Optimal Project Management *

Alessandro Bonatti† Doruk Cetemen‡ Juuso Toikka§

September 7, 2026

Abstract

We analyze optimal dynamic contracts for long-term projects where a risk-averse agent exerts hidden effort to increase the drift of a Brownian process toward a completion threshold. Relative to the first best, the optimal contract slows progress, makes success less likely, induces earlier termination, and narrows project scope. Allowing for costly project resets generates further distortions: resets may be excessive or insufficient depending on the residual effort required afterwards. With endogenous risk-taking, optimal incentives introduce penalties for failure and reduce rewards for effort above a critical progress level, resulting in discontinuous and non-monotonic effort in progress. In the baseline model, the optimal contract can be implemented with linear spot contracts that require no long-term commitment.

Keywords: Dynamic contracting, Optimal Stopping, Project Management

JEL codes: D86, M12, J33, J41

*We would like to thank Matteo Escudé, Johannes Hörner, Aaron Kolb, and Can Urgun for helpful comments and discussions. We also thank seminar participants at City University of London, Collegio Carlo Alberto, Essex, JHU, London Business School, Miami Business School, MIT, US Naval Academy, Nottingham, Penn, Princeton, Southeast Theory Festival, Surrey, and Toulouse.

†Sloan School of Management, MIT, bonatti@mit.edu

‡LUISS, RHUL and EIEF, dcetemen@luiss.it

§The Wharton School, University of Pennsylvania, toikka@wharton.upenn.edu

1 Introduction

Many long-term projects—from scientific research and drug development to complex engineering and software initiatives—progress through uncertain paths and often require previously unanticipated steps. Success typically depends on a sequence of decisions made under limited observability and evolving conditions. Managing such projects involves dynamic trade-offs between motivating effort, tolerating setbacks, and making strategic choices about when to persevere, reset, or take on risk.

In this paper, we develop a flexible dynamic contracting framework to gain insight into optimal project management. We consider a principal who must decide how to incentivize an agent to exert effort; how tolerant she should be of failure before pulling the plug or shifting the approach; how ambitious the project should be in the first place; and when to allow the agent to take on additional risk in order to speed up progress.

To capture the rich uncertainty inherent to complicated projects, we model a project's progress by a Brownian motion that must reach a threshold, or target, for completion. The agent can increase the project's drift by exerting hidden, costly effort. At any moment in time, the distance to the target measures the expected remaining total effort required to complete the project, but the actual effort needed may turn out to be smaller or larger depending on the resolution of uncertainty.

We assume that the agent has constant absolute risk aversion (CARA) preferences and can privately save and borrow on the same terms as the principal. This allows us to abstract from wealth effects. Further, it eliminates the principal's ability to provide incentives via distorting the agent's intertemporal consumption smoothing—a feature of optimal contracts in models without private savings which seems of limited relevance for the applications we have in mind. Moreover, contracts then have a natural implementation where (history-dependent) payments only take place when the project succeeds or is abandoned.

At a more technical level, CARA preferences with hidden savings imply that instead of two state variables—progress and continuation utility—the principal's problem can be formulated recursively with progress as the only state variable. This makes our non-stationary dynamic contracting problem remarkably tractable. Specifically, in most variants we study, the principal's expected payoff is characterized by the value of an auxiliary first-best problem where effort is contractible but its cost to the agent is scaled up by a constant. The effect of agency frictions thus becomes a matter of a comparative static with respect to a cost parameter. The exception is the extension to risk taking where the agent's hidden action is two-dimensional and where the first-best and second-best solutions differ qualitatively.

Main Results In the baseline model, the principal must only decide how to reward the agent and when to stop the project. The optimal contract then features increasing effort in the progress level and a constant termination cutoff. That is, it is optimal to make the agent work harder the closer the target, and abandon the project in favor of an outside option if progress falls sufficiently far from the target due to adverse shocks. Relative to the first best (which shares these qualitative features), the optimal contract induces the agent to exert less effort and exhibits lower failure tolerance by adopting a higher termination cutoff. Thus, moral hazard extends expected completion time and makes it more likely that the project be abandoned.

The optimal contract can be implemented by a sequence of spot contracts, each of which is the optimal linear contract of the static model of Holmström and Milgrom (1987) for the current marginal value of progress. Linearity is thus a result rather than an assumption. Moreover, under this implementation the agent is indifferent between continuing and leaving at every history, and the principal gains nothing from revising the terms, so the relationship requires no commitment on either side.

The baseline model takes the target as given, but in many settings it is itself a choice variable. For instance, goals have to be set for research and development, and firms have to decide how good a new product should be before it is launched. Accordingly, we consider the case of endogenous project scope where the principal can choose the target for the project, with more ambitious targets yielding larger benefits upon completion. The findings from the baseline model carry over readily to this case. Moreover, we show that the principal optimally pursues less ambitious projects than in the first best, which points to a simple but novel form of agency cost.

We then enrich the model to incorporate endogenous project resets to capture settings where there are multiple alternative approaches to tackling the project. We assume that the principal can always abandon a failing approach and switch to a new ex ante identical one. Formally, this means that progress can be reset to a fixed restarting point by paying a fixed cost. Motivating examples here include abandoned early-stage AI models that are retrained from scratch following poor performance, software teams discarding early versions of a system, or clinical trials restarted after unfavorable interim results.[1]

The optimal contract in the model with resets implements a threshold policy: below a certain progress level, the principal pays the fixed cost to reset; otherwise, the project continues. Relative to the first best, the second-best policy may reset too aggressively when the restarting point is close to the target, and too conservatively when the restarting point is far. This result reflects an intuitive tension: using resets substitutes for effort, which is more costly to implement under moral hazard; but resets alone do not complete

[1] For illustrative examples, see Cusumano (1995), and Bubeck et al. (2023).

the project. The more a reset requires exerting further effort to complete the project, the more the value of resetting is depressed when incentives are costly to provide.

Finally, we study endogenous risk-taking. To do so, we extend the model to include a Poisson shock that terminates the project. In addition to controlling the drift, the agent chooses the riskiness of his action unobservably to the principal. The riskier action increases the project's drift relative to the safe action, but comes at the cost of a higher hazard rate of a breakdown. This structure reflects dynamic moral hazard in settings like early-phase pharmaceutical development, where firms may fast-track compounds into clinical trials to compress timelines, accepting higher risk of regulatory failure. Similarly, technology startups may try to push unstable features to market in pursuit of growth.[2]

The optimal contract allows the agent to take additional risk to accelerate development when the project is far from completion. By contrast, when success is close, the downside risk of a breakdown outweighs the gains from faster progress. Under additional assumptions, we show that the contract has a cutoff structure: once progress exceeds a threshold, the contract switches from a risky to a safe regime. In the risky regime, the agent is not penalized for breakdowns. In the safe regime, the principal deters risk taking by imposing penalties for breakdowns and reducing rewards for progress.

Interestingly, the interaction between hidden effort and risk choices implies that effort is non-monotonic in progress: it rises as progress approaches the cutoff, drops discontinuously at the regime switch, and then increases again until completion. By contrast, first-best effort is always increasing in progress.

A technical contribution of our analysis is verifying global incentive compatibility for a CARA agent with hidden savings subject to a standard pathwise no-Ponzi condition. We use localization arguments that allow us to dispense with the ad hoc integrability conditions (that imply implicit restrictions on deviations) imposed in the literature. Bloedel, Krishna, and Strulovici (2023) show an analogous result for the pure consumption-savings problem when income follows an Ornstein-Uhlenbeck process; our proof can be seen as a generalization to an arbitrary controlled income process with a finite stopping time. The approach could prove useful in other contracting environments with unbounded states.

Related Literature This paper contributes to the study of dynamic contracting in continuous time following Sannikov (2008). As noted above, we abstract from wealth effects and assume the agent can privately save and borrow on the same terms as the principal. Fudenberg, Holmström, and Milgrom (1990) show (in discrete time) that under these

[2]The effects of risk-taking behavior can be observed in practice: overall approval rates (starting from the Phase I clinical trials) for drugs developed between 2011 and 2020 averaged just 7.9%, with oncology programs below 5.3%. See Kola and Landis (2004), Pallmann et al. (2018), and Hay, Thomas, et al. (2021).

assumptions short-term contracts are without loss of optimality. Most importantly for us, the principal's value function is then separable in progress and continuation utility, rendering the state effectively one-dimensional. These assumptions are also adopted by He (2011) to study the relationship between pay-performance sensitivity and firm size; our model can be viewed as an extension of his to a non-stationary setting.[3] While this case is known to be tractable, the general case of hidden borrowing and savings is challenging—see, e.g., Di Tella and Sannikov (2021) and the references therein.

In terms of substantive questions, the most closely related papers are those considering contracting for project completion. Mason and Välimäki (2015) analyze a dynamic moral hazard problem in which an agent exerts costly effort to increase the arrival rate of a Poisson event that corresponds to project completion. Green and Taylor (2016) extend the model to a setting where the friction is cash-flow diversion rather than hidden effort and where two Poisson arrivals are required for completion, with the agent privately observing the first one. Feng, Taylor, Westerfield, and Zhang (2024) use the cash-flow diversion approach to study a project with deterministic progress subject to random, privately observed Poisson setbacks that erase all accumulated progress. In turn, Feng, Taylor, Westerfield, and Zhang (2025) consider partial setbacks in a two-stage project. Our analysis complements these papers by introducing endogenous project scope, resets, and risk taking, and by allowing for a rich stochastic model of progress similar to that in Georgiadis's (2015) team-production game. (Georgiadis (2015) considers contracting with limited instruments as an extension, but does not solve for the fully optimal contract.) For other works considering incentives for project completion, see, e.g., Shavell and Weiss (1979), Hopenhayn and Nicolini (1997), Kocherlakota (2004), Toxvaerd (2006), and Georgiadis, Lippman, and Tang (2014).

Risk-taking has been considered in a stationary dynamic contracting context by DeMarzo, Livdan, and Tchistyi (2014) who assume that, in addition to diverting funds, the agent can boost short-term profits by exposing the firm to the risk of a terminal breakdown. The authors show that when the agent's continuation value falls sufficiently low due to poor performance, he takes excessive risks. Wong (2019) analyzes the case in which taking risks increases the chance of a large loss rather than termination.[4] Finally, Biais, Mariotti, Rochet, and Villeneuve (2010) and Myerson (2015) consider models in which the agent exerts hidden effort to prevent the arrival of disasters. In contrast to our model, the moral hazard problem in these two papers is one-dimensional and there is no interaction between exerting effort and taking risks.

[3] Other papers adopting a similar approach include Holmström and Milgrom (1987), Williams (2015), He, Wei, Yu, and Gao (2017), Marinovic and Varas (2019), and Cetemen, Feng, and Urgun (2023).

[4] Bromberg-Silverstein, Moreno-Bromberg, and Roger (2021) introduce scale into the setup of Wong (2019). See also Rochet and Roger (2016) and Li and Williams (2025).

Our analysis is related to contracting in Brownian settings where non-stationarity arises because of learning, such as in Madsen (2022) or DeMarzo and Sannikov (2016). In these models the agent has private information about the profitability of the venture either on path or following a deviation, whereas here progress is publicly observable.

Also related is the literature on contracting for experimentation starting from Bergemann and Hege (1998, 2005)—see, e.g., Manso (2011), Hörner and Samuelson (2013), Guo (2016), Halac, Kartik, and Liu (2016), and McClellan (2022). The key difference relative to these papers is that while in our model the amount of effort needed to complete the project is uncertain and only learned over time, it is nevertheless always feasible (albeit not optimal) to reach the target almost surely in finite time.

Outline The rest of the paper is organized as follows. Section 2 presents the baseline model; Section 3 solves the first-best benchmark; Section 4 derives the optimal contract; Section 5 examines resetting; Section 6 analyzes endogenous risk-taking; and Section 7 concludes. All proofs are in the Appendix and in the Supplementary Appendix.

2 Baseline Model

A risk-neutral principal (she) hires a risk-averse agent (he) to complete a project. The project's progress is given by a publicly observable stochastic process, which the agent controls by expending costly, hidden effort. The events unfold in continuous time, indexed by $t \in [0, \infty)$.

Project Technology The project's progress X is a real-valued, controlled process, whose starting value is normalized to zero (i.e., $X_0 = 0$) and which evolves according to

$$dX_t = A_t\, dt + \sigma\, dZ_t, \tag{1}$$

where $A_t \in \mathbb{R}_+$ is the agent's hidden effort at time t; $\sigma > 0$ is a constant parameterizing volatility; and dZ_t is the increment of a standard Brownian motion Z, unobservable to both parties.

The project is completed when progress reaches the target $\bar{x} > 0$, i.e., the first time $X_t = \bar{x}$. The project can be terminated at any time before completion.

Beyond its literal interpretation, the project's progress can be viewed as the gradual resolution of uncertainty about the remaining effort required for completion, starting from an expected requirement of $\bar{x} - X_t$ at time t. (The target $\bar{x}$ is then the time-0 expected requirement.) In this interpretation, upward and downward movements in

progress reflect the arrival of information about the realized residual work associated with completing intermediate tasks. The Brownian shocks capture incremental learning about this residual requirement as individual steps are undertaken.[5]

Payoffs The risk-neutral principal receives a lump-sum benefit $b > 0$ upon project completion. If instead the project is terminated before it is completed, she receives a salvage value $\underline{v} \in (0, b)$. Thus, if the project is stopped (i.e., completed or terminated) at a random time τ, her payoff is given by

$$e^{-r\tau}\left(\mathbf{1}_{\{X_\tau \geq \bar{x}\}} b + \mathbf{1}_{\{X_\tau < \bar{x}\}} \underline{v}\right) - \int_0^\infty e^{-rt} C_t dt,$$

where $r > 0$ is the interest rate and C_t is the time-t payment to the agent, which may be positive or negative. (Payments can continue even after the project ends.)

The agent is risk-averse with CARA flow utility. His flow cost of effort is measured in monetary terms by a cost function $h : \mathbb{R}_+ \to \mathbb{R}_+$ with $h(a) = \phi a^2/2$ for some $\phi > 0$, assumed quadratic for simplicity.[6] The agent can save and borrow, and thus his time-t consumption, $\hat{C}_t$, may in general be different from the payment C_t received from the principal. Denoting his time-t effort by $\hat{A}_t$, the agent's payoff is given by

$$\int_0^\infty e^{-rt} u(\hat{A}_t, \hat{C}_t) dt,$$

where

$$u(\hat{a}, \hat{c}) := -\frac{1}{\eta} \exp\{-\eta(\hat{c} - h(\hat{a}))\} = -\frac{1}{\eta} \exp\left\{-\eta\left(\hat{c} - \frac{\phi}{2}\hat{a}^2\right)\right\}. \tag{2}$$

Both parties can save and borrow at the rate r. The agent's consumption and savings decisions are unobservable to the principal.

Contracting The parties commit at time 0 to a long-term contract that specifies monetary compensation for the agent as a function of the project's progress. The contract also specifies the conditions (if any) under which the project be terminated.

Formally, a contract is a triple (A, C, τ) consisting of a recommended effort process A taking values in $\mathbb{R}_+$, a payment process C taking values in $\mathbb{R}$, and a stopping time τ taking values in $\mathbb{R}_+$, all progressively measurable with respect to progress X. The stopping time τ captures both project completion and early termination. It is thus required to satisfy $\tau \leq \inf\{t \geq 0 : X_t = \bar{x}\}$ almost surely.

[5] For a related model of continuous-time gradual learning, see Branco, Sun, and Villas-Boas (2012).

[6] Our results extend qualitatively to any effort cost h with $h(0) = h_a(0) = 0$, $h_{aa} > 0$, and $h_{aaa} \geq 0$.

Given a contract (A, C, τ), the agent's problem is to choose an effort process $\hat{A}$ and a consumption process $\hat{C}$, both progressively measurable with respect to progress, to maximize his utility:

$$U(A, C, \tau) := \max_{\hat{A}, \hat{C}} \mathbb{E}^{\hat{A}} \left[\int_0^\infty e^{-rt} u(\hat{A}_t, \hat{C}_t)\, dt \right], \tag{3}$$

where the expectation is conditional on the effort process $\hat{A}$. The agent's choice of actions and consumption is subject to the laws of motion for progress and savings, given by

$$\begin{aligned} dX_t &= \hat{A}_t\, dt + \sigma\, dZ_t, \quad X_0 = 0, \\ dS_t &= (rS_t + C_t - \hat{C}_t)\, dt, \quad S_0 = 0, \end{aligned}$$

and to the transversality condition $\lim_{t\to\infty} e^{-rt} S_t = 0$ (a.s.) ruling out Ponzi schemes. Savings increase through interest income rS_t and payments C_t, and decrease through consumption $\hat{C}_t$; initial savings are normalized to zero without loss of generality.

The agent has an outside option u_0 available to him at time 0, which means that the contract (A, C, τ) satisfies the participation constraint if

$$U(A, C, \tau) \geq u_0. \tag{4}$$

We impose the participation constraint at time 0 only to simplify the analysis. However, we show in Section 4.4 that the optimal contract can be taken to satisfy it at all times.

It is convenient to focus on contracts (A, C, τ) under which it is optimal for the agent to obey the effort recommendation A and neither borrow nor save, i.e., to consume exactly C. In such a contract, the payment process C also acts as an incentive compatible consumption recommendation. Formally, we define these contracts as follows.

Definition 1 *A contract (A, C, τ) is **incentive compatible** if $(\hat{A}, \hat{C}) = (A, C)$ is a solution to the agent's problem (3).*

The following lemma is standard in the literature on dynamic contracts.[7]

Lemma 1 *For any contract with finite payoffs to the agent and the principal, there is an incentive-compatible contract that delivers the same payoffs to both parties as the original contract.*

That is, focusing on incentive-compatible contracts is without loss of generality as far as payoffs are concerned. Armed with this result, the principal's problem of designing

[7] The idea is similar to the revelation principle. For a proof, see, e.g., He (2011).

an optimal contract can be formulated as

$$\max_{A,C,\tau} \mathbb{E}^A \left[e^{-r\tau} \left(\mathbf{1}_{\{X_\tau \geq \bar{x}\}} b + \mathbf{1}_{\{X_\tau < \bar{x}\}} \underline{v} \right) - \int_0^\infty e^{-rt} C_t \, dt \right], \tag{5}$$

where the maximum is taken over incentive-compatible contracts that satisfy the agent's participation constraint (4).

It is worth noting that while the solution to (5) pins down the agent's consumption, the optimal payment policy is unique only up to the net present value because the agent can save and borrow. For instance, all payments could be deferred until the project is completed or terminated.

