minimum prior-weighted wage costs with idle retention and termination on track θ, allowing the menu and requirements to reoptimize. Define

$$\tau_\theta^j(H) = \frac{\mathcal{C}_H^{I,\theta;j} - \mathcal{C}_H^{T,\theta;j}}{m_\theta q_\theta}, \qquad g = r_2(F,F) + \big[p_1[r_2(F,S) - r_2(F,F)] - c\big]^+.$$

The main-text threshold τ_θ is $\tau_\theta^j(H)$, with the service set H and the other track's policy j left implicit. At the characterized incentive-minimal contracts, replacing productive continuation after failure by idle retention and removing its incentive bonus preserves incentives and both types' best-response utilities. It reduces the intended worker's expected continuation wage by c. The best retained use is productive when $p_1[r_2(F,S) - r_2(F,F)] \geq c$ and idle otherwise. The optimized retention-minus-termination profit difference is $m_\theta q_\theta[g - \tau_\theta^j(H)]$. Termination is therefore chosen below the threshold and retention above it; at equality both are optimal. The post-selection benchmark retains when $g \geq \underline{w}$. These comparisons are conditional on H and j, not choices among service sets.

For exploitation, Online Appendix B.1.2 decomposes the cost saving into the nonnegative direct intended-worker saving and the nonnegative saving from relaxed explorer self-selection or exclusion. Thus $\tau_{\mathrm{X}}^j(H) \geq 0$. If idle retention is infeasible while termination is feasible, the threshold is $+\infty$. With both types served and exploration retained, the idle-retention optimum is finite and its normalized exploiter requirement is at least $2\underline{w}$; the direct saving is then zero and the threshold consists entirely of the screening saving.

When first-period failure is at least as diagnostic of exploration as double success,

$$\frac{1 - E[p_2]}{1 - p_1} \geq \frac{E[p_2]E[p_2 \mid S, 2]}{p_1^2},$$

the retained exploration optimum lies in the severance region when

$$\overline{U}_{\mathrm{E}}^{N;j} \geq \overline{U}_{N}^{*},$$

where $\overline{U}_{\mathrm{E}}^{N;j}$ is its explorer utility requirement with the other track's policy j fixed, and $\overline{U}_{N}^{*}$ is the threshold defined above.

For exploration in the severance region, the joint-branch schedules satisfy

$$\tilde{w}_{F}^{T} = \tilde{w}_{F}^{I} + \underline{w}, \qquad \tilde{w}_{SS}^{T} = \tilde{w}_{SS}^{I}.$$

Hence the lost continuation floor is paid as severance. Both types' values and the exploration item's cost are unchanged. The equality extends to optimized menu costs, giving $\tau_{\mathrm{E}}^{j}(H) = 0$. The online proof uses convexity for an interior optimum, and participation and the unchanged exploitation envelope at the boundary; the fixed-requirement swap alone would not establish the optimized-cost equality.

These arguments give parts (i) and (ii) of Corollary 2, including excessive continuation on exploration when $0 < g < \underline{w}$. For part (iii), the online severance-transfer lemma carries the severance premise from terminated to retained exploitation. Exploration retention then beats termination under either exploitation policy when $g > 0$. If also $g < \tau_{\mathrm{X}}^{\mathrm{N}}(H)$, terminating exploiters beats retaining them while exploration is retained. These comparisons rank the candidate pair (N, T) above all three alternatives. Online Appendix B.1.2 supplies the complete proof, including the boundary, singleton, and extended-value cases.