3 Observable Effort

Before deriving the optimal contract, it is useful to consider the first-best benchmark where the agent's effort is observable. (Whether consumption is also observable is immaterial.) In this case, the least expensive way to provide utility to the agent is by giving him constant consumption. This can be achieved by paying a flow wage C_t that has a variable part that exactly compensates for the required effort A_t, and a constant part that ensures the agent's expected utility matches his outside option u_0, i.e.,

$$C_t = \frac{\phi}{2} A_t^2 - \frac{1}{\eta} \ln(-\eta r u_0) \quad \text{for all } t \in [0, \infty).$$

It is of course optimal to ask for zero effort after the project is stopped.[8] Moreover, the constant part of C_t is irrelevant for any decisions. Thus, at any time t and at any current progress level $x \leq \bar{x}$, the first-best effort policy and stopping time solve

$$v(x) := \max_{A,\tau} \mathbb{E}^A_{t,x} \left[\mathbf{1}_{\{X_\tau \geq \bar{x}\}} e^{-r(\tau - t)} b + \mathbf{1}_{\{X_\tau < \bar{x}\}} e^{-r(\tau - t)} \underline{v} - \frac{\phi}{2} \int_t^\tau e^{-r(s-t)} A_s^2 ds \right] \tag{6}$$

subject to the law of motion for progress given in (1) with $X_t = x$.

This first-best optimal control and stopping problem is a one-player version of the team-production game in Georgiadis (2015).[9] As such, part of the solution is a termination threshold $\underline{x}^{FB} < \bar{x}$ at which the project is optimally abandoned. In the continuation region $[\underline{x}^{FB}, \bar{x}]$, the value function v defined by (6) is infinitely differentiable and satisfies

[8] If the agent's outside option is available at every t, then he should take it when the project is stopped, at which point payments can cease as the agent then receives u_0 from the outside option.

[9] See Remark 4 in Georgiadis (2015) for the case of a positive salvage value as in our model.

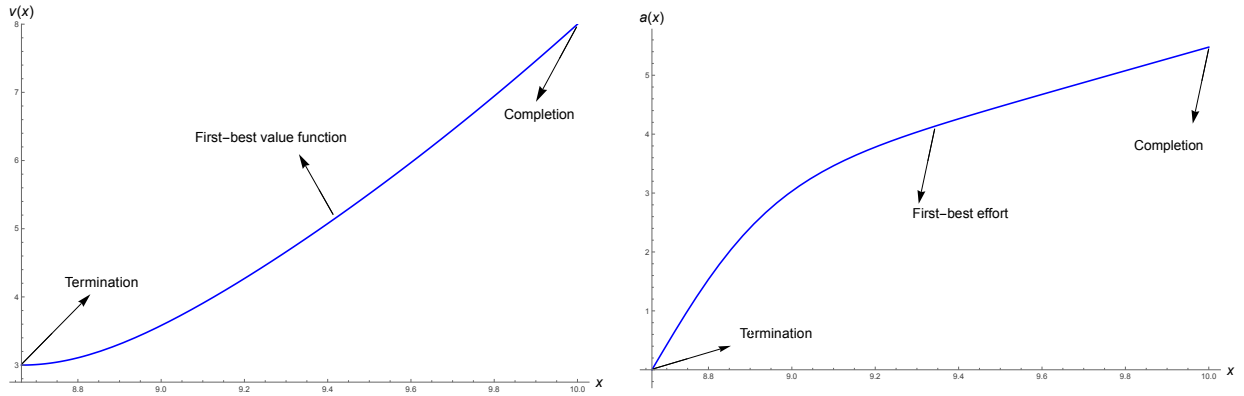


Figure 1: First-best value function (left) and effort profile (right) for parameter values $(b, \bar{x}, \sigma, \underline{v}, r, \eta, \phi) = (8, 10, 1, 3, 2, 2, 1)$.

the Hamilton-Jacobi-Bellman (HJB) equation

$$rv(x) = \max_a \left\{ -\frac{\phi}{2}a^2 + av_x(x) + \frac{1}{2}\sigma^2 v_{xx}(x) \right\} \tag{7}$$

subject to value matching ($v(\bar{x}) = b$, $v(\underline{x}^{FB}) = \underline{v}$) and smooth pasting ($v_x(\underline{x}^{FB}) = 0$). The first-best effort in turn is pinned down by the first-order condition

$$a^{FB}(x) = \phi^{-1} v_x(x), \tag{8}$$

which can be substituted back into (7) to obtain a simple ODE for v.

The following proposition summarizes properties of the first-best solution.

Proposition 1 *Let v be the first-best value function and $\underline{x}^{FB}$ the first-best termination threshold. In the continuation region $[\underline{x}^{FB}, \bar{x}]$, the following properties hold:*

(i) Value is increasing in progress: $v_x(x) \geq 0$ with strict inequality if and only if $x > \underline{x}^{FB}$.

(ii) Value is convex in progress: $v_{xx}(x) > 0$.

(iii) Effort is increasing in progress: $a_x^{FB}(x) > 0$.

(iv) Effort grows in expectation at rate r: the discounted effort process $e^{-rt}a^{FB}(X_t)$ is a martingale up to the first-best stopping time τ.

The first-best policy has an intuitive structure. If progress is sufficiently far from the target, the project is terminated, as completing it would be too costly. At the termination threshold, zero effort is implemented. Everywhere else effort is positive and increasing in progress. This arises from the payoff structure in the model: The cost of exerting effort

is incurred continuously, while the reward is received only upon project completion. Therefore, the principal has a stronger incentive to accelerate progress as the project nears completion. The same logic operates along the path. Since progress yields no payoff until the project is stopped, the envelope theorem implies that the marginal value of progress today equals the expected discounted marginal value of progress at any later date. It therefore grows in expectation at the discount rate, and so does effort, which is proportional to it. Figure 1 illustrates the first-best value function and effort policy.

Next, we establish some comparative statics for the first-best solution.

Proposition 2 *Consider two projects, labeled 1 and 2, that differ in exactly one parameter.*

(i) If $b_1 > b_2$, *then* $\underline{x}_1^{FB} < \underline{x}_2^{FB}$. *Furthermore,* $a_1^{FB}(x) > a_2^{FB}(x)$ *for all* $x \in (\underline{x}_1^{FB}, \bar{x}]$.

(ii) If $r_1 > r_2$, *then* $\underline{x}_1^{FB} > \underline{x}_2^{FB}$. *Furthermore, there exists an interior threshold* $x_r^\dagger$ *such that* $a_1^{FB}(x) > a_2^{FB}(x)$ *if* $x \in (x_r^\dagger, \bar{x}]$ *whereas* $a_1^{FB}(x) < a_2^{FB}(x)$ *if* $x \in (\underline{x}_2^{FB}, x_r^\dagger)$.

(iii) If $\sigma_1 > \sigma_2$, *then* $\underline{x}_1^{FB} < \underline{x}_2^{FB}$. *Furthermore, there exist interior thresholds* $\{x_\sigma^\dagger, x_\sigma^{\dagger\dagger}\}$ *such that* $a_1^{FB}(x) > a_2^{FB}(x)$ *if* $x \in (\underline{x}_1^{FB}, x_\sigma^\dagger)$ *whereas* $a_1^{FB}(x) < a_2^{FB}(x)$ *if* $x \in (x_\sigma^{\dagger\dagger}, \bar{x}]$.

(iv) If $\phi_1 > \phi_2$, *then* $\underline{x}_1^{FB} > \underline{x}_2^{FB}$. *Furthermore,* $a_1^{FB}(x) < a_2^{FB}(x)$ *for all* $x \in (\underline{x}_2^{FB}, \bar{x}]$.

Part (i) is intuitive: the higher the reward, the more valuable the project and the stronger the incentive to exert effort.

Part (ii) states that a less patient principal terminates the project earlier than a more patient one. Further, the less patient principal asks for more effort if and only if the project is sufficiently close to completion. This is because she values early completion more highly than a more patient principal. Since the reward is realized only upon completion, and additional effort is costly upfront, the net benefit of accelerating completion outweighs its cost only when the project is near the threshold.

Part (iii) shows that an increase in the volatility σ lowers the termination threshold. This is because completing the project is a real-option, whose value increases with volatility. This part also shows that incentives for effort become stronger with volatility when the project is far from completion, and weaker as it nears completion. To see this, observe that when the project is near the termination threshold, the downside is negligible since $v_x(x) \approx 0$ and $v(x) \approx \underline{v}$. In contrast, the upside is significant because $v_{xx}(\underline{x}) > 0$. The opposite occurs as the project approaches the completion point $\bar{x}$.

Finally, part (iv) establishes that higher effort costs lead to earlier termination and uniformly lower effort. The intuition is direct: when effort is more expensive, the principal requires less of it at every progress level and abandons projects sooner since the cost of reaching completion increases.

4 Unobservable Effort

In this section, we begin the analysis of the optimal contracting problem under moral hazard. The relevant state variables for the problem are the project's current progress level and the agent's continuation utility. We first note some necessary conditions for incentive compatibility and use them to formulate the principal's problem recursively. We then characterize the optimal contract.

4.1 Local Incentive Compatibility

Given an incentive-compatible contract (A, C, τ), the agent's continuation utility is defined by

$$W_t := \mathbb{E}_t^A \left[\int_t^\infty e^{-r(s-t)} u(A_s, C_s) ds \right]. \tag{9}$$

Its evolution takes the following familiar form.

Lemma 2 (Sannikov, 2008; He, 2011) *The agent's continuation utility under any incentive-compatible contract (A, C, τ) satisfies*

$$dW_t = \big(rW_t - u(A_t, C_t)\big)\, dt + \beta_t(-\eta r W_t)(dX_t - A_t dt) \tag{10}$$

for some progressively measurable process β.

The proof is standard and hence omitted. We have adopted He's (2011) convention of normalizing the process β controlling the strength of incentives by $-\eta r W_t$. (Note that $W_t < 0$ because of CARA utility.) This way, β_t measures directly the sensitivity of the agent's monetary compensation to his performance.

The following lemma gives necessary conditions for incentive compatibility.

Lemma 3 (He, 2011) *Let (A, C, τ) be an incentive-compatible contract and let β be the process from (10). Then $rW_t = u(A_t, C_t)$ and thus the agent's continuation utility W_t is a martingale. Moreover, the agent's consumption is given by*

$$C_t = \frac{\phi}{2} A_t^2 - \frac{1}{\eta} \ln(-\eta r W_t), \tag{11}$$

and effort is given by

$$A_t = \phi^{-1} \beta_t. \tag{12}$$

The above conditions follow readily from first-order conditions to the agent's problem as in He (2011). To sketch the argument, suppose the agent has deviated under the

incentive-compatible contract (A, C, τ) and has accumulated private savings $S_t = s \in \mathbb{R}$ at time t. Then his continuation utility, $W_t(s)$, depends only on the public history of progress (as the contract conditions on it) and the savings s. Using the properties of CARA utility, it is straightforward to verify that $W_t(s) = e^{-\eta r s} W_t$, where $W_t = W_t(0)$ is the continuation utility at this same progress history under no savings (see Lemma A.1 in the Appendix). This is because under CARA, the optimal use of the savings s is to simply consume the interest rs at every time instant in perpetuity, and the corresponding constant part of the consumption utility can be factored out of the agent's problem. This implies that, holding effort fixed, the Euler equation for optimal time-t consumption takes the form

$$u_c(A_t, C_t) = \frac{d}{ds} W_t(s) = -\eta r e^{-\eta r s} W_t.$$

Noting that $u_c(A_t, C_t) = -\eta u(A_t, C_t)$ by (2) and that $s = 0$ by incentive compatibility, this Euler equation simplifies to $u(A_t, C_t) = r W_t$, which is the first claim in Lemma 3. Using CARA utility again, this is seen to be equivalent to (11).

The effort condition (12) in turn is simply the condition for A_t to be optimal in the agent's one-shot-deviation problem

$$\max_a a\beta_t(-\eta r W_t) + u(a, C_t),$$

where the first term captures the expected impact of effort on continuation utility and the second term captures the flow disutility of effort (cf. Sannikov, 2008). Taking the first-order condition and evaluating it at $a = A_t$ gives $\beta_t(\eta r W_t) = u_a(A_t, C_t)$. But by CARA utility, $u_a(A_t, C_t) = u_c(A_t, C_t)(-\phi A_t) = \eta r W_t \phi A_t$, where the second equality follows by the consumption Euler evaluated at $s = 0$. Thus, $\beta_t = \phi A_t$ as claimed.

The upshot of Lemma 3 for the principal's problem is that the level of the agent's continuation utility cannot play a role in providing incentives. By (12), the sensitivity β_t of the agent's monetary compensation to performance pins down his effort A_t independent of W_t, as expected under CARA preferences. Furthermore, because the agent can freely save and borrow, toying with the timing of when W_t is delivered does not change what the agent does. One way to deliver W_t is by using an incentive-compatible contract under which the agent does not save or borrow, in which case the payments to the agent are additively separable in the cost of effort and the interest on the certainty equivalent of W_t as shown in equation (11). This means that the agent's continuation utility W_t enters the principal's value function as an additive cost term; we verify this formally in the next section.

4.2 Recursive Formulation

We consider a relaxation of the second-best problem (5) where incentive compatibility is replaced with the local conditions from Lemma 3. That is, at any time t, given current progress $x \leq \bar{x}$ and promised continuation utility w, we solve

$$f(x,w) := \max_{A,C,\tau} \mathbb{E}^A_{t,(x,w)} \left[\mathbf{1}_{\{X_\tau \geq \bar{x}\}} e^{-r(\tau-t)} b + \mathbf{1}_{\{X_\tau < \bar{x}\}} e^{-r(\tau-t)} \underline{v} - \int_t^\infty e^{-r(s-t)} C_s ds \right] \tag{13}$$

subject to the laws of motion (1) and (10) with $(X_t, W_t) = (x,w)$ as well as the local incentive compatibility conditions (11) and (12). Then, by construction, $f(0,u_0)$ is an upper bound on the principal's maximal payoff in (5).

To formulate the relaxed problem recursively, we guess that the value function f is additively separable with the following functional form:

$$f(x,w) = \pi(x) + \frac{1}{\eta r} \ln(-\eta r w). \tag{14}$$

We refer to the function π as the project value, while the second term is the negative of the agent's certainty equivalent for the continuation utility w (optimally consumed as an annuity). Given the additively separable form (14), the HJB equation for the relaxed problem is

$$r f(x,w) = \max_{a,c,\beta} \left\{ -c + a f_x(x,w) + \frac{1}{2}\sigma^2 f_{xx}(x,w) + \frac{1}{2}[\beta \eta r (-w)\sigma]^2 f_{ww}(x,w) \right\},$$

because the drift of continuation utility is zero by Lemma 3, and $f_{xw} = 0$. Substituting for consumption c and sensitivity β using (11) and (12), and using the functional form of f, this becomes

$$r\pi(x) + \frac{1}{\eta}\ln(-\eta r w) = \max_a \left\{ -\frac{\phi}{2}a^2 + \frac{1}{\eta}\ln(-\eta r w) + a\pi_x(x) + \frac{1}{2}\sigma^2 \pi_{xx}(x) - \frac{1}{2}\eta r \phi^2 a^2 \sigma^2 \right\},$$

which simplifies to

$$r\pi(x) = \max_a \left\{ -\frac{\hat{\phi}}{2}a^2 + a\pi_x(x) + \frac{1}{2}\sigma^2 \pi_{xx}(x) \right\}, \tag{15}$$

where we have defined $\kappa := 1 + \phi\eta r\sigma^2$ and $\hat{\phi} := \phi\kappa$.

Note that equation (15) only involves the project value π, with progress the only state variable. By inspection of (7), it is of the same form as the first-best HJB, the only

difference being the adjusted cost parameter $\hat{\phi}$. It follows that here, too, the solution consists of a termination threshold $\underline{x}$ and a value function π which is infinitely differentiable on the continuation region $[\underline{x}, \bar{x}]$ and satisfies the HJB equation (15) subject to value matching

$$\pi(\bar{x}) = b, \quad \pi(\underline{x}) = \underline{v}, \tag{16}$$

and smooth pasting

$$\pi_x(\underline{x}) = 0. \tag{17}$$

This solution then defines the value f of the relaxed problem (13) via (14).

We show next that the bound is tight: $f(0, u_0)$ is the value to the principal's (unrelaxed) problem (5), and π is the second-best project value function. To this end, note that the solution to the HJB equation (15) subject to (16) and (17) defines a contract as follows: Set effort $A_t = \hat{\phi}^{-1}\pi_x(X_t)$ and let $\beta_t = \phi A_t$; determine consumption C_t and continuation utility W_t via (10) and (11); terminate the project when X_t first hits $\underline{x}$.

Proposition 3 *Let (A, C, τ) be the contract defined by the solution to (15) subject to (16)–(17). Then (A, C, τ) is incentive compatible, and thus it is an optimal contract.*

Proposition 3 follows from Proposition 12, which gives an analogous result under endogenous risk taking. We comment on our novel proof strategy after the latter result.

4.3 The Optimal Contract

By Proposition 3, an optimal contract can be found by solving the auxiliary planner's problem (15). This problem differs from the first-best problem in section 3 only in that the cost parameter ϕ is replaced by the incentive-adjusted cost parameter $\hat{\phi} = \kappa\phi > \phi$, where $\kappa = 1 + \phi\eta r\sigma^2$.[10] The adjustment is a risk premium. Inducing effort a requires tying the agent's compensation to progress with sensitivity $\beta = \phi a$ (Lemma 3). Compensating the CARA agent for the resulting flow variance $\beta^2\sigma^2$ costs $\frac{\eta r}{2}\beta^2\sigma^2$ per unit of time; the coefficient of risk aversion is ηr rather than η because the agent spreads any change in wealth over his future consumption.[11] Hence, the incentive-adjusted cost

$$\frac{\hat{\phi}}{2}a^2 = \frac{\phi}{2}a^2 + \frac{\eta r}{2}\beta^2\sigma^2$$

[10] The reduction to an auxiliary planner's problem is not specific to quadratic costs. For a general effort cost h, the sensitivity required to induce effort a is $\beta = h_a(a)$, and the incentive-adjusted cost is $\hat{h}(a) := h(a) + \frac{\eta r\sigma^2}{2}h_a(a)^2$ (cf. He, 2011, eq. (13)), which reduces to the multiplicative form $\hat{h} = \kappa h$ when h is quadratic. Our results carry over to a general h as long as $\hat{h}$ is more convex than h in the sense that $\hat{h}_{aa} \geq h_{aa}$, which is ensured by $h_{aaa} \geq 0$ (footnote 6).

[11] See Lemma A.1 in the Appendix.

is simply the cost of effort plus the risk premium for the exposure needed to induce it. The upshot is that the effect of dynamic moral hazard on effort and the principal's value has thus been reduced to a comparative static with respect to ϕ in the first-best problem. (The agent's compensation and consumption of course differ qualitatively from the first best as he now has skin in the game.)

The first-order condition to problem (15) gives the optimal effort

$$a(x) = \hat{\phi}^{-1}\pi_x(x).$$

By Lemma 3, the pay-performance sensitivity required to induce it is

$$\beta(x) = \phi a(x) = \frac{\pi_x(x)}{1 + \phi\eta r\sigma^2}. \tag{18}$$

Thus, at each progress level, incentives are exactly as in the static linear-CARA-normal model of Holmström and Milgrom (1987) for a project with marginal product $\pi_x(x)$ and an agent with risk-aversion coefficient ηr. Substituting the first-order condition back into the HJB equation (15) yields a second-order ODE for the project value π:

$$r\pi(x) = \frac{1}{2\hat{\phi}}\pi_x(x)^2 + \frac{1}{2}\sigma^2\pi_{xx}(x). \tag{19}$$

The problem thus boils down to finding a termination threshold $\underline{x}$ and a project value π that solve this ODE subject to value matching (16) and smooth pasting (17).

The following proposition summarizes the optimal contract.

Proposition 4 *Let π be the optimal project value and $\underline{x}$ the optimal termination threshold. In the continuation region $[\underline{x}, \bar{x}]$, the following properties hold:*

(i) Value is increasing in progress: $\pi_x(x) \geq 0$ with strict inequality if and only if $x > \underline{x}$.

(ii) Value is convex in progress: $\pi_{xx}(x) > 0$.

(iii) Effort and incentives are increasing in progress: $a_x(x) > 0$ and $\beta_x(x) > 0$.

(iv) Effort grows in expectation at rate r: the discounted effort process $e^{-rt}a(X_t)$ is a martingale up to the second-best stopping time τ.

Since the optimal effort and project value are derived from the auxiliary planner's problem, the result and the intuition are analogous to Proposition 1.

A closed-form example When $\underline{v} = \hat{\phi}\sigma^2/2$, the project value ODE can be solved in closed form to illustrate Proposition 4. In this case, value is quadratic in progress, with

$$\pi(x) = \underline{v} + \frac{r\hat{\phi}}{2}(x - \underline{x})^2, \qquad \underline{x} = \bar{x} - \sqrt{\frac{2(b - \underline{v})}{r\hat{\phi}}}.$$

Effort and incentives are linear in the distance from the termination threshold:

$$a(x) = r(x - \underline{x}), \qquad \beta(x) = \phi r(x - \underline{x}).$$

Since $a(X_t) = r(X_t - \underline{x})$, effort has drift $r\,a(X_t)$ along the path. □

The next proposition provides comparative statics of the optimal contract.

Proposition 5 *Consider two projects, labeled 1 and 2, that differ in exactly one parameter.*

(i) If $b_1 > b_2$, then $\underline{x}_1 < \underline{x}_2$ and $a_1(x) > a_2(x)$ for all $x \in (\underline{x}_1, \bar{x}]$.

(ii) If $\eta_1 > \eta_2$, then $\underline{x}_1 > \underline{x}_2$ and $a_1(x) < a_2(x)$ for all $x \in (\underline{x}_2, \bar{x}]$.

(iii) If $\phi_1 > \phi_2$, then $\underline{x}_1 > \underline{x}_2$ and $a_1(x) < a_2(x)$ for all $x \in (\underline{x}_2, \bar{x}]$.

(iv) If $r_1 > r_2$, then $\underline{x}_1 > \underline{x}_2$, and there exists $x^\ddagger \in (\underline{x}_1, \bar{x}]$ such that $a_1(x) < a_2(x)$ for all $x \in (\underline{x}_2, x^\ddagger)$ and $a_1(x) > a_2(x)$ for all $x \in (x^\ddagger, \bar{x}]$.

The reward b affects the optimal contract as it does the first best. Risk aversion η and the cost parameter ϕ enter the second-best problem only through the incentive-adjusted cost parameter $\hat{\phi}$, which is increasing in each, so a more risk-averse agent is treated like one with a higher cost of effort: the principal tolerates less failure and induces less effort. The discount rate r enters twice, as the rate proper and through $\hat{\phi}$. Both channels make the principal less tolerant of failure, but they conflict for effort: impatience raises effort near the target (Proposition 2(ii)), whereas higher $\hat{\phi}$ lowers it everywhere, so the contract with the higher discount rate implements less effort near the threshold and either less or more near the target. Volatility is absent from the proposition because it, too, enters twice, and the two channels conflict already for the threshold: a higher σ raises the option value of continuing but also the cost of incentives, and either may dominate.

Figure 2 plots two sample paths of progress under the optimal contract, together with the agent's continuation utility, effort, and consumption (which under an incentive-compatible contract coincides with compensation). In the left panel the project is completed; in the right panel it is terminated by the principal. Along the successful path, effort and consumption rise as progress nears the target. At completion, effort drops

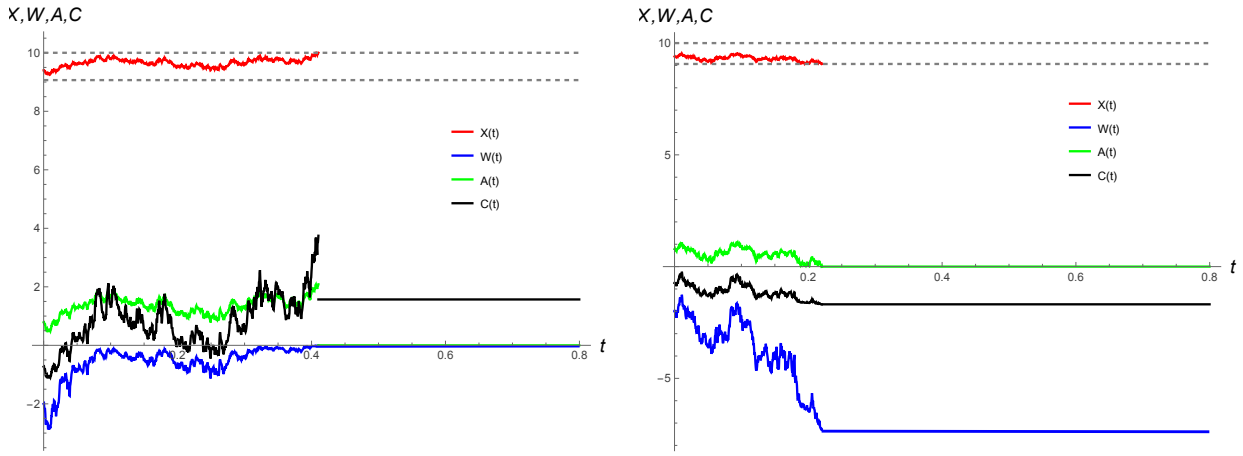


Figure 2: Sample paths for a complete project (left panel) and a terminated project (right panel). Parameter values: $(b, \bar{x}, \sigma, \underline{v}, r, \eta, \phi, X_0) = (8, 10, 1, 3, 2, 2, 1, 9.4)$.

to zero and consumption drops by the cost of effort and stays constant thereafter. This downward jump does not arise in the unsuccessful project, since effort converges to zero as the termination threshold approaches, so that consumption converges continuously to its post-termination level.

Note that, as the project evolves, the optimal contract remains renegotiation-proof: the principal's full value $\pi(X_t) + \frac{1}{\eta r} \ln(-\eta r W_t)$ is decreasing in the agent's continuation value W_t at every (X_t, W_t), so at no history can both parties gain by replacing the contract (cf. Fudenberg et al., 1990).

We then compare the optimal contract to the first best.

Proposition 6 *Let π, $\underline{x}$, and a be the project value, termination threshold, and effort under the optimal contract, and let v, $\underline{x}^{FB}$, and a^{FB} be their first-best counterparts.*

(i) Value is lower under moral hazard, with the difference vanishing with the risk premium:

$$0 < v(x) - \pi(x) \leq \phi\eta r\sigma^2 \left(b - v(x)\right) \quad \text{for all } x \in (\underline{x}^{FB}, \bar{x}).$$

(ii) Moral hazard makes the principal less tolerant of failure: $\underline{x} > \underline{x}^{FB}$.

(iii) Effort is uniformly lower under moral hazard: $a(x) < a^{FB}(x)$ for all $x \in [\underline{x}, \bar{x}]$.

Figure 3 plots the second-best value function and effort profile against their first-best counterparts. Since the second-best problem is the first-best problem with $\hat{\phi} > \phi$, the comparisons follow as comparative statics of the first-best problem. The upper bound on the value gap in turn obtains by implementing the first-best policy under moral hazard. It shows that the cost of moral hazard vanishes if effort costs vanish ($\phi \to 0$), the agent

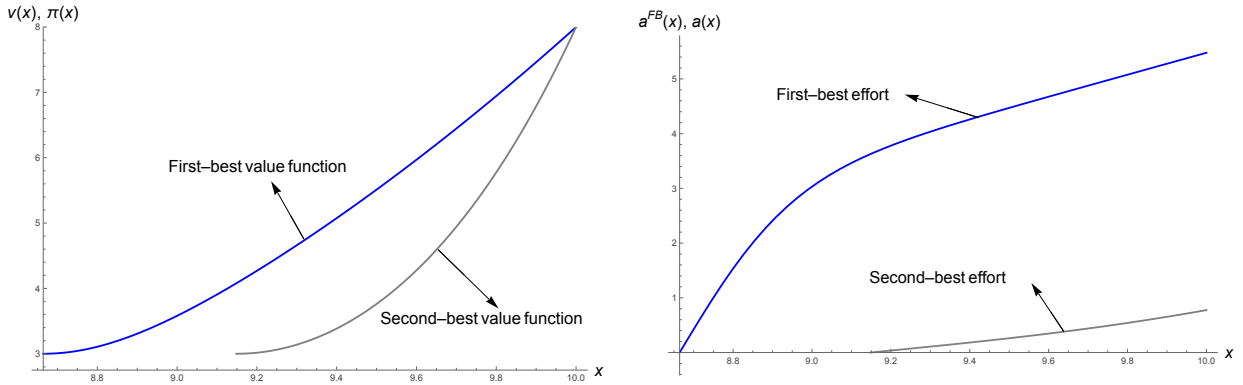


Figure 3: Second-best value function (left) and effort profile (right) compared to the first-best. Parameter values: $(b, \bar{x}, \sigma, \underline{v}, r, \eta, \phi) = (8, 10, 1, 3, 2, 2, 1)$.

approaches risk neutrality ($\eta \to 0$), noise vanishes ($\sigma^2 \to 0$), or the parties become perfectly patient ($r \to 0$) making waiting free, with both policies relying on mere luck.

The ranking of efforts and thresholds has a sharp implication for progress itself, which can be seen by using a coupling argument.

Proposition 7 *Let X^{FB} and X denote progress under the first-best and the second-best policy when both are driven by the same Brownian motion, and let τ^{FB} and τ be the corresponding stopping times. Then, almost surely, $X^{FB}_{t\wedge\tau^{FB}} > X_t$ for all $0 < t < \tau$. Consequently, whenever the second best completes the project, the first best has completed it strictly earlier; whenever the first best terminates the project, the second best has terminated it strictly earlier; and the probability of completion is higher under the first best.*

That is, moral hazard not only makes the project more likely to fail, it also extends the completion time for any realized path of the Brownian shocks. In particular, if we identify termination with completion at $+\infty$, then for any path of shocks, the first-best project is completed weakly earlier than the second best, and strictly earlier whenever the first-best completion time is finite. Conversely, identifying completion with termination at $+\infty$, the second best leads to weakly earlier termination than the first best along every realized shock path, strictly so when the second-best termination time is finite. Project durations are not ranked, however, if we do not condition on the outcome: the project can end earlier under the second best, when it is terminated while the more failure-tolerant first-best policy still continues, or later, when the first best has already completed it while the second best is still at work.

4.4 Implementation

The optimal contract pins down the agent's consumption, but the payments are unique only up to their present value along each path, because the agent can privately save and borrow. We say that a payment scheme implements the optimal contract if, given the principal's stopping rule, it is optimal for the agent to follow the optimal effort policy $A_t = a(X_t)$, and this yields both parties the same expected payoffs as the optimal contract. We note here three natural implementations. The first is the incentive-compatible contract itself, under which the agent consumes what he is paid: by (11), $C_t = \frac{\phi}{2}A_t^2 + r\,CE_t$, where $CE_t := -\frac{1}{\eta r}\ln(-\eta r W_t)$ is the certainty equivalent of his continuation utility.

Second, all payments can be deferred until the project is stopped. Paying at τ the lump sum $\int_0^\tau e^{r(\tau-t)}C_t\,dt + CE_\tau$, the compounded value of the interim payments plus the certainty equivalent of the remaining promise, leaves the present value of payments (and hence the agent's feasible set) unchanged along every path. In this implementation, the agent either borrows against the lump sum payment to finance consumption along the way, or, after poor performance, saves in anticipation of a small or even negative payment—a penalty akin to liquidated damages or a malus on deferred pay. Payment then occurs only upon completion or termination, though the amount still depends on the path of progress, and not just the stopping time.

Third, the optimal contract can be implemented via a sequence of spot contracts (cf. Fudenberg et al., 1990). By Itô's formula, $dCE_t = \beta_t(dX_t - A_t\,dt) + \frac{\eta r}{2}\beta_t^2\sigma^2\,dt$. Consider the payment scheme Γ that pays at each $t \le \tau$ (i.e., while the project is active) the flow

$$d\Gamma_t := \left[r\,CE_0 + \frac{\phi}{2}A_t^2 + \frac{\eta r}{2}\beta_t^2\sigma^2\right]dt + \beta_t\,(dX_t - A_t\,dt), \tag{20}$$

with $A_t = a(X_t)$ and $\beta_t = \beta(X_t)$, and nothing thereafter. Note that $CE_0 = -\frac{1}{\eta r}\ln(-\eta r u_0)$ is the certainty equivalent of the agent's outside option, because the participation constraint (4) binds. The first term in (20) is a fixed wage that compensates the agent for forgoing his outside option while the project is active; this term vanishes if the outside option is to live off savings since we normalized $S_0 = 0$ (without loss). By construction, $d\Gamma_t = r\,CE_0\,dt + \frac{\phi}{2}A_t^2\,dt + dCE_t$, so by inspection of (11), the scheme pays the increment dCE_t of the agent's certainty equivalent in place of its annuity $rCE_t dt$, leaving to the agent the consumption smoothing the incentive-compatible contract performs on his behalf.

Proposition 8 *For every $t \ge 0$, almost surely, the agent's continuation problem at t under the scheme Γ with savings S coincides with his continuation problem under the optimal contract with*

savings $S - (CE_t - CE_0)$. *Hence,* Γ *implements the optimal contract, with the agent consuming* C_t *and holding savings* $CE_t - CE_0$ *while the project is active, and his continuation utility at every* t *with savings* S *equals* $e^{-\eta r S} u_0$*, the value of his outside option given savings* S*.*

The idea is simple: rather than smoothing the agent's consumption, the principal can just pay him what he is owed as it accrues, and the agent smooths consumption himself by consuming only the interest on his savings. That is, whereas under the incentive-compatible contract the agent's financial account is kept with the principal, the spot implementation Γ has the agent bank with the outside credit market instead.

Remarkably, the spot implementation Γ requires no commitment on either side. The agent's outside option scales with his wealth: leaving with savings S is worth $e^{-\eta r S} u_0$, since the savings are optimally consumed as an annuity (Lemma A.1). By Proposition 8, this is exactly his continuation utility at every history, whatever his savings, so the participation constraint holds at all times. The scheme is likewise sequentially rational for the principal: delivering $e^{-\eta r S_t} u_0$ to an agent with savings S_t costs the same as delivering u_0 to an agent with no savings, namely CE_0, so her continuation value is $\pi(X_t) - CE_0$, the value of the optimal contract against an agent with no rent. Thus, at no history can the principal gain by altering the spot terms or the stopping rule. She also owes nothing once the project stops, because the agent then takes the outside option, which together with his savings delivers the promised continuation utility W_τ.

It is also worth highlighting that each spot contract in Γ is an optimal contract in the static linear-CARA-normal model of Holmström and Milgrom (1987): by (20) and (18), it pays the agent his outside option, his effort cost, and a risk premium, plus $\beta_t = \pi_x(X_t)/(1 + \phi\eta r\sigma^2)$ times the innovation in progress, which is their contract for output dX_t with marginal product $\pi_x(X_t)$, risk aversion ηr, and outside option $r\,CE_0\,dt$. Importantly, here linearity is not assumed but it is shown to be fully optimal. Unlike in their stationary model, however, here the slope varies with progress, so the spot contracts do not aggregate into a contract that is linear in cumulative progress.[12]

4.5 Endogenous Scope and Flow Payoffs

We now discuss two immediate extensions of the baseline model.

[12] In the stationary version of our model, in which the principal receives the flow dX_t rather than a reward upon completion, the optimal contract is stationary, with $\beta = 1/\kappa$ and effort $1/\hat{\phi}$. In this version the linear spot contracts aggregate and the agent's compensation is linear in cumulative progress over any horizon, reproducing the solution in Section 5 of Holmström and Milgrom (1987) with ηr as the coefficient of risk aversion. This gives an infinite-horizon model with private intermediate consumption and savings in which linearity is fully optimal, with no need for limiting arguments (cf. Hellwig and Schmidt, 2002).

Endogenous Scope Suppose that instead of the target $\bar{x}$ and the benefit b being given, the principal can choose the scope of the project. Specifically, assume that a project completed at progress level x yields a benefit $b(x)$, where b is strictly increasing and concave with $b(0) < \underline{v} < \lim_{x\to\infty} b(x)$ and $b_x(x) \to 0$ as $x \to \infty$. The principal can then end the project at any time τ and either abandon it for the salvage value $\underline{v}$ or declare it complete for the benefit $b(X_\tau)$. The solution again takes the form of a project value function π that solves the project value ODE (19) on the continuation region $[\underline{x}, \bar{x}]$, with the project terminated at $\underline{x}$ and completed at $\bar{x}$, but now both thresholds are endogenous. Thus, value matching and smooth pasting must hold at both ends: $\pi(\underline{x}) = \underline{v}$, $\pi_x(\underline{x}) = 0$, $\pi(\bar{x}) = b(\bar{x})$, and $\pi_x(\bar{x}) = b_x(\bar{x})$ (see the Supplementary Appendix for details). The first-best problem is analogous, with ϕ in place of $\hat{\phi}$; its continuation region is $[\underline{x}^{FB}, \bar{x}^{FB}]$.

Proposition 9 *With endogenous scope, the optimal contract is less ambitious and less tolerant of failure than the first best:* $\bar{x} < \bar{x}^{FB}$ *and* $\underline{x} > \underline{x}^{FB}$.

That is, agency frictions may distort the goal setting and design of projects in addition to their execution.