C.3 Feedback

Without feedback, the common first-period payment cannot reveal the interim outcome. For Proposition 8, terminal additive bonuses

$$\tilde{w}_{SS} = 2c/p_1, \qquad \tilde{w}_{SF} = \tilde{w}_{FS} = c/p_1, \qquad \tilde{w}_{S} = \tilde{w}_{F} = \tilde{w}_{FF} = 0$$

implement exploitation with expected bonuses $2c$ and utility $2\underline{w}$. An explorer can shirk, explore once, or repeat blindly. The respective rents above the floor are

$$0, \qquad c(E[p_2]/p_1 - 1), \qquad 2c(E[p_2]/p_1 - 1),$$

so he shirks. Unlike exploration with feedback, blind repetition provides no recognized-failure fallback: its unconditional second-period success probability remains $E[p_2]$. Thus

the baseline preserves the exploiter's cost while reducing the explorer's mimicking payoff to the global floor.

For higher requirements put $\Delta_{\mathrm{X}} = (\overline{U}_{\mathrm{X}} - 2\underline{w})^+$ and $H = 2c + \Delta_{\mathrm{X}}$. The full proof in Online Appendix B.2 eliminates common transfers and FF bonuses, averages SF and FS, and uses convexity to show that overpaying the intended exploiter cannot improve screening. Writing the remaining bonuses as $a = \tilde{w}_{SS}$ and $b = \tilde{w}_{SF} = \tilde{w}_{FS}$, the exact problem is

$$
\begin{aligned}
a(b) &= \frac{H - 2p_1(1-p_1)b}{p_1^2}, \\
0 \le b \le B(\Delta_{\mathrm{X}}) &= \min\left\{\frac{H}{2p_1(1-p_1)}, \frac{H-c}{p_1}\right\}, \\
F_\Delta(b) &= E[p_2]E[p_2 \mid S,2]a(b) + 2E[p_2](1 - E[p_2 \mid S,2])b - 2c, \\
R_{\mathrm{X}}^{\mathrm{NF}}(\Delta_{\mathrm{X}}) &= \min_{0 \le b \le B(\Delta_{\mathrm{X}})} \max\{0, E[p_2]b - c, F_\Delta(b)\}.
\end{aligned}
$$

The one-period deviation value rises with b, while the blind-repetition value falls. The minimizer is their crossing clipped to the feasible interval, possibly with an interval of minimizers at zero rent. Conditional feasibility is $2\underline{w} + R_{\mathrm{X}}^{\mathrm{NF}}(\Delta_{\mathrm{X}}) \le \overline{U}_{\mathrm{E}}$.

To see that screening improves strictly at every common exploiter requirement, consider the feasible additive candidate $b = H/(2p_1)$, $a = 2b$. Its explorer rent is

$$
\left[2c(E[p_2]/p_1 - 1) + (E[p_2]/p_1)\Delta_{\mathrm{X}}\right]^+.
$$

With feedback the minimum is instead $V_{\mathrm{E}}^{\mathrm{X},0} - 2\underline{w} + (E[p_2]/p_1)\Delta_{\mathrm{X}}$, which is strictly greater because $V_{\mathrm{E}}^{\mathrm{X},0} > 2\underline{w}$ and $E[p_2] < p_1$. The optimal no-feedback item does at least as well as this candidate. The comparison preserves intended output and cost; its value to the menu depends on whether self-selection or exclusion is constraining.

For Proposition 9, feedback replaces blind repetition after failure by the fallback, raising expected revenue by

$$
(1 - E[p_2])\left[p_1 - E[p_{2^{\mathrm{NF}}} \mid F,2]\right]\left[r_2(F,S) - r_2(F,F)\right] > 0,
$$

and the revenue condition ranks adaptive exploration above every plan that explores without feedback. Under the proposition's hypothesis a blind exploration item earns less, net of effort, than the two wage-floor payments it owes, so it loses money. With retention fixed, every type can secure $2\underline{w}$ by shirking. For a served type a requirement below $2\underline{w}$ is therefore equivalent to $2\underline{w}$, and every retained item gives the other type at least $2\underline{w}$, so the retained cases of the exploitation-track step of the proof of Lemma 4

(Online Appendix B.3), which do not use $\underline{w} \leq c$, and its exploration-track and blind-track steps hold without the lemma's outside-option bounds. Beyond the maintained assumptions these steps use only the blind-loss condition, and their replacement-or-closure argument proves the remaining claims of Proposition 9.