Flow Payoffs Suppose that, while the project is active, the principal receives a flow payoff that is increasing in progress. As the project gets closer to target, the flow payoff makes completion relatively less important. We show in the Supplementary Appendix that this may lead to hump-shaped effort under the optimal contract, with the agent working the hardest at intermediate progress levels. In terms of Proposition 4(iv), progress now yields a payoff before the project is stopped, so the envelope theorem equates the marginal value of progress today to the expected discounted marginal value at a later date plus the marginal flow payoffs accrued in between. The marginal value of progress, and hence effort, therefore grows in expectation at a rate below r, and effort can fall as progress approaches the target.

5 Endogenous Resets

In the baseline model, there was only a single viable approach to reaching the target, and abandoning it thus amounted to terminating the entire project. Here, we analyze the opposite extreme, in which an infinite number of ex ante identical approaches is available, and the principal can switch to a fresh one instead of terminating the project.

Specifically, the principal can reset the project at any time by incurring a fixed cost $L > 0$, upon which progress restarts from an exogenously given level $x^R < \bar{x}$. We allow

for $x^R \neq 0$, so that the original approach may be superior (or inferior) to the others.[13]

A contract now also specifies the increasing sequence of stopping times $T_1 \leq T_2 \leq \cdots$ at which progress is reset to x^R. The principal's problem is (5) with the reset costs subtracted:

$$\max_{A,C,\tau,(T_i)} \mathbb{E}^A \left[e^{-r\tau} \left(\mathbf{1}_{\{X_\tau \geq \bar{x}\}} b + \mathbf{1}_{\{X_\tau < \bar{x}\}} \underline{v} \right) - \int_0^\infty e^{-rt} C_t \, dt - \sum_i \mathbf{1}_{\{T_i \leq \tau\}} e^{-rT_i} L \right], \tag{21}$$

where the maximum is taken over incentive-compatible contracts that satisfy the participation constraint (4).

Let $\pi^{NR}(\cdot;\hat{\phi})$ denote the project value of the baseline model, that is, without resets (with $\hat{\phi}$ replaced by ϕ, this is the first-best value). The following lemma gives the condition for resets to be used at all and shows how it depends on the cost parameter.

Lemma 4 *Let $\bar{L}(\hat{\phi}) := \pi^{NR}(x^R;\hat{\phi}) - \underline{v}$. The optimal contract uses resets if $L < \bar{L}(\hat{\phi})$ and does not if $L > \bar{L}(\hat{\phi})$. If $\bar{L}(\hat{\phi}) > 0$, then it is strictly decreasing in $\hat{\phi}$.*

The intuition behind the lemma is straightforward. If the condition $\pi^{NR}(x^R;\hat{\phi}) - L > \underline{v}$ holds, then whenever the principal is on the brink of abandoning the project, she will prefer to invest L and to jump to the point x^R where her continuation value is higher. Conversely, if the fixed cost of resetting is too large relative to the gain in continuation value, it is better to abandon the project. The comparative static in $\hat{\phi}$ follows from the mimicking argument in the proof of Proposition 2(iv): a higher cost parameter reduces the value function $\pi^{NR}(x^R;\hat{\phi})$ at every point of the continuation region, and hence lowers the threshold reset cost $\bar{L}(\hat{\phi})$ below which resetting is optimal.

We now examine how the principal should optimally utilize resetting, and how the second-best resetting policy differs from the first-best. To that end, we assume $L < \bar{L}(\hat{\phi})$, so that resets are used. The derivation of Section 4 applies, since resets do not involve the agent, and the project value π now satisfies

$$r\pi(x) = \frac{1}{2\hat{\phi}} \pi_x(x)^2 + \frac{1}{2}\sigma^2 \pi_{xx}(x), \quad \forall x \in [\underline{x}^R, \bar{x}] \tag{22}$$

$$\pi(\underline{x}^R) = \pi(x^R) - L, \tag{23}$$

$$\pi_x(\underline{x}^R) = 0. \tag{24}$$

Resetting the project at a given progress level is akin to exercising an option. The

[13] Note that resetting is different from replacing the agent by paying a fixed cost and letting the new agent start the project from scratch. In the case of replacement, the principal also needs to pay the agent's continuation utility. Hence, given the structure of the model, the principal is better off keeping the agent but resetting the project.

return on this option (which costs L to exercise) is endogenous, as it depends both on the structure of the contract and on the resetting policy itself—because there is a positive probability that the project will drift back down to the reset threshold $\underline{x}^R$. Therefore, both $\underline{x}^R$ and the value of the contract at that point must be pinned down by the value matching and smooth pasting conditions. This characterization also immediately implies that $\underline{x}^R$ lies above the second-best termination threshold $\underline{x}$, since it must be the case that $\pi(\underline{x}^R) \geq \underline{v}$.

The first-best policy is given by the solution to (22)–(24) with $\hat{\phi}$ replaced by ϕ. As in Section 3, we denote the value function in the first-best problem by $v(x)$. In Proposition 10 below, we contrast the optimal resetting policies in the first- and in the second-best. We refer to a policy as more *failure tolerant* if it uses a lower (i.e., more permissive) resetting point $\underline{x}$.

Proposition 10

(i) For any restart point $x^R > \underline{x}$, there exist constants $0 < L_1 < L_2$ such that:

- *for all $L \in (L_1, L_2)$, the first-best policy uses resetting, but the second-best does not.*
- *for all $L \in (0, L_1)$ both policies use resetting, and for $L > L_2$, neither policy does.*

(ii) There exists $\varepsilon > 0$ such that for all $|x^R - \bar{x}| < \varepsilon$, the second-best policy is less failure-tolerant than the first-best, i.e., $\underline{x}^R > \underline{x}^{R,FB}$.

(iii) There exist $\varepsilon > 0$ and $\delta > 0$ such that if $\underline{v} < \varepsilon$ and $x^R \in (x^\dagger, x^\dagger + \delta)$, where $x^\dagger$ solves $\pi^{NR}(x^\dagger; \hat{\phi}) = \underline{v} + L$, then both policies use resetting and the second-best policy is more failure tolerant than the first best, i.e., $\underline{x}^R < \underline{x}^{R,FB}$.

(iv) When $\underline{x}^R < \underline{x}^{R,FB}$, there exists $x^* \in (\underline{x}^{R,FB}, \bar{x}]$ such that for all $x \in (\underline{x}^{R,FB}, x^*)$ the second-best policy implements higher effort than the first-best policy, i.e., $a(x) > a^{FB}(x)$; and for all $x > x^*$, the first-best policy implements higher effort than the second-best policy, i.e., $a(x) < a^{FB}(x)$.

Part (i) follows from Lemma 4 and standard comparative statics: for any restart point $x^R < \bar{x}$, the value function in the first-best problem is strictly above that in the second-best. It is then always possible to find a resetting cost large enough that the second-best policy never exercises the resetting option, and an even larger one such that the first-best policy does not exercise it either.

Parts (ii) and (iii) capture a more subtle logic. One could have imagined that the second-best policy is more likely to use resetting, which allows the project to jump to an exogenous level and to save the larger incentive-adjusted cost of effort. Part (ii) of

Proposition 10 captures this intuition: if the restart point x^R is close to the completion point $\bar{x}$, then the project is nearly complete after each reset. In this case, the second-best policy resets more often than the first-best, as resetting substitutes for effort and allows the principal to economize on variable costs. However, if the restart point is far from $\bar{x}$, then a new force dominates: even after a reset, much of the work remains to be done. Since the principal must induce effort at higher cost under hidden action, the value of resetting in the second-best solution is diminished. As a result, the principal uses the reset option less frequently than in the first-best policy.

Figure 4 plots the value function under the optimal contract and compares it to the first-best policy in two different cases.

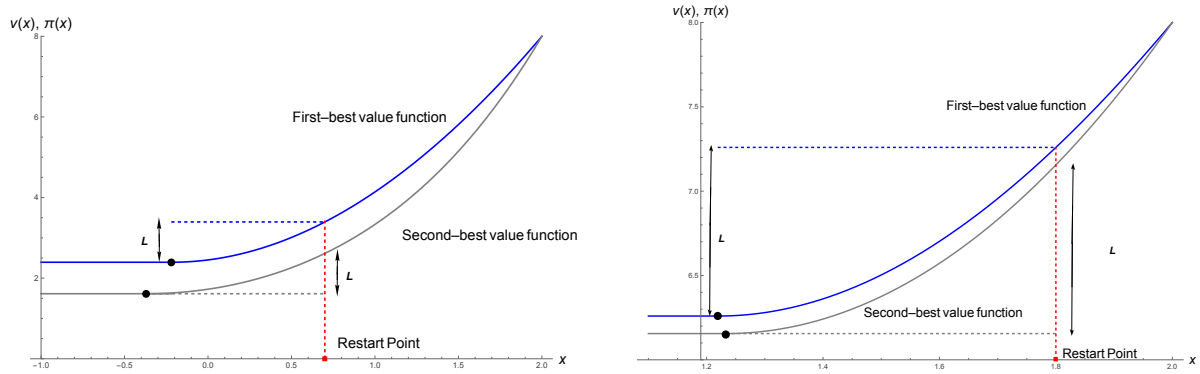


Figure 4: Restarting point: $x^R = 0.7$ (left panel), $x^R = 1.8$ (right panel). Other parameter values: $(b, L, \bar{x}, \sigma, \underline{v}, r, \eta, \phi) = (8, 1, 2, 2, 2, 2, 1/2, 1)$.

Therefore, Proposition 10 illustrates that a higher (incentives-adjusted) marginal cost of effort can be associated with lower expected fixed costs of resetting. The second-best policy finds resetting too costly not because the fixed cost L is larger, but because the complementary effort cost that must be incurred after a reset is higher under the second-best policy than in the first-best. This force discourages resets, even when they are socially efficient. As a result, the total expected cost under the principal's (second-best) contract may involve lower expenditure in resets, despite the higher marginal cost per unit of effort.

Finally, the ranking of the optimal effort levels between the first- and second-best policies also depends on the reset point: since effort is zero at the reset thresholds, it is possible for first-best effort to be lower than second-best effort near that point under the conditions of part (iii). Moral hazard scales the principal's marginal benefit of effort by $\phi/\hat{\phi} \in (0, 1)$, so relative to the first-best the principal has (i) a lower effective marginal return to raising a but (ii) possibly a lower reset point. As x grows, the first-best solution's slope catches up and then dominates because her HJB has the larger drift term (no κ penalty) while both problems share the same convexity structure. Once the first-best

effort policy crosses the second-best policy, the gap keeps widening: a second reversal would contradict the monotone and convex evolution implied by the HJBs.

6 Risk of Breakdowns

Many real-life innovation projects, particularly in the life sciences, face discrete, high-impact, and often irreversible setbacks. Clinical trials may be terminated due to adverse events; early-stage biotech programs may collapse upon a single scientific failure. These risks are endogenous: they may be influenced by the agent's choices, such as accelerating timelines or pursuing more aggressive protocols. In this section, we extend the baseline model to allow for such endogenous breakdowns.

We assume the project's progress X_t follows a Lévy process. In addition to choosing an effort level A_t, the agent can continuously adjust the riskiness of their approach $Q_t \in \{0,1\}$. We denote the low-risk action as $Q_t = 0$ and the high-risk action as $Q_t = 1$. The risky action boosts the drift term of the progress X_t by a constant $g > 0$ but increases the risk of a breakdown. There is a baseline arrival rate of a terminal breakdown $\lambda_0 > 0$. If the agent chooses the risky action, the arrival rate is $\lambda_0 + \lambda_r > 0$. Thus, breakdowns follow a Poisson process, denoted by N, with arrival intensity $\lambda_0 + \lambda_r Q_t$. After a breakdown, the project is terminated. Finally, we assume that the choice of risk is not observed by the principal, but the arrival of a breakdown is publicly observed. Thus, the principal can impose a penalty on the agent upon the arrival of a breakdown, which can help deter risk-taking behavior.

6.1 Observable Actions

We begin our analysis with the first-best benchmark. Thus, we assume that both the agent's effort and their choice of risk are observable to the principal. The principal's problem can be written as

$$\max_{A,Q,\tau} \mathbb{E}^{A,Q}\left[e^{-r(\tau\wedge\zeta)}\left(\mathbf{1}_{\{X_\tau\geq\bar{x},\ \tau<\zeta\}}b + \mathbf{1}_{\{X_\tau<\bar{x},\ \tau<\zeta\}}\underline{v} + \mathbf{1}_{\{\tau\geq\zeta\}}\underline{v}\right) - \frac{\phi}{2}\int_0^{\tau\wedge\zeta} e^{-rt}A_t^2 dt\right],$$

subject to

$$dX_t = (A_t + gQ_t)dt + \sigma dZ_t,$$

where the breakdown time ζ follows a Poisson process with instantaneous intensity $\Lambda_t = \lambda_0 + \lambda_r Q_t$.

The principal's problem is thus a standard optimal stopping problem. The value

function satisfies the following Hamilton-Jacobi-Bellman (HJB) equation:

$$rv(x) = \max_{a,q} \left\{ -\frac{\phi}{2}a^2 + (a+qg)v_x(x) + \frac{1}{2}\sigma^2 v_{xx}(x) + (\lambda_0 + \lambda_r q)(\underline{v} - v(x)) \right\}. \tag{25}$$

with value b if the project succeeds and $\underline{v}$ if it is terminated.

In order to characterize the first-best policy, we define the following operators:

$$\begin{aligned}
\mathcal{L}_{FB}^R v(x) &:= \frac{1}{r}\max_a \left\{ -\frac{\phi}{2}a^2 + (a+g)v_x(x) + \frac{1}{2}\sigma^2 v_{xx}(x) + (\lambda_0 + \lambda_r)(\underline{v} - v(x)) \right\}, \\
\mathcal{L}_{FB}^S v(x) &:= \frac{1}{r}\max_a \left\{ -\frac{\phi}{2}a^2 + av_x(x) + \frac{1}{2}\sigma^2 v_{xx}(x) + \lambda_0(\underline{v} - v(x)) \right\}.
\end{aligned}$$

Here, $\mathcal{L}_{FB}^R$ captures the evolution of the value function if the principal implements the risky action, and $\mathcal{L}_{FB}^S$ captures the evolution if she implements the safe action. Therefore, the value function satisfies

$$v(x) = \max\left\{ \mathcal{L}_{FB}^R v(x), \mathcal{L}_{FB}^S v(x), \underline{v} \right\},$$

with the boundary conditions $v(\underline{x}) = \underline{v}$, $v(\bar{x}) = b$, and the smooth pasting condition $v_x(\underline{x}) = 0$.

The optimal policy implements the risky action if and only if $\mathcal{L}_{FB}^R v(x) > \mathcal{L}_{FB}^S v(x)$, that is,

$$gv_x(x) + \lambda_r\left(\underline{v} - v(x)\right) > 0. \tag{26}$$

In general, the left-hand side of this inequality need not be monotonic in x, so the solution may switch between the risky and safe actions multiple times. For the rest of this subsection, we assume the following condition holds:

$$\lambda_r \geq \frac{g}{\sigma}\sqrt{2(\lambda_0 + r)}. \tag{27}$$

Condition (27) guarantees the cutoff structure, as we show in Proposition 11 below.[14]

Proposition 11 *Assume condition (27) holds.*

(i) The first-best policy is characterized by two regions: a high-risk region $x \in [\underline{x}^{FB}, x_c^{FB}]$, where $q(x) = 1$; and a low-risk region $x \in [x_c^{FB}, \bar{x}]$, where $q(x) = 0$.

(ii) Given the other parameters, if b is high enough, the cutoff x_c^{FB} is interior.

[14]Condition (27) is merely sufficient for the cutoff structure. Numerical analysis suggests that this property holds for a broader set of parameters.

(iii) Effort $a(x)$ is increasing in progress.

To understand the forces that determine the first-best risk-taking policy, consider the inequality (26). The principal trades off the additional drift benefit $gv_x(x)$ from taking the risky action against the higher expected loss from a breakdown, given by $\lambda_r(\underline{v} - v(x))$. As x increases, the marginal value of progress $v_x(x)$ rises because the project approaches its successful completion payoff b, making the drift boost more attractive. However, the loss from a breakdown, $(\underline{v} - v(x))$, also becomes larger as the project becomes more valuable. Under condition (27), this effect dominates as the project progresses: below x_c^{FB}, the drift boost is sufficiently valuable that the principal chooses the risky action; above x_c^{FB}, the marginal cost of failure becomes too large, and the safe action is preferred.

6.2 Unobservable Actions

We now derive the principal's optimal contract under endogenous risk-taking and unobservable actions.

The Martingale Representation Theorem for Lévy processes allows us to write the agent's continuation utility W under an incentive-compatible contract (A, C, Q, τ), which evolves according to the following stochastic differential equation (SDE)

$$dW_t = \big(rW_t - u(A_t, C_t)\big)\,dt + \beta_t(-\eta r W_t)(dX_t - (A_t + gQ_t)dt) \\ + \psi_t(-\eta r W_{t-})(dN_t - (\lambda_0 + \lambda_r Q_t)\,dt), \quad (28)$$

where β_t and ψ_t are progressively measurable processes.[15]

The following proposition characterizes the incentive compatibility constraint of the agent.

Proposition 12 *The agent's no-savings condition implies that $rW_t = u(A_t, C_t)$ and W_t is a martingale. The risky action is optimal if and only if*

$$-r\eta W_t\left(g\beta_t + \lambda_r\psi_t\right) \geq 0 \quad \Longleftrightarrow \quad \beta_t \geq -\frac{\lambda_r\psi_t}{g}. \quad (29)$$

The incentive compatibility condition for A_t is given by $A_t = \phi^{-1}\beta_t$. For the candidate optimal contract characterized in Proposition 13, these local conditions are also sufficient for global incentive compatibility.

[15] For example, see Theorem 11.2.8.1 of Jeanblanc, Yor, and Chesney (2009).

The incentive compatibility constraint ties the diffusion and the jump parts of the progress process. It is intuitive that the agent chooses the risky action if and only if the potential boost coming from the drift is higher than the potential punishment from the occurrence of the breakdown. A higher β_t—which induces higher effort—makes risk taking more valuable; hence, to deter the agent from taking risks, the principal needs to punish the agent after a breakdown occurs.

As in the baseline model, the principal's value function is additively separable as in (14), with the project value π now characterized as follows. Analogous to the first-best case, we define the operators $\mathcal{L}^R$ and $\mathcal{L}^S$, with the difference that incentive constraints for (A, Q) are embedded in these operators.[16] The principal's value function satisfies

$$\pi(x) = \max\left\{\mathcal{L}^S\pi(x), \mathcal{L}^R\pi(x), \underline{v}\right\},$$

with the boundary conditions $\pi(\underline{x}) = \underline{v}$, $\pi(\bar{x}) = b$, and the smooth pasting condition $\pi_x(\underline{x}) = 0$.

In our characterization result (Proposition 13 below), we assume that λ_r is high enough or σ is high enough. Formally, fixing the model parameters $(\lambda_0, \sigma, r, g, b, \eta)$, we show that there exists $\bar{\lambda}_r$ such that, for all $\lambda_r \geq \bar{\lambda}_r$, we have the cutoff structure. Likewise, fixing the model parameters $(\lambda_0, \lambda_r, r, g, b, \eta)$, we show that there exists $\bar{\sigma}$ such that for all $\sigma \geq \bar{\sigma}$, the cutoff structure also obtains.

Proposition 13 *Assume that λ_r is high enough or σ is high enough.*

(i) The optimal contract is characterized by two regions: a high-risk region $x \in [\underline{x}, x_c^{SB}]$, where $q(x) = 1$ and $\psi = 0$; and a low-risk region $x \in [x_c^{SB}, \bar{x}]$, where $q(x) = 0$ and $\psi < 0$. In the low-risk region, the punishment for breakdown is proportional to the optimal effort,

$$a = \phi^{-1}\beta = -\frac{\lambda_r \psi}{\phi g}.$$

(ii) The effort profile is discontinuous and non-monotonic in x: it is continuous and increasing up to the switching point, then jumps down, then is continuous and increasing up to $\bar{x}$.

(iii) Given the other parameters, if b is high enough, the cutoff x_c^{SB} is interior.

In the high-risk region, the agent is encouraged to take risks. The principal does not punish the agent for the occurrence of a breakdown. The principal could potentially incentivize risk taking by offering a bonus after the breakdown, yet does not use such a bonus in the optimal contract.

[16] For the formal definitions, see the proof of Proposition 13 in the Appendix.

The agent is punished for being unlucky in the low-risk region, even though he is known not to be taking the risky action. This punishment is necessary to deter the agent from taking risks. Penalties for breakdowns become more severe as the project gets closer to completion. This is because the principal prefers to implement higher effort near completion.

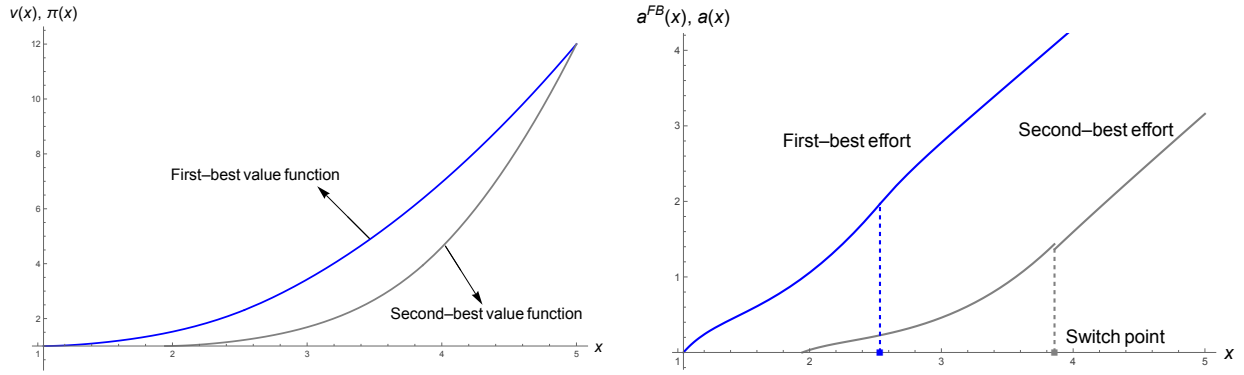


Figure 5: Value function (left) and effort profile (right) under endogenous risk-taking, comparing the first-best and second-best. Parameter values: $(b, \bar{x}, \sigma, \underline{v}, r, \eta, \phi, \lambda_0, \lambda_r, g) = (12, 5, 1, 1, 2, 1, 1, 1/2, 3, 2)$.

The optimal effort profile features a discontinuity at the switching point. Stronger incentives (higher β) are required to implement higher effort, but that also makes risk taking more attractive. The principal has two instruments to deter the agent from risk taking: (i) punishing the agent after a breakdown, and (ii) implementing lower effort. In the low-risk region, the principal uses both instruments jointly, leading to a non-monotonic effort path. Unlike in the first-best solution, the joint use of these instruments results in a discontinuous effort path. Due to the downward jump in action at the switching point, the principal may implement the highest level of effort in the interior of the domain. The final property is intuitive: as the project approaches completion, the potential loss from a breakdown is maximized. Figure 5 plots the value function and the effort profile of the second-best contract and compares them to the first-best.

7 Conclusion

This paper studies optimal contracting for projects with gradual progress and lumpy success. As such, our framework provides a tractable platform for studying dynamic incentive design in complex projects. Our results highlight several novel features of optimal project management. First, the presence of moral hazard leads the principal to be less failure tolerant and less ambitious than under observable actions. Second, the principal substitutes a fixed resetting cost for the high agency costs of recovering from

downturns and setbacks, but moral hazard can lead to more frequent or less frequent resets, depending on the expected amount of residual effort. Third, the principal uses discontinuous effort profiles and non-monotonic incentive schemes to deter risk taking or to speed up the project's progress.

Future research could extend our analysis along several promising directions. One possibility is to explore richer informational frictions, such as noisy or delayed feedback on project progress, which are common in real-world settings like R&D, drug development, and engineering design. Another is to introduce uncertainty over project quality, where the underlying drift is initially unknown and both parties learn from the realized progress. In such a setting, incentivizing effort, and hence progress, serves a dual role: it brings the project closer to completion and generates signals about the technology's viability.

Appendix

Savings and Continuation Values

We state and prove the analogue of Lemma 3 of He (2011) in the general model with endogenous breakdown risk, together with a companion statement used in the proof of Proposition 8.

Lemma A.1 *Fix a time $t \geq 0$. Let $W_t(S;P)$ denote the agent's continuation value when he holds savings $S \in \mathbb{R}$ and faces the remaining payments of a payment process P together with the effort and risk recommendations of the contract, that is,*

$$W_t(S;P) := \sup_{(\hat{A}_s,\hat{C}_s,\hat{Q}_s)_{s\geq t}} \mathbb{E}\left[\int_t^{\tilde{\tau}} e^{-r(s-t)}\, u(\hat{A}_s,\hat{C}_s)\, ds + e^{-r(\tilde{\tau}-t)}\, e^{-\eta r S_{\tilde{\tau}}}\, W_{\tilde{\tau}} \,\middle|\, \mathcal{F}_t\right], \tag{A.1}$$

where $W_{\tilde{\tau}}$ is the continuation value promised by the contract at termination, the supremum is over policies satisfying the savings equation and the no-Ponzi condition, and $S_t = S$. Then

(i) almost surely, $W_t(S;P) = e^{-\eta r S}\, W_t(0;P)$ for all S;

(ii) if P' is another payment process with $\int_t^\infty e^{-r(s-t)}\, dP'_s = \int_t^\infty e^{-r(s-t)}\, dP_s - D$ almost surely, for some $\mathcal{F}_t$-measurable D, then, almost surely, $W_t(S;P') = W_t(S-D;P)$ for all S.

Proof Part (i). The map $(\hat{A},\hat{C},\hat{Q}) \mapsto (\hat{A},\hat{C}+rS,\hat{Q})$ is a bijection between the policies feasible from savings 0 and those feasible from savings S: the two savings paths differ by the constant S, since their difference minus S solves $dG = rG\,dt$ with $G_t = 0$, so the no-Ponzi condition holds for one if and only if it holds for the other. The map leaves the law of $(X, N, \tilde{\tau}, W_{\tilde{\tau}})$ unchanged, as effort and risk choices are the same, and by

$$u(\hat{a},\hat{c}+\delta) = e^{-\eta\delta}\, u(\hat{a},\hat{c}) \qquad \text{for all } \delta \in \mathbb{R}, \tag{A.2}$$

together with $e^{-\eta r(S_{\tilde{\tau}}+S)} = e^{-\eta r S} e^{-\eta r S_{\tilde{\tau}}}$, it multiplies the payoff of every policy by $e^{-\eta r S}$. Taking suprema over the two feasible sets gives the claim.

Part (ii). The flow savings equation together with the no-Ponzi condition is equivalent to the intertemporal budget constraint, under which a consumption-effort policy is feasible from (S,P') if and only if it is feasible from $(S-D,P)$, path by path. The payoff of any given policy is the same in the two problems, so the suprema coincide.

Proof of Proposition 1

The following observation will be used repeatedly:

Lemma A.2 *If $\hat{x} \in [\underline{x}^{FB}, \bar{x}]$ satisfies $v_x(\hat{x}) = 0$ and $v_x(x) \geq 0 \; \forall x \in [\underline{x}^{FB}, \hat{x}]$, then $v_{xx}(\hat{x}) > 0$.*

Proof We have $v(\hat{x}) \geq v(\underline{x}^{FB}) = \underline{v} > 0$, where the first inequality follows since $v_x \geq 0$ on $[\underline{x}^{FB}, \hat{x}]$, and the equality is by value matching at $\underline{x}^{FB}$. Evaluating the HJB equation (7) at $\hat{x}$ gives $v_{xx}(\hat{x}) = 2\sigma^{-2} r v(\hat{x}) > 0$, since $a^{FB}(\hat{x}) = \phi^{-1} v_x(\hat{x}) = 0$ by the FOC (8).

Proof of part (i) Smooth pasting implies that $v_x(\underline{x}^{FB}) = 0$. Taking $\hat{x} = \underline{x}^{FB}$ in Lemma A.2 then shows that $v_x > 0$ on $(\underline{x}^{FB}, \underline{x}^{FB} + \epsilon)$ for some $\epsilon > 0$. Suppose toward a contradiction that the set $O := \{x \in [\underline{x}^{FB} + \epsilon, \bar{x}] : v_x(x) \leq 0\}$ is non-empty. Let $x^* := \inf O$. By continuity, $v_x(x^*) = 0$. Moreover, $v_x \geq 0$ on $[\underline{x}^{FB}, x^*]$ by definition of x^*. Taking $\hat{x} = x^*$ in Lemma A.2 then implies that v_x is strictly increasing at x^*, which contradicts x^* being the greatest lower bound of O. We conclude that O is empty. Hence, $v_x > 0$ on $(\underline{x}^{FB}, \bar{x}]$.

Proof of part (ii) Lemma A.2 and smooth pasting imply that $v_{xx}(\underline{x}^{FB}) > 0$. Suppose toward a contradiction that the set $O := \{x \in [\underline{x}^{FB}, \bar{x}] : v_{xx}(x) \leq 0\}$ is non-empty. Let $x^* := \inf O$. By continuity, $v_{xx}(x^*) = 0$ and thus $x^* > \underline{x}^{FB}$. Applying the envelope theorem of Milgrom and Segal (2002) to the HJB equation (7) and evaluating at x^* gives

$$r v_x(x^*) = \frac{1}{\phi} v_x(x^*) v_{xx}(x^*) + \frac{1}{2}\sigma^2 v_{xxx}(x^*) = \frac{1}{2}\sigma^2 v_{xxx}(x^*). \tag{A.3}$$

Therefore, we have $v_{xxx}(x^*) > 0$ by part (i). That is, v_{xx} is strictly increasing at x^*, contradicting x^* being the greatest lower bound of O. We conclude that O is empty.

Proof of part (iii) The result is immediate from the FOC (8) and part (ii).

Proof of part (iv) The first equality in (A.3) holds at every x in the continuation region, that is, $r v_x = \phi^{-1} v_x v_{xx} + \frac{1}{2}\sigma^2 v_{xxx}$. Under the first-best policy, X has drift $a^{FB}(X_t) = \phi^{-1} v_x(X_t)$, so Itô's formula gives, for $t < \tau$,

$$da^{FB}(X_t) = \phi^{-1}\left[v_{xx}(X_t)\, dX_t + \frac{1}{2}\sigma^2 v_{xxx}(X_t)\, dt\right] = r\, a^{FB}(X_t)\, dt + \phi^{-1}\sigma v_{xx}(X_t)\, dZ_t.$$

Hence $d\left(e^{-rt} a^{FB}(X_t)\right) = e^{-rt}\phi^{-1}\sigma v_{xx}(X_t)\, dZ_t$. Since v_{xx} is bounded on $[\underline{x}^{FB}, \bar{x}]$, the stopped process is a martingale.

Proof of Proposition 2

Proof of part (i) Fix parameters such that $b_1 > b_2$. The HJB equation (7) with the FOC (8) substituted in is autonomous in x: the right-hand side depends on v, v_x, and v_{xx} but not on x directly. Let g denote the unique solution to

$$rg = \frac{1}{2\phi}(g')^2 + \frac{1}{2}\sigma^2 g''$$

with initial conditions $g(0) = \underline{v}$ and $g'(0) = 0$. By Proposition 1, g is strictly increasing and strictly convex on $(0, \infty)$, and $g' > 0$ is strictly increasing.

By autonomy, each first-best value function is a horizontal translate of g: For $i = 1, 2$, we have $v_i(x) = g(x - \underline{x}_i^{FB})$ for $x \in [\underline{x}_i^{FB}, \bar{x}]$. Value matching $v_i(\bar{x}) = b_i$ requires

$$g(\bar{x} - \underline{x}_i^{FB}) = b_i.$$

Since g is strictly increasing and $b_1 > b_2$, it follows that $\bar{x} - \underline{x}_1^{FB} > \bar{x} - \underline{x}_2^{FB}$, i.e., $\underline{x}_1^{FB} < \underline{x}_2^{FB}$.

For the effort ranking, note that $a_i^{FB}(x) = \phi^{-1} v_{i,x}(x) = \phi^{-1} g'(x - \underline{x}_i^{FB})$. Since g' is strictly increasing and $x - \underline{x}_1^{FB} > x - \underline{x}_2^{FB}$, we have $a_1^{FB}(x) > a_2^{FB}(x)$ for all $x \in (\underline{x}_2^{FB}, \bar{x}]$. The comparison extends to $(\underline{x}_1^{FB}, \bar{x}]$ since project 2 is terminated in this region.

Proof of part (ii) Fix parameters such that $r_1 > r_2$. A mimicking argument implies that $v_1(x) \leq v_2(x)$ for all x and $\underline{x}_1^{FB} \geq \underline{x}_2^{FB}$.[17] We first show $\underline{x}_1^{FB} > \underline{x}_2^{FB}$. If $\underline{x}_1^{FB} = \underline{x}_2^{FB} =: \underline{x}$, then $v_{i,xx}(\underline{x}) = 2r_i\underline{v}/\sigma^2$ by the HJB. Since $r_1 > r_2$, we have $v_{1,xx}(\underline{x}) > v_{2,xx}(\underline{x})$, so $v_1 > v_2$ in a right neighborhood, contradicting $v_1 \leq v_2$.

Define the gap function $D(x) := v_2(x) - v_1(x)$. This function is continuous and smooth on $[\underline{x}_2^{FB}, \bar{x}]$, with the following properties: $D(\bar{x}) = 0$, $D(x) = 0$ for all $x \leq \underline{x}_2^{FB}$, $D(x) \geq 0$ for all $x \in (\underline{x}_2^{FB}, \bar{x})$, and $D(\underline{x}_1^{FB}) > 0$.

We show that D is single-peaked. Suppose not. Given the properties above, there must exist a local maximum at some point $x^\dagger$ followed by a local minimum at some point $x^\ddagger > x^\dagger$. At these points, $D_x(x^\dagger) = 0$, $D_x(x^\ddagger) = 0$, $D_{xx}(x^\dagger) \leq 0$, and $D_{xx}(x^\ddagger) \geq 0$. These conditions imply that $v_{2,x}(x^\dagger) = v_{1,x}(x^\dagger)$ with $v_{2,xx}(x^\dagger) \leq v_{1,xx}(x^\dagger)$, and $v_{2,x}(x^\ddagger) = v_{1,x}(x^\ddagger)$ with $v_{2,xx}(x^\ddagger) \geq v_{1,xx}(x^\ddagger)$.

Subtracting the HJB equations at $x^\dagger$ yields $r_1 v_1(x^\dagger) > r_2 v_2(x^\dagger)$. Applying the same logic at the local minimum gives $r_1 v_1(x^\ddagger) < r_2 v_2(x^\ddagger)$. Combining these two inequalities,

[17] Realized payoffs are not ranked path by path, as a higher discount rate also reduces the present value of the effort costs. However, because the expected continuation payoff under project 1's policy is bounded below by the salvage value $\underline{v} > 0$ at every history, a decrease in the discount rate raises the expected payoff at every history, and in particular at time zero.

we obtain

$$r_2\Big(v_2(x^\ddagger) - v_2(x^\dagger)\Big) > r_1\Big(v_1(x^\ddagger) - v_1(x^\dagger)\Big).$$

However, since $x^\dagger$ is a local maximum and $x^\ddagger$ is a subsequent local minimum of D, we have $D(x^\ddagger) < D(x^\dagger)$, which implies

$$0 \leq v_2(x^\ddagger) - v_2(x^\dagger) < v_1(x^\ddagger) - v_1(x^\dagger),$$

where non-negativity follows by monotonicity of v_2. Together the last two inequalities imply $r_1 < r_2$, a contradiction. We conclude that D is single-peaked on $[\underline{x}_2^{FB}, \bar{x}]$.

Since D is single-peaked with $D(\underline{x}_2^{FB}) = D(\bar{x}) = 0$ and $D > 0$ in between, D_x must change sign. At the peak, $D_x = v_{2,x} - v_{1,x} = 0$. By single-peakedness, $D_x > 0$ below the peak and $D_x < 0$ above. The cutoff structure for effort now follows from the FOC (8).

Proof of part (iii) Fix parameters such that $\sigma_1 > \sigma_2$. Since the value function is convex by Proposition 1, higher volatility raises the value of the continuation option, so $v_1(x) \geq v_2(x)$ and $\underline{x}_1^{FB} \leq \underline{x}_2^{FB}$ (see, e.g., Theorem 2.6′ in Camboni and Durandard, 2024).

Define the gap function $D(x) := v_1(x) - v_2(x)$ as in part (ii).

We first show $\underline{x}_1 < \underline{x}_2$, omitting the superscripts to simplify notation. Suppose toward a contradiction that $\underline{x}_1 = \underline{x}_2 =: \underline{x}$. Then $v_i(\underline{x}) = \underline{v}$ and $v_{i,x}(\underline{x}) = 0$ for $i = 1, 2$. Evaluating the HJB (7) at $\underline{x}$ gives $v_{i,xx}(\underline{x}) = 2r\underline{v}/\sigma_i^2$. Since $\sigma_1 > \sigma_2$, we have $v_{1,xx}(\underline{x}) < v_{2,xx}(\underline{x})$, so $v_2 > v_1$ in a right neighborhood of $\underline{x}$, contradicting $v_1 \geq v_2$.