Corollary 3 follows. Proposition 9 gives feedback on the exploration track. On the exploitation track every optimum implements $\langle {}_1{}^1_1 \rangle$; the no-feedback item gives a mimicking explorer strictly less at every exploiter utility, so the first step of Section II.B withholds feedback. The exploitation-track steps of the proof of Lemma 4 do not use the blind-loss condition, so they prove retention under the corollary's own hypotheses.

C.4 Target Plans

This subsection proves that the plans on which the paper focuses are the ones an optimal allocation implements. The lemma concerns the allocation problem of Section II, in which each item $z = (\tilde{w}, \phi, D)$ chooses its feedback and dismissal policies, and uses the envelopes of Appendices B, C.2, and C.3. The accounting is that of Sections V and VI: wages, effort costs, the wage floor, and revenue count employed periods only; when the first-period outcome remains hidden, that is, when $D = \varnothing$ and $\phi =$ withhold, the second-period action cannot depend on it and $w_S = w_F$ is required; after dismissal there is no second-period action. The agent observes whether his employment continues, so a known outcome-dependent dismissal rule reveals the first-period outcome to a retained agent: the feedback policy matters only when $D = \varnothing$, and when $D = \{S, F\}$ there is no second-period action.

The revenue condition of Section II implies the comparisons the proofs below use, each strict and in expected revenue net of effort: retained $\langle {}_1{}^1_1 \rangle$ dominates every other retained exploiter plan and every exploiter plan dismissing after failure, after success, or after both; retained-feedback $\langle {}_2{}^2_2 \rangle$ dominates every other retained-feedback explorer plan and every retained no-feedback explorer plan with at most one exploratory period, and $\langle {}_2{}^2_t \rangle$ dominates every other explorer plan dismissing after failure; retained-feedback $\langle {}_2{}^2_2 \rangle$ dominates every explorer item dismissing after success, and $\langle {}_2{}^2_t \rangle$ every explorer item dismissing after both outcomes; and retained-feedback $\langle {}_2{}^2_2 \rangle$ dominates blind repetition whenever that replacement is used.

LEMMA 4 (Target plans)*: Suppose, in addition to the maintained assumptions, that* $\underline{w} \leq c$*, the blind-loss condition holds, and* $\overline{U}_{\mathrm{E}}^{\mathrm{com}}, \overline{U}_{\mathrm{X}}^{\mathrm{com}} \geq 2\underline{w}$*. For every* $\mu \in (0,1)$*, the allocation problem has a solution. Every optimal allocation that serves explorers implements* $\langle {}_2{}^2_2 \rangle$ *or* $\langle {}_2{}^2_t \rangle$ *on an item under which the explorer observes the first-period outcome*

whenever he is retained. Every optimal allocation that serves exploiters retains them after both outcomes and implements $\langle {}_1{}^1_1 \rangle$*, with feedback provided or withheld. Dismissal, when it occurs, is therefore of an explorer after a first-period failure.*

Proof. For fixed policies, target plans, and service set, the menu problem is a finite linear program whose objective is bounded above by the target revenue; the finite collection of such problems and the null allocation give existence. The rest of the proof compares items that deliver the same utility to their intended type: one that gives the other type a weakly lower best-response value and has weakly higher surplus preserves participation, self-selection, and exclusion and weakly raises profit. Three comparisons, carried out in Online Appendix B.3, leave only the stated plans: one on the exploitation track, one on the exploration track, and one for exploration without feedback. The condition $\underline{w} \leq c$ is used only for items that dismiss, in the first comparison and through the low-floor condition $\underline{w} < c/p_1$ that it implies, and the blind-loss condition only in the third. Under the revenue condition, every omitted productive margin is strictly costly, so every optimal allocation is of the stated form. Each such allocation exerts productive effort in every employed period and ends employment only after a first-period failure, which proves Proposition 1. □