Given $\underline{x}_1 < \underline{x}_2$, we have $D(x) = 0$ for all $x \leq \underline{x}_1$, $D(\underline{x}_2) > 0$ (since $v_1(\underline{x}_2) > \underline{v} = v_2(\underline{x}_2)$), $D(x) \geq 0$ for all x, and $D(\bar{x}) = 0$. Note that $D_x = v_{1,x} - v_{2,x} = \phi(a_1 - a_2)$ by (8). Since $v_{2,x}(\underline{x}_2) = 0$, we have $D_x(\underline{x}_2) = v_{1,x}(\underline{x}_2) > 0$, hence $a_1(\underline{x}_2) > a_2(\underline{x}_2) = 0$.

We next show $D_x(\bar{x}) < 0$. Since $D \geq 0$ and $D(\bar{x}) = 0$, we have $D_x(\bar{x}) \leq 0$. If $D_x(\bar{x}) = 0$, then $v_{1,x}(\bar{x}) = v_{2,x}(\bar{x})$ and the HJB equations at $\bar{x}$ give $\sigma_1^2 v_{1,xx}(\bar{x}) = \sigma_2^2 v_{2,xx}(\bar{x})$. Since $\sigma_1 > \sigma_2$, this implies $D_{xx}(\bar{x}) = v_{1,xx}(\bar{x}) - v_{2,xx}(\bar{x}) < 0$. Combined with $D(\bar{x}) = D_x(\bar{x}) = 0$, this gives $D < 0$ in a left neighborhood of $\bar{x}$, contradicting $D \geq 0$. Hence $D_x(\bar{x}) < 0$, i.e., $a_1(\bar{x}) < a_2(\bar{x})$.

Define $x^\dagger$ as the infimum of $\{x > \underline{x}_2 : D_x(x) = 0\}$ and $x^{\dagger\dagger}$ as the supremum of $\{x < \bar{x} : D_x(x) = 0\}$. Since $D_x(\underline{x}_2) > 0$ and $D_x(\bar{x}) < 0$, both thresholds exist and are interior. By construction, $D_x > 0$ on $(\underline{x}_2, x^\dagger)$ and $D_x < 0$ on $(x^{\dagger\dagger}, \bar{x})$, which gives $a_1(x) > a_2(x)$ for $x \in (\underline{x}_2, x^\dagger)$ and $a_1(x) < a_2(x)$ for $x \in (x^{\dagger\dagger}, \bar{x})$. For $x \in (\underline{x}_1, \underline{x}_2)$, $a_1(x) > 0 = a_2(x)$ since project 2 is terminated in this region.

Proof of part (iv) Fix parameters such that $\phi_1 > \phi_2$. A standard mimicking argument implies that $v_1(x) \leq v_2(x)$ for all x and $\underline{x}_1 \geq \underline{x}_2$. Moreover, the value ranking is strict

wherever project 1 continues: for $x > \underline{x}_1$, implementing project 1's optimal effort and stopping policy for project 2 yields

$$v_2(x) \;\geq\; v_1(x) + \mathbb{E}_x\left[\int_0^{\tau_1} e^{-rt}\,\frac{\phi_1 - \phi_2}{2}\left(a_1^{FB}(X_t)\right)^2 dt\right] \;>\; v_1(x),$$

where τ_1 denotes project 1's stopping time, and the second inequality holds because, started at $x > \underline{x}_1$, the process spends positive expected discounted time in the region where $a_1^{FB} > 0$ before τ_1.

We next show $\underline{x}_1 > \underline{x}_2$. Suppose toward a contradiction that $\underline{x}_1 = \underline{x}_2 =: \underline{x}$. For $i = 1, 2$, v_i is infinitely differentiable on $[\underline{x}, \bar{x}]$ and satisfies the HJB equation (7) with the FOC (8) substituted in; moreover, $v_i(\underline{x}) = \underline{v}$ and $v_{i,x}(\underline{x}) = 0$ by value matching and smooth pasting. Evaluating the resulting ODE and its first two derivatives at $\underline{x}$ gives

$$v_{i,xx}(\underline{x}) = \frac{2r\underline{v}}{\sigma^2}, \qquad v_{i,xxx}(\underline{x}) = 0, \qquad v_{i,xxxx}(\underline{x}) = \frac{2}{\sigma^2}\left(r\,v_{i,xx}(\underline{x}) - \frac{v_{i,xx}(\underline{x})^2}{\phi_i}\right).$$

Thus v_1 and v_2 agree at $\underline{x}$ up to third order, while their fourth derivatives satisfy

$$v_{2,xxxx}(\underline{x}) - v_{1,xxxx}(\underline{x}) = -\frac{2}{\sigma^2}\left(\frac{2r\underline{v}}{\sigma^2}\right)^2\left(\frac{1}{\phi_2} - \frac{1}{\phi_1}\right) < 0.$$

By Taylor's theorem, $v_2(x) < v_1(x)$ for all x in a right neighborhood of $\underline{x}$, contradicting the strict ranking established above. We conclude that $\underline{x}_1 > \underline{x}_2$.

Given $\underline{x}_1 > \underline{x}_2$, at $\underline{x}_1$ we have $a_1(\underline{x}_1) = 0 < a_2(\underline{x}_1)$. Suppose toward a contradiction that there exists a first crossing point $x^* > \underline{x}_1$ where $a_1(x^*) = a_2(x^*) =: a$. Since a_1 approaches a_2 from below, $a_{1,x}(x^*) \geq a_{2,x}(x^*)$. Equal efforts imply $v_{1,x}(x^*) = \phi_1 a > \phi_2 a = v_{2,x}(x^*)$. Subtracting the HJB equations at x^*,

$$r\Big(v_1(x^*) - v_2(x^*)\Big) = \tfrac{1}{2}(\phi_1 - \phi_2)\,a^2 + \tfrac{1}{2}\sigma^2\Big(v_{1,xx}(x^*) - v_{2,xx}(x^*)\Big).$$

The left-hand side is nonpositive (since $v_1 \leq v_2$) and the first term on the right is strictly positive (since $\phi_1 > \phi_2$ and $a > 0$), so $v_{1,xx}(x^*) < v_{2,xx}(x^*)$. Hence $a_{1,x}(x^*) = \phi_1^{-1} v_{1,xx}(x^*) < \phi_2^{-1} v_{2,xx}(x^*) = a_{2,x}(x^*)$, contradicting $a_{1,x}(x^*) \geq a_{2,x}(x^*)$.

Proof of Proposition 3

This is the special case of the global verification argument in the proof of Proposition 12 with no risk choice and no jump term.

Proof of Proposition 4

The second-best problem is the first-best problem with $\hat{\phi}$ in place of ϕ, and the proof of Proposition 1 uses only the form of the HJB equation. Parts (i), (ii), and (iv) therefore follow from the corresponding parts of Proposition 1, and part (iii) from $a = \hat{\phi}^{-1}\pi_x$ and $\beta = \phi a$.

Proof of Proposition 5

Proof of parts (i)–(iii) The reward enters the auxiliary problem (15) as it enters the first-best problem, and η and ϕ enter only through $\hat{\phi}$, which is increasing in each, so the three parts are Proposition 2(i) and (iv) applied to the auxiliary problem.

Proof of part (iv) Let $\underline{x}(r, \hat{\phi})$ denote the termination threshold of the auxiliary problem with discount rate r and cost parameter $\hat{\phi}$, so that $\underline{x}_i = \underline{x}(r_i, \hat{\phi}_i)$ with $\hat{\phi}_1 > \hat{\phi}_2$. Holding $\hat{\phi}$ fixed, a higher discount rate raises the threshold by the argument in the proof of Proposition 2(ii); holding r fixed, a higher cost parameter raises it by Proposition 2(iv). Hence $\underline{x}_1 = \underline{x}(r_1, \hat{\phi}_1) > \underline{x}(r_1, \hat{\phi}_2) > \underline{x}(r_2, \hat{\phi}_2) = \underline{x}_2$.

For effort, let $\theta_i := \pi_i / \hat{\phi}_i$, so that $a_i = \theta_{i,x}$ and, dividing the project value ODE by $\hat{\phi}_i$,

$$r_i \theta_i = \frac{1}{2} a_i^2 + \frac{1}{2}\sigma^2 a_{i,x} \qquad \text{on } [\underline{x}_i, \bar{x}]. \tag{A.4}$$

Both projects are active on $[\underline{x}_1, \bar{x}]$, where we define $\Lambda := r_1\theta_1 - r_2\theta_2$. By (A.4), at any point where $a_1 = a_2$,

$$a_{1,x} - a_{2,x} = \frac{2}{\sigma^2}\Lambda, \tag{A.5}$$

and differentiating (A.4) gives $a_{i,xx} = \frac{2}{\sigma^2}(r_i a_i - a_i a_{i,x})$, so that at any point where $a_1 = a_2 =: a$ and $a_{1,x} = a_{2,x}$,

$$a_{1,xx} - a_{2,xx} = \frac{2}{\sigma^2}(r_1 - r_2)\, a. \tag{A.6}$$

Finally, $\Lambda_x = r_1 a_1 - r_2 a_2 > 0$ wherever $a_1 \geq a_2 > 0$.

Since $a_1(\underline{x}_1) = 0 < a_2(\underline{x}_1)$, let $x^\ddagger := \inf\{x \in (\underline{x}_1, \bar{x}] : a_1(x) \geq a_2(x)\}$, with $x^\ddagger := \bar{x}$ if the set is empty. Then $a_1 < a_2$ on $(\underline{x}_1, x^\ddagger)$, hence on $(\underline{x}_2, x^\ddagger)$ since $a_1 = 0 < a_2$ on $(\underline{x}_2, \underline{x}_1]$, and it remains to show that $a_1 > a_2$ on $(x^\ddagger, \bar{x}]$ when $x^\ddagger < \bar{x}$. At $x^\ddagger$, $a_1 = a_2 =: a > 0$, and since $a_1 < a_2$ to the left, $a_{1,x} \geq a_{2,x}$, that is, $\Lambda \geq 0$ by (A.5). If $\Lambda(x^\ddagger) = 0$, then (A.6) implies $a_1 > a_2$ on a punctured neighborhood of $x^\ddagger$, contradicting $a_1 < a_2$ to the left; hence $\Lambda(x^\ddagger) > 0$ and $a_1 > a_2$ immediately to the right. Suppose $a_1 > a_2$ fails on $(x^\ddagger, \bar{x}]$, and let x^{**} be the infimum of the points where it fails. Then $a_1 > a_2$ on $(x^\ddagger, x^{**})$, so Λ is strictly increasing there and $\Lambda(x^{**}) > \Lambda(x^\ddagger) > 0$; but $a_1(x^{**}) = a_2(x^{**})$ with

$a_{1,x}(x^{**}) \leq a_{2,x}(x^{**})$ requires $\Lambda(x^{**}) \leq 0$ by (A.5), a contradiction.

Proof of Proposition 6

Proof of parts (ii) and (iii) The second-best problem is the first-best problem with $\hat{\phi} > \phi$, so both parts are Proposition 2(iv).

Proof of part (i) The proof of Proposition 2(iv) gives $\pi < v$ on $(\underline{x}, \bar{x})$, and on $(\underline{x}^{FB}, \underline{x}]$ we have $\pi(x) = \underline{v} < v(x)$. For the upper bound, implementing the first-best policy in the auxiliary problem (15) gives

$$\pi(x) \geq v(x) - \frac{\hat{\phi} - \phi}{2} \mathbb{E}_x \left[\int_0^{\tau^{FB}} e^{-rt} a^{FB}(X_t)^2 \, dt \right],$$

where τ^{FB} is the first-best stopping time. By the definition of v,

$$\frac{\phi}{2} \mathbb{E}_x \left[\int_0^{\tau^{FB}} e^{-rt} a^{FB}(X_t)^2 \, dt \right] = \mathbb{E}_x \left[e^{-r\tau^{FB}} \left(b \, \mathbf{1}_{\{X_{\tau^{FB}} = \bar{x}\}} + \underline{v} \, \mathbf{1}_{\{X_{\tau^{FB}} < \bar{x}\}} \right) \right] - v(x) \leq b - v(x).$$

Combining, $v(x) - \pi(x) \leq \frac{\hat{\phi} - \phi}{\phi} (b - v(x)) = (\kappa - 1)(b - v(x)) = \phi \eta r \sigma^2 (b - v(x))$.

Proof of Proposition 7

Set $Y_t := X_t^{FB} - X_t$ on $[0, \tau \wedge \tau^{FB}]$. The Brownian increments cancel, so Y is differentiable with $Y_0 = 0$ and $\dot{Y}_t = a^{FB}(X_t^{FB}) - a(X_t)$. If $Y_t = 0$, then $X_t^{FB} = X_t \in [\underline{x}, \bar{x}]$ and $\dot{Y}_t = a^{FB}(X_t) - a(X_t) > 0$ by Proposition 6(iii). Hence Y becomes positive immediately after $t = 0$ and cannot return to zero, which would require a nonpositive derivative at the point of return; so $Y_t > 0$ on $(0, \tau \wedge \tau^{FB}]$. Two consequences follow. First, $\tau^{FB} \geq \tau$ unless the first best completes before τ, because termination before τ would require $X_t^{FB} = \underline{x}^{FB} < \underline{x} \leq X_t < X_t^{FB}$. This gives $X^{FB}_{t \wedge \tau^{FB}} > X_t$ for $0 < t < \tau$. Second, on $\{X_\tau = \bar{x}\}$, if $\tau^{FB} \geq \tau$ then $X_\tau^{FB} > X_\tau = \bar{x}$, which is impossible; so the first best completed strictly before τ. On $\{X^{FB}_{\tau^{FB}} = \underline{x}^{FB}\}$, the first best did not complete before τ, so $\tau^{FB} \geq \tau$, with equality impossible since $X_\tau^{FB} > X_\tau \geq \underline{x} > \underline{x}^{FB}$; thus $\tau < \tau^{FB}$, and $X_\tau = \bar{x}$ is excluded by the previous case, so the second best terminated at τ. The ranking of completion probabilities follows from either inclusion.

Proof of Proposition 8

Under the scheme, the agent receives $d\Gamma_s$ until τ and his outside option thereafter, which is worth CE_0 at τ; his receipts thus have the same present value as the flow $r\,CE_0\,ds + \frac{\phi}{2}A_s^2\,ds + dCE_s$ paid forever. Integration by parts from t, using that CE is constant after τ, gives, for every $t \geq 0$, almost surely,[18]

$$\int_t^\infty e^{-r(s-t)}\left(r\,CE_0\,ds + \frac{\phi}{2}A_s^2\,ds + dCE_s\right) = \int_t^\infty e^{-r(s-t)}C_s\,ds - (CE_t - CE_0).$$

The first claim is Lemma A.1(ii) with $D = CE_t - CE_0$. At $t = 0$, with zero savings, the agent's problem thus coincides with his problem under the optimal contract, so the recommended effort and the consumption C are optimal for him. Given that effort, the present values of the principal's payments coincide along every path, so her expected payoff is unchanged. Following the recommended effort and consuming $C_s = \frac{\phi}{2}A_s^2 + r\,CE_s$ from zero savings, the agent's savings satisfy $dS_s = r(S_s - CE_s + CE_0)\,ds + dCE_s$, so $S_s = CE_s - CE_0$ while the project is active. Finally, by Lemma A.1, his continuation utility with savings S is $W_t(S - CE_t + CE_0; C) = e^{-\eta r(S - CE_t + CE_0)}W_t = e^{-\eta r S}u_0$, using $W_t = -\frac{1}{\eta r}e^{-\eta r CE_t}$ and $u_0 = -\frac{1}{\eta r}e^{-\eta r CE_0}$.

Proof of Proposition 9

Let $D := v - \pi \geq 0$, where the inequality holds by a mimicking argument. First, $\bar{x} \leq \bar{x}^{FB}$: otherwise $\bar{x}^{FB} \in (\underline{x}, \bar{x})$, so $\pi(\bar{x}^{FB}) > b(\bar{x}^{FB}) = v(\bar{x}^{FB})$, contradicting $D \geq 0$. Next, $\bar{x} < \bar{x}^{FB}$: if $\bar{x} = \bar{x}^{FB}$, then value matching and smooth pasting give $D(\bar{x}) = D_x(\bar{x}) = 0$, and the two HJB equations at $\bar{x}$ give $\frac{1}{2}\sigma^2(\pi_{xx} - v_{xx})(\bar{x}) = \frac{1}{2}b_x(\bar{x})^2(\phi^{-1} - \hat{\phi}^{-1}) > 0$, where the inequality uses $b_x > 0$, so $D_{xx}(\bar{x}) < 0$ and $D < 0$ to the left of $\bar{x}$, a contradiction. Finally, $\underline{x} > \underline{x}^{FB}$ follows from the proof of Proposition 2(iv), which does not use the upper endpoint of the continuation region, since $\hat{\phi} > \phi$.

Proof of Lemma 4

The first two claims follow from a verification argument. Write π^{NR} for $\pi^{NR}(\cdot;\hat{\phi})$, extended by $\underline{v}$ below $\underline{x}$ and by b above $\bar{x}$. Then $\pi^{NR} \geq \underline{v}$, and π^{NR} is continuously differentiable on $(-\infty, \bar{x}]$ and satisfies

$$r\pi^{NR}(x) \geq \max_a \left\{ -\frac{\hat{\phi}}{2}a^2 + a\pi_x^{NR}(x) + \frac{1}{2}\sigma^2\pi_{xx}^{NR}(x)\right\},$$

[18] By Girsanov's theorem, the probability measures induced by different effort processes are mutually absolutely continuous, so the identity holds almost surely under every admissible effort strategy.

with equality on $[\underline{x}, \bar{x}]$. Consider the discounted value along the path net of the costs paid so far,

$$M_t := e^{-rt}\pi^{NR}(X_t) - \int_0^t e^{-rs}\frac{\hat{\phi}}{2}A_s^2\,ds.$$

By Itô's formula, between resets,

$$dM_t = e^{-rt}\left[-r\pi^{NR}(X_t) + A_t\pi_x^{NR}(X_t) + \frac{1}{2}\sigma^2\pi_{xx}^{NR}(X_t) - \frac{\hat{\phi}}{2}A_t^2\right]dt + e^{-rt}\pi_x^{NR}(X_t)\sigma\,dZ_t,$$

where the drift is nonpositive by the inequality above, and zero under the optimal no-reset contract, while the stochastic integral is a martingale as π_x^{NR} is bounded. Hence M is a supermartingale between resets for any effort process, and a martingale under the optimal no-reset contract.

No resets if $L > \bar{L}(\hat{\phi})$. Then $\pi^{NR}(x^R) - L < \underline{v} \leq \pi^{NR}(x)$ for all x, so at every reset time $T_i < \tau$, $\pi^{NR}(X_t)$ falls by more than L. Thus, for any contract, $M_t - \sum_{T_i \leq t} e^{-rT_i}L$ is a supermartingale. Since $\pi^{NR}(X_\tau)$ dominates the stopping payoff, the principal's payoff is at most $M_0 = \pi^{NR}(x)$, with strict inequality if the contract resets before τ with positive probability. As $\pi^{NR}(x)$ is attained without resets, the optimal contract does not reset.

Resets if $L < \bar{L}(\hat{\phi})$. Modify the optimal no-reset contract so that, when progress first reaches $\underline{x}$, the project is reset rather than terminated, and the optimal no-reset contract is followed from x^R thereafter. This changes the principal's payoff by

$$\mathbb{E}\left[e^{-r\tau}\left(\pi^{NR}(x^R) - L - \underline{v}\right)\mathbf{1}_{\{X_\tau = \underline{x}\}}\right] > 0,$$

as $\underline{x}$ is reached with positive probability. Hence every contract without resets is strictly suboptimal.

Comparative static. If $\bar{L}(\hat{\phi}) > 0$, then x^R lies in the continuation region of the no-reset problem, where $\pi^{NR}(x^R; \cdot)$ is strictly decreasing in the cost parameter by the mimicking argument in the proof of Proposition 2(iv). Hence so is $\bar{L}(\cdot)$.

Proof of Proposition 10

Proof of part (i) Define $L_1 := \pi^{NR}(x^R; \hat{\phi}) - \underline{v}$ and $L_2 := \pi^{NR}(x^R; \phi) - \underline{v}$. Because $x^R > \underline{x}$, the restart point lies in the second-best no-reset continuation region, so $L_1 > 0$. The first-best no-reset value strictly exceeds the second-best no-reset value at x^R, and therefore $L_2 > L_1$. Lemma 4 implies that the second-best reset option is used when $L < L_1$, while the first-best reset option is used when $L < L_2$. The three stated ranges follow immediately.

Proof of part (ii) As $x^R \to \bar{x}$, the value matching condition $\pi(\underline{x}^R) = \pi(x^R) - L$ becomes

$\pi(\underline{x}^R) = b - L$, and likewise for v. The limiting problems are therefore the baseline problems of Sections 3 and 4 with salvage value $b - L$ in place of $\underline{v}$, for which $\underline{x} > \underline{x}^{FB}$ by Proposition 6(ii). Since the reset thresholds depend continuously on x^R (Theorem 2.8 of Teschl, 2012), the strict inequality persists for x^R close to $\bar{x}$.

Proof of part (iii) Consider a sequence of problems indexed by $n \in \mathbb{N}$ such that $x_n^R = (\pi^{NR})^{-1}(\underline{v}_n + L)$ and $\lim_{n\to\infty} \underline{v}_n = 0$. By Lemma 4, the first-best must be using resetting at $\underline{v} = 0$. Define $x^{R,*} := \lim_{n\to\infty}(\pi^{NR})^{-1}(\underline{v}_n + L)$. First we show that $\bar{x} - x^{R,*} < K$, that is, $x^{R,*}$ does not diverge to $-\infty$. To find the upper bound K, consider an auxiliary problem

$$rF(x) = \frac{1}{2\hat{\phi}}[F_x(x)]^2 + \frac{1}{2}\sigma^2 F_{xx}(x) \tag{A.7}$$

with the boundary conditions $F(0) = L$, $F_x(0) = 0$, and $K := F^{-1}(b)$, which is a finite number. At $x^{R,*}$, $\pi^{NR}(x^{R,*}) = L$ and $\pi_x^{NR}(x^{R,*}) > 0$. Therefore, given that the evolution of F and π is the same, we have $0 < \bar{x} - x^{R,*} < K$. To see this, note that since F is independent of the levels, consider the new problem starting at $x^{R,*}$ rather than 0. Suppose toward a contradiction that there exists $x^\dagger > x^{R,*}$ at which $F_x(x^\dagger) = \pi_x^{NR}(x^\dagger)$ for the first time. At $x^\dagger$, since F_x crosses $\pi_{,x}^{NR}$ from below, we have $F_{xx}(x^\dagger) \geq \pi_{xx}^{NR}(x^\dagger)$. Moreover, since $F_x < \pi_x^{NR}$ on $(x^{R,*}, x^\dagger)$ and $F(x^{R,*}) = \pi^{NR}(x^{R,*})$, we have $F(x^\dagger) < \pi^{NR}(x^\dagger)$. However, (A.7) evaluated at $x^\dagger$ with equal slopes gives

$$r\Big(\pi^{NR}(x^\dagger) - F(x^\dagger)\Big) = \frac{\sigma^2}{2}\Big(\pi_{xx}^{NR}(x^\dagger) - F_{xx}(x^\dagger)\Big),$$

so $\pi_{xx}^{NR}(x^\dagger) > F_{xx}(x^\dagger)$, contradicting $F_{xx}(x^\dagger) \geq \pi_{xx}^{NR}(x^\dagger)$. Hence no such $x^\dagger$ exists, and $F_x < \pi_x^{NR}$ on $(x^{R,*}, \bar{x})$, which implies π^{NR} reaches b before F does, giving $\bar{x} - x^{R,*} < K$.

At the exact second-best indifference locus $\pi_n^{NR}(x_n^R; \hat{\phi}) = \underline{v}_n + L$, the no-reset policy is optimal. Let $\underline{x}_n^{NR}$ denote its termination threshold. We first show that

$$\underline{x}_n^{NR} \longrightarrow -\infty. \tag{A.8}$$

Suppose instead that a subsequence of $\underline{x}_n^{NR}$ were bounded below. Since every threshold is below $\bar{x}$, a further subsequence would converge to a finite point $\underline{x}^*$. The corresponding value functions solve the autonomous ODE

$$r\pi = \frac{1}{2\hat{\phi}}\pi_x^2 + \frac{\sigma^2}{2}\pi_{xx}$$

with $\pi_n^{NR}(\underline{x}_n^{NR}) = \underline{v}_n \to 0$, and $\pi_{n,x}^{NR}(\underline{x}_n^{NR}) = 0$. Continuous dependence on initial data (see Theorem 2.8 of Teschl, 2012) implies that the limiting solution satisifes $\pi(\underline{x}^*) = 0$,

$\pi_x(\underline{x}^*) = 0$. The associated first-order ODE system is locally Lipschitz, so the unique solution with these initial conditions is $\pi \equiv 0$. This contradicts the completion boundary condition $\pi(\bar{x}) = b > 0$. Hence (A.8) holds.

By construction and continuity, $\pi^{NR}(x^{R,*}; \hat{\phi}) = L$ in the limiting problem with $\underline{v} = 0$. Since $L > 0$, $x^{R,*}$ lies strictly inside the second-best continuation region. The strict effort-cost comparison then gives

$$v^{NR}(x^{R,*}; \phi) > \pi^{NR}(x^{R,*}; \hat{\phi}) = L.$$

Consequently, $v^{NR}(x^{R,*}; \phi) - L > 0 = \underline{v}$, so Lemma 4 implies that the first-best strictly uses resetting in the limiting problem. Denote its reset threshold by $\underline{x}_*^{R,FB}$. This threshold is finite: its value-matching level is strictly positive, and the ODE solution starting from that positive level with zero slope reaches the restart and completion values over finite distances. Thus the two limiting boundaries satisfy $\underline{x}_n^{NR} \to -\infty$, and $\underline{x}_n^{R,FB} \to \underline{x}_*^{R,FB} > -\infty$. It follows that, for all sufficiently large n, $\underline{x}_n^{NR} < \underline{x}_n^{R,FB}$.

Finally, continuous dependence of the reset boundary-value problems on $(\underline{v}, x^R)$ implies that this strict inequality persists in a right neighborhood of the second-best reset-use locus. Therefore, there exist $\varepsilon > 0$ and $\delta > 0$ such that, whenever $\underline{v} \in (0, \varepsilon)$ and

$$0 < x^R - \left(\pi^{NR}(\cdot; \hat{\phi})\right)^{-1}(\underline{v} + L) < \delta,$$

both reset options are used and $\underline{x}^R < \underline{x}^{R,FB}$.

Proof of part (iv) When $\underline{x}_R^{FB} > \underline{x}_R$, we have $a(\underline{x}_R^{FB}) > a^{FB}(\underline{x}_R^{FB}) = 0$. The continuity of HJB equations implies that there exist a range $(\underline{x}_R, x^*)$, $a(x) > a^{FB}(x)$ for all $x \in (\underline{x}_R, x^*)$. At x^*, $a(x^*) = a^{FB}(x^*)$, hence we have

$$\hat{\phi}^{-1}\pi_x(x^*) = \phi^{-1}v_x(x^*).$$

Given that $\hat{\phi} > \phi$, we have $\pi_x(x^*) > v_x(x^*)$, hence $\hat{\phi}^{-1}(\pi_x(x^*))^2 > \phi^{-1}(v_x(x^*))^2$. Since $v(x^*) > \pi(x^*)$, then by inspection of the HJB equations at (x^*), it must be the case that $v_{xx}(x^*) > \pi_{xx}(x^*)$. Therefore, second-best effort can not cross the first-best effort. By the same HJB comparison, any subsequent crossing point $x^{**} > x^*$ would again yield $v_{xx}(x^{**}) > \pi_{xx}(x^{**})$. Hence no second crossing exists, and the first-best solution's effort crosses the second-best's effort at (x^*) and stays above for all $x > x^*$. Note that x^* is not necessarily interior.

Proof of Proposition 11

Proof of part (i) We start by defining the gap function $D^{FB}(x) := \mathcal{L}^R_{FB} v(x) - \mathcal{L}^S_{FB} v(x) = g v_x(x) + \lambda_r(\underline{v} - v(x))$. The first-best solution chooses the risky action at point x if and only if $D^{FB}(x) \geq 0$.

For the cutoff structure, suppose at x^* the first-best solution switches to the safe action, that is, $D^{FB}(x^*) = 0$ and the derivative with respect to x must be weakly negative:

$$g v_{xx}(x^*) - \lambda_r v_x(x^*) \leq 0, \tag{A.9}$$

otherwise the first-best solution would not switch to the safe action.

Next, we prove that after switching to the safe action, the first-best solution never switches back to the risky action. Using the safe-region HJB to substitute for $v_{xx}(x^*)$ and the condition $D^{FB}(x^*) = 0$ (i.e., $g v_x(x^*) = \lambda_r(v(x^*) - \underline{v})$) to eliminate $v_x(x^*)$, the left-hand side of (A.9) can be rewritten as

$$\frac{1}{g\sigma^2}\left(2g^2 v(x^*) r + \left(2g^2\lambda_0 - \lambda_r^2\sigma^2\right)\left(v(x^*) - \underline{v}\right) - \frac{\lambda_r^2}{\phi}(v(x^*) - \underline{v})^2\right).$$

Taking the derivative with respect to $v(x^*)$ yields

$$2g^2(\lambda_0 + r) - \lambda_r^2\sigma^2 - 2\frac{\lambda_r^2}{\phi}(v(x^*) - \underline{v}). \tag{A.10}$$

Therefore, $2g^2(\lambda_0 + r) - \lambda_r^2\sigma^2 < 0$ is a sufficient condition for (A.10) to be negative, since $v(x^*) - \underline{v} > 0$. This condition is sufficient because the value function is strictly increasing and convex; therefore, the above condition immediately implies that for a higher $x^{**} > x^*$, if $D^{FB}(x^{**}) = 0$, then $D^{FB}_x(x^{**}) < D^{FB}_x(x^*) < 0$. Hence, the first-best solution does not switch back to the risky action.

Proof of part (ii) We show that if b is high enough, the first-best solution implements the safe action near $\bar{x}$. Recall that the first-best solution implements the safe action if and only if $D^{FB}(x) < 0$, equivalently $g v_x(x) + \lambda_r(\underline{v} - v(x)) \leq 0$. Assume, toward a contradiction, that near $\bar{x}$ the first-best solution implements the risky action, which implies

$$g v_x(x) + \lambda_r(\underline{v} - v(x)) > 0.$$

From the risky-region HJB equation, using $v_{xx} \geq 0$ and $g v_x \geq 0$, we have the bound

$$v_x(x) \leq \sqrt{2\phi(r + \lambda_0 + \lambda_r) v(x)}.$$

Plugging it into the risk-taking condition yields

$$g\sqrt{2\phi(r+\lambda_0+\lambda_r)v(x)}+\lambda_r(\underline{v}-v(x)).$$

The first term grows as $O(\sqrt{v(x)})$ while the second term is $\lambda_r(\underline{v}-v(x))$, which is $O(-v(x))$ as $v(x)\to\infty$. Since $v(\bar{x})=b$, for b sufficiently large $v(x)$ is large near $\bar{x}$, making the expression negative. This is a contradiction.

Proof of part (iii) By the HJB equation, whenever $v_x(x)=0$, it follows that $v_{xx}(x)>0$.

In the risky region, by the envelope theorem, we have

$$(r+\lambda_0+\lambda_r)v_x(x)=(\frac{1}{\phi}v_x(x)+g)v_{xx}(x)+\frac{1}{2}\sigma^2 v_{xxx}(x). \tag{A.11}$$

Hence, whenever $v_{xx}(x)=0$, it must be the case that $v_{xxx}(x)>0$.

In the safe region, by the envelope theorem, we have

$$(r+\lambda_0)v_x(x)=\frac{1}{\phi}v_x(x)v_{xx}(x)+\frac{1}{2}\sigma^2 v_{xxx}(x). \tag{A.12}$$

Hence, whenever $v_{xx}(x)=0$, it must be the case that $v_{xxx}(x)>0$. Therefore, $a(x)$ must be increasing in x.

Proof of Proposition 12

Proof *Necessity.* Let ζ be the first jump time of N and let $\tilde{\tau}:=\tau\wedge\zeta$ denote the time at which the project ends. On $[0,\tilde{\tau}]$, the agent's lifetime utility conditional on time-t information is

$$V_t=\mathbb{E}^{A,Q}\left[\int_0^{\tilde{\tau}}e^{-rs}u(A_s,C_s)ds+e^{-r\tilde{\tau}}W_{\tilde{\tau}}\,\middle|\,\mathcal{F}_t\right]=\int_0^t e^{-rs}u(A_s,C_s)ds+e^{-rt}W_t.$$

Define the compensated processes $\tilde{N}_t:=N_t-\int_0^t(\lambda_0+\lambda_r Q_u)du$ and $dZ_t^{A,Q}:=dX_t-(A_t+gQ_t)dt$. By the martingale representation theorem for Lévy processes,[19] there exist $\mathcal{F}_t$-predictable processes (β_t,ψ_t), normalized by $-\eta r W_t$, such that

$$V_t=V_0+\int_0^t e^{-rs}\beta_s(-\eta r W_s)\,dZ_s^{A,Q}+\int_0^t e^{-rs}\psi_s(-\eta r W_{s-})\,d\tilde{N}_s.$$

[19] See Theorem 11.2.8.1 of Jeanblanc et al. (2009).

Taking the time derivative of V_t in both representations and solving for dW_t yields

$$dW_t = \big(rW_t - u(A_t, C_t)\big)\,dt + \beta_t(-\eta r W_t)(dX_t - (A_t + gQ_t)dt) + \psi_t(-\eta r W_{t-})(dN_t - (\lambda_0 + \lambda_r Q_t)dt).$$

For the no-savings condition, suppose the agent is given extra savings S at time t. By Lemma A.1, the optimal continuation policy is to follow the original effort choices and consume an extra flow rS, so $W_t(S) = e^{-\eta r S} W_t(0)$. The no-savings condition requires $u_C(A_t, C_t) = \frac{\partial}{\partial S} W_t(S)\Big|_{S=0} = -\eta r W_t(0)$, which gives $rW_t = u(A_t, C_t)$, i.e., condition (11).

For the IC, consider the recommended policy (A_t, Q_t) and any deviation $(\hat{A}_t, \hat{Q}_t)$. Under the deviation, the conditional drift of W_t is

$$(-\eta r W_t)\left[\beta_t(\hat{A}_t - A_t + g(\hat{Q}_t - Q_t)) + \lambda_r \psi_t(\hat{Q}_t - Q_t)\right] dt.$$

The agent maximizes the sum of flow utility and the drift of continuation utility. The best deviation in $\hat{Q}_t$ is to set $\hat{Q}_t = 1$ if and only if $(-\eta r W_t)(g\beta_t + \lambda_r \psi_t) \geq 0$. The FOC for $\hat{A}_t$ gives $u_A(\hat{A}_t, C_t) + \beta_t(-\eta r W_t) = 0$. Since $u_A = -\phi \hat{A}\, u_C$ and the Euler equation gives $u_C = -\eta r W_t$, we obtain $\hat{A}_t = \phi^{-1}\beta_t$.

Sufficiency. Fix the candidate optimal contract characterized in Proposition 13. Under any deviation $(\hat{A}, \hat{C}, \hat{Q})$, N has intensity $\lambda_0 + \lambda_r \hat{Q}_t$ and the state (X, S, W) evolves according to

$$\begin{aligned}
dX_t &= (\hat{A}_t + g\hat{Q}_t)\,dt + \sigma\, dZ_t, \\
dS_t &= (rS_t + C_t - \hat{C}_t)\,dt, \\
dW_t &= \beta_t(-\eta r W_t)\Big(\hat{A}_t - A_t + g(\hat{Q}_t - Q_t)\Big)\,dt + \beta_t(-\eta r W_t)\sigma\, dZ_t \\
&\qquad + \psi_t(-\eta r W_{t-})\Big(dN_t - (\lambda_0 + \lambda_r Q_t)\,dt\Big).
\end{aligned}$$

At termination (whether by success, failure, or breakdown), the no-Ponzi condition and consumption smoothing imply the agent's payoff given savings $S_{\tilde{\tau}}$ and contract value $W_{\tilde{\tau}}$ is $e^{-\eta r S_{\tilde{\tau}}} W_{\tilde{\tau}}$.

Consider the candidate value function $J(x, s, w) = e^{-\eta r s} w$. Since J is independent of x and linear in w, the generator associated with these state equations simplifies to

$$\mathcal{L}^{\hat{A},\hat{C},\hat{Q}} J = (rs + C - \hat{C})\, J_s + (-\eta r w)\Big[\beta(\hat{A} - A + g(\hat{Q} - Q)) + \lambda_r(\hat{Q} - Q)\psi\Big]\, J_w.$$

The agent's HJB equation is $rJ = \sup_{\hat{A},\hat{C},\hat{Q}}\{u(\hat{A}, \hat{C}) + \mathcal{L}^{\hat{A},\hat{C},\hat{Q}} J\}$. Substituting and optimizing:

FOC in $\hat{C}$: Setting $u_C - J_s = 0$ gives $\hat{C}^* = rs - \eta^{-1}\log(-\eta r w) + \frac{\phi}{2}(\hat{A})^2$.
FOC in $\hat{A}$: Setting $u_A(\hat{A}, \hat{C}) + \beta(-\eta r w)J_w = 0$ and using $u_C = J_s = -\eta r e^{-\eta r s} w$ gives $\hat{A}^* = \beta/\phi$.
FOC in $\hat{Q}$: The marginal benefit of choosing $\hat{Q} = 1$ is $-\eta r w(g\beta + \lambda_r \psi)J_w$. Since $J_w > 0$, its sign is the sign of $-\eta r w(g\beta + \lambda_r \psi)$.
The objective is concave in $(\hat{A}, \hat{C})$ and linear in the binary choice $\hat{Q}$, so these conditions give its global maximum. Under the local IC conditions, $(A, C + rs, Q)$ is a maximizer. The generator is then zero and, by (11), $u(A, C + rs) = e^{-\eta r s} u(A, C) = rJ$. Thus J satisfies the HJB equation. It remains to justify the verification argument for arbitrary deviations.

To establish global IC, consider any progressively measurable deviation $(\hat{A}, \hat{C}, \hat{Q})$ for which the state equations are well defined and

$$\mathbb{E}\left[\int_0^{\tilde{\tau}} e^{-rt} u(\hat{A}_t, \hat{C}_t) dt\right] > -\infty, \qquad \mathbb{E}\left[e^{-r\tilde{\tau}} e^{-\eta r S_{\tilde{\tau}}} W_{\tilde{\tau}}\right] > -\infty.$$

A deviation that violates either condition yields payoff $-\infty$ and is never optimal.

We first establish the bound needed for the localization argument. The recommended effort is bounded on the continuation region. In particular, there is a finite constant $\bar{A}$ such that

$$0 \le A_t \le \bar{A}, \qquad 0 \le \beta_t = \phi A_t \le \phi \bar{A}, \qquad -\bar{\psi} \le \psi_t \le 0, \qquad \bar{\psi} := \frac{\phi g}{\lambda_r} \bar{A}.$$

Indeed, the effort rules in the risky and safe regions are continuous functions of π_x, which is bounded on $[\underline{x}, \bar{x}]$. Moreover, $\psi = 0$ in the risky region and $\psi = -\phi g A/\lambda_r$ in the safe region. Define $m := \frac{\eta r \phi \bar{A}^2}{2} + \eta r g \phi \bar{A} + \eta r \bar{\psi}(\lambda_0 + \lambda_r) + \frac{\eta^2 r^2 \phi^2 \bar{A}^2 \sigma^2}{2}$ and

$$K := \exp\left\{\frac{2}{e} + \frac{\eta \phi \bar{A}^2}{2} + \frac{m}{r}\right\}. \tag{A.13}$$

We show below that, for every stopping time $t \le \tilde{\tau}$,

$$e^{-rt} e^{-\eta r S_t} W_t \ge K\,\mathbb{E}\left[\int_t^{\tilde{\tau}} e^{-rs} u(\hat{A}_s, \hat{C}_s) ds + e^{-r\tilde{\tau}} e^{-\eta r S_{\tilde{\tau}}} W_{\tilde{\tau}} \,\middle|\, \mathcal{F}_t\right]. \tag{A.14}$$

The state equation for W is multiplicative, and a breakdown multiplies W_{t-} by $1 - \eta r \psi_t \ge 1$. Hence W remains negative and $L_t := \log(-\eta r W_t)$ is well defined. By the consumption rule (11), $C_t = \frac{\phi}{2} A_t^2 - \frac{1}{\eta} L_t$.

First suppose that savings at time t are zero, and set $p := 1 - e^{-r(\tilde{\tau} - t)}$. Since $\hat{A}_s + g\hat{Q}_s \ge 0$ for all s, Lemma S.1 implies $\tilde{\tau} < \infty$ almost surely, so $p < 1$. On the event $p = 0$,

termination is immediate and the desired bound follows from $K \geq 1$. Consider the event $p \in (0,1)$. Integrating the savings equation gives

$$\int_t^{\tilde{\tau}} e^{-r(s-t)} \left(\hat{C}_s - \frac{\phi}{2} \hat{A}_s^2 \right) ds = \int_t^{\tilde{\tau}} e^{-r(s-t)} C_s ds - (1-p) S_{\tilde{\tau}} - \int_t^{\tilde{\tau}} e^{-r(s-t)} \frac{\phi}{2} \hat{A}_s^2 ds.$$

Since $\int_t^{\tilde{\tau}} e^{-r(s-t)} ds = p/r$, Jensen's inequality and the CARA utility function imply

$$\begin{aligned} &\int_t^{\tilde{\tau}} e^{-r(s-t)} u(\hat{A}_s, \hat{C}_s) ds \\ &\quad \leq -\frac{p}{\eta r} \exp\left\{ -\frac{\eta r}{p} \left[\int_t^{\tilde{\tau}} e^{-r(s-t)} C_s ds - (1-p) S_{\tilde{\tau}} - \int_t^{\tilde{\tau}} e^{-r(s-t)} \frac{\phi}{2} \hat{A}_s^2 ds \right] \right\}. \end{aligned}$$

Adding the terminal utility and maximizing the right-hand side over $S_{\tilde{\tau}}$ yields

$$\begin{aligned} &\int_t^{\tilde{\tau}} e^{-r(s-t)} u(\hat{A}_s, \hat{C}_s) ds + (1-p) e^{-\eta r S_{\tilde{\tau}}} W_{\tilde{\tau}} \\ &\quad \leq -\left(\frac{p}{\eta r} \right)^p \left(\frac{(1-p) e^{L_{\tilde{\tau}}}}{\eta r} \right)^{1-p} \exp\left\{ -\eta r \int_t^{\tilde{\tau}} e^{-r(s-t)} C_s ds + \eta r \int_t^{\tilde{\tau}} e^{-r(s-t)} \frac{\phi}{2} \hat{A}_s^2 ds \right\}. \end{aligned} \tag{A.15}$$

To verify this step, write the expression being maximized as $-ae^{by} - de^{-\eta r y}$, where $y = S_{\tilde{\tau}}$, $b = \eta r(1-p)/p$, $d = (1-p)e^{L_{\tilde{\tau}}}/(\eta r)$, and a is the remaining positive factor. The minimum of the corresponding positive sum is

$$a^p d^{1-p} \left[\left(\frac{p}{1-p} \right)^{1-p} + \left(\frac{1-p}{p} \right)^p \right],$$

where $p = \eta r/(b + \eta r)$. The expression in brackets is at least one: if $(1-p)/p \geq 1$, its second term is at least one, and otherwise its first term is at least one. Taking the negative and dropping this term gives (A.15).

Substituting the consumption rule into (A.15) and noting that $p^p \geq e^{-1/e}$ and $(1-p)^{1-p} \geq e^{-1/e}$, gives

$$\begin{aligned} &\int_t^{\tilde{\tau}} e^{-r(s-t)} u(\hat{A}_s, \hat{C}_s) ds + e^{-r(\tilde{\tau}-t)} e^{-\eta r S_{\tilde{\tau}}} W_{\tilde{\tau}} \\ &\quad \leq -\frac{e^{-2/e}}{\eta r} e^{-\eta \phi \bar{A}^2/2} \exp\left\{ r \int_t^{\tilde{\tau}} e^{-r(s-t)} L_s ds + (1-p) L_{\tilde{\tau}} + \eta r \int_t^{\tilde{\tau}} e^{-r(s-t)} \frac{\phi}{2} \hat{A}_s^2 ds \right\}. \end{aligned} \tag{A.16}$$

The factor $e^{-\eta\phi\bar{A}^2/2}$ follows from $\eta r\int_t^{\tilde{\tau}} e^{-r(s-t)}\frac{\phi}{2}A_s^2 ds \le \frac{\eta\phi\bar{A}^2}{2}$.

It remains to bound the exponent in (A.16). Before termination, N has not jumped. Applying Itô's formula to $L_s = \log(-\eta r W_s)$ using the state equation for W gives

$$dL_s = \left[-\eta r\beta_s\Big(\hat{A}_s - A_s + g(\hat{Q}_s - Q_s)\Big) + \eta r\psi_s(\lambda_0 + \lambda_r Q_s) - \frac{\eta^2 r^2\beta_s^2\sigma^2}{2}\right] ds \\ - \eta r\beta_s\sigma\, dZ_s.$$

If termination is caused by a breakdown, then $\Delta L_{\tilde{\tau}} = \log(1-\eta r\psi_{\tilde{\tau}}) \ge 0$; otherwise L does not jump at termination. Integration by parts and $1-p = e^{-r(\tilde{\tau}-t)}$ give

$$r\int_t^{\tilde{\tau}} e^{-r(s-t)}L_s ds + (1-p)L_{\tilde{\tau}} - L_t = \int_t^{\tilde{\tau}} e^{-r(s-t)}dL_s.$$

After adding the effort term in (A.16), the drift in the last integral is bounded below by

$$\eta r\left(\frac{\phi}{2}\hat{A}_s^2 - \beta_s\hat{A}_s\right) + \eta r\beta_s A_s - \eta r g\beta_s(\hat{Q}_s - Q_s) + \eta r\psi_s(\lambda_0 + \lambda_r Q_s) - \frac{\eta^2 r^2\beta_s^2\sigma^2}{2} \ge -m.$$

Indeed, $\frac{\phi}{2}\hat{A}_s^2 - \beta_s\hat{A}_s \ge -\beta_s^2/(2\phi)$, the term $\beta_s A_s$ is non-negative, $|\hat{Q}_s - Q_s| \le 1$, and the remaining bounds follow from $\beta_s \le \phi\bar{A}$ and $-\bar{\psi} \le \psi_s \le 0$. Hence

$$r\int_t^{\tilde{\tau}} e^{-r(s-t)}L_s ds + (1-p)L_{\tilde{\tau}} + \eta r\int_t^{\tilde{\tau}} e^{-r(s-t)}\frac{\phi}{2}\hat{A}_s^2 ds \\ \ge L_t - \frac{m}{r} - \eta r\sigma\int_t^{\tilde{\tau}} e^{-r(s-t)}\beta_s\, dZ_s. \tag{A.17}$$

The stochastic integral in (A.17) has zero conditional mean. In fact, its conditional quadratic variation is bounded by $\eta^2 r^2\sigma^2\phi^2\bar{A}^2\int_t^{\infty} e^{-2r(s-t)}ds = \frac{\eta^2 r\sigma^2\phi^2\bar{A}^2}{2} < \infty$. Combining (A.16) and (A.17) and applying conditional Jensen's inequality therefore gives

$$\begin{aligned}
&\mathbb{E}\left[\int_t^{\tilde{\tau}} e^{-r(s-t)}u(\hat{A}_s,\hat{C}_s)ds + e^{-r(\tilde{\tau}-t)}e^{-\eta r S_{\tilde{\tau}}}W_{\tilde{\tau}}\,\middle|\,\mathcal{F}_t\right] \\
&\le -\frac{e^{-2/e}}{\eta r}e^{-\eta\phi\bar{A}^2/2}e^{L_t - m/r}\mathbb{E}\left[\exp\left\{-\eta r\sigma\int_t^{\tilde{\tau}} e^{-r(s-t)}\beta_s\, dZ_s\right\}\,\middle|\,\mathcal{F}_t\right] \\
&\le e^{-2/e-\eta\phi\bar{A}^2/2-m/r}W_t = \frac{W_t}{K},
\end{aligned}$$

where the last equality follows from (A.13). This proves the required bound when $S_t = 0$.

For $S_t \neq 0$, shift $(S_s, \hat{C}_s)$ to $(S_s - S_t, \hat{C}_s - rS_t)$. By Lemma A.1(i), this leaves the

other states unchanged and multiplies continuation utility by $e^{-\eta r S_t}$. Thus the preceding inequality, multiplied by $Ke^{-rt}e^{-\eta r S_t}$, gives (A.14).

We now complete the localization argument. Applying the Itô–Lévy formula and the HJB inequality shows that

$$\int_0^{t\wedge\tilde{\tau}} e^{-rs}u(\hat{A}_s,\hat{C}_s)ds + e^{-r(t\wedge\tilde{\tau})}J(X_{t\wedge\tilde{\tau}},S_{t\wedge\tilde{\tau}},W_{t\wedge\tilde{\tau}})$$

is a local supermartingale. More explicitly, let

$$\rho_n := \inf\{t \geq 0 : |S_t| \geq n \text{ or } |W_t| \geq n\} \wedge n.$$

On $[0,\rho_n]$, the stochastic integrands are bounded, so the stopped stochastic integrals are true martingales. Thus the process stopped at ρ_n is a supermartingale.

Set

$$Y := \int_0^{\tilde{\tau}} e^{-rs}u(\hat{A}_s,\hat{C}_s)ds + e^{-r\tilde{\tau}}e^{-\eta r S_{\tilde{\tau}}}W_{\tilde{\tau}}.$$

The finite-value condition gives $Y \in L^1$. The absolute value of the accumulated flow utility is bounded by $|Y|$, while (A.14) bounds the absolute value of the continuation term by $K\mathbb{E}[|Y| \mid \mathcal{F}_t]$ at every stopping time $t \leq \tilde{\tau}$. Since $Y \in L^1$, $\{\mathbb{E}[|Y| \mid \mathcal{F}_t] : t \leq \tilde{\tau}\}$ is uniformly integrable. Thus the local supermartingale is of class (D), and the inequalities stopped at ρ_n converge in L^1 as $n \to \infty$. Thus, for $0 \leq s \leq t$,

$$\begin{aligned}\mathbb{E}\left[\int_0^{t\wedge\tilde{\tau}} e^{-rv}u(\hat{A}_v,\hat{C}_v)dv + e^{-r(t\wedge\tilde{\tau})}J(X_{t\wedge\tilde{\tau}},S_{t\wedge\tilde{\tau}},W_{t\wedge\tilde{\tau}}) \,\middle|\, \mathcal{F}_s\right] \\ \leq \int_0^{s\wedge\tilde{\tau}} e^{-rv}u(\hat{A}_v,\hat{C}_v)dv + e^{-r(s\wedge\tilde{\tau})}J(X_{s\wedge\tilde{\tau}},S_{s\wedge\tilde{\tau}},W_{s\wedge\tilde{\tau}}).\end{aligned}$$

Finally, since $\tilde{\tau} < \infty$ almost surely, the same class (D) bound gives L^1 convergence as $t \to \infty$. Therefore,

$$\mathbb{E}\left[\int_0^{\tilde{\tau}} e^{-rs}u(\hat{A}_s,\hat{C}_s)ds + e^{-r\tilde{\tau}}e^{-\eta r S_{\tilde{\tau}}}W_{\tilde{\tau}}\right] \leq W_0. \tag{A.18}$$

Obedience attains equality by promise keeping, so the contract is globally incentive compatible.

Proof of Proposition 13

Throughout, let $\hat{\phi} := \phi\kappa$, $\Lambda := r + \lambda_0 + \lambda_r$, and $\lambda(q) := \lambda_0 + \lambda_r q$, write $z := \pi_x(x)$ whenever convenient, and abbreviate the binding IC ratio as $\mu := \phi g/\lambda_r$. We first establish the properties of π that the argument relies on, mirroring Proposition 4(i)–(ii).

Write the HJB equation below as $r\pi = H(\pi, \pi_x) + \frac{1}{2}\sigma^2\pi_{xx}$, where $H(\pi, z)$ is the maximum over (a, ψ, q), subject to incentive compatibility, of the terms not involving π_{xx}. Three properties of H will be used. First, $H(\pi, 0) = \lambda_0(\underline{v} - \pi)$ for $\pi \geq \underline{v}$: at $z = 0$ the effort and punishment terms are maximized at $a = \psi = 0$, and the lower hazard is then preferred. Second, H is a maximum of finitely many continuously differentiable functions, so it has one-sided partial derivatives, with $H_\pi = -\lambda(q^*)$ and $H_z = a^* + q^* g \geq 0$ at the maximizer (a^*, ψ^*, q^*) by the envelope theorem. Third, π is twice continuously differentiable on $[\underline{x}, \bar{x}]$, as both branches of the HJB equation contain the same second-order term.

Lemma A.3 *Under the optimal contract, $\pi_x > 0$ on $(\underline{x}, \bar{x}]$ and $\pi_{xx} > 0$ on $[\underline{x}, \bar{x}]$.*

Proof If $\hat{x} \in [\underline{x}, \bar{x}]$ satisfies $\pi_x(\hat{x}) = 0$ and $\pi_x \geq 0$ on $[\underline{x}, \hat{x}]$, then $\pi(\hat{x}) \geq \underline{v}$, and the HJB equation at $\hat{x}$ gives $\frac{1}{2}\sigma^2\pi_{xx}(\hat{x}) = r\pi(\hat{x}) + \lambda_0(\pi(\hat{x}) - \underline{v}) \geq r\underline{v} > 0$ by the first property of H. This is the analog of Lemma A.2, and $\pi_x > 0$ on $(\underline{x}, \bar{x}]$ follows as in the proof of Proposition 1(i). For convexity, taking $\hat{x} = \underline{x}$ gives $\pi_{xx}(\underline{x}) > 0$. Suppose that $O := \{x \in [\underline{x}, \bar{x}] : \pi_{xx}(x) \leq 0\}$ is nonempty and let $x^* := \inf O$, so that $x^* > \underline{x}$, $\pi_{xx}(x^*) = 0$, and $\pi_{xx} > 0$ on $[\underline{x}, x^*)$. Differentiating the HJB equation from the left at x^* and using the second property of H,

$$\tfrac{1}{2}\sigma^2\pi_{xxx}(x^*-) = r\pi_x(x^*) - H_\pi\pi_x(x^*) - H_z\pi_{xx}(x^*) = (r + \lambda(q^*))\,\pi_x(x^*) > 0,$$

so $\pi_{xx} < 0$ on a left neighborhood of x^*, a contradiction.

Proof of parts (i) and (ii) Using the guess $f(x, w) = \pi(x) + \frac{1}{\eta r}\ln(-\eta r w)$, the consumption rule, and $f_{xw} = 0$, the principal's HJB becomes

$$r\pi(x) = \max_{a,\beta,\psi,q} \left\{ -\tfrac{1}{2}\hat{\phi}\,a^2 + (a + qg)\pi_x + (\lambda_0 + \lambda_r q)\left(\underline{v} - \pi + \psi + \tfrac{1}{\eta r}\ln(1 - \eta r\psi)\right) + \tfrac{1}{2}\sigma^2\pi_{xx} \right\},$$

subject to $a = \phi^{-1}\beta$ and the IC constraint $g\beta + \lambda_r\psi \geq 0$ when $q = 1$, with $g\beta + \lambda_r\psi \leq 0$ when $q = 0$.

We define the following operators :

$$\mathcal{L}^R\pi := \frac{1}{r}\max_{a,\psi}\left\{-\tfrac{\hat{\phi}}{2}a^2 + (a+g)z + \tfrac{1}{2}\sigma^2\pi_{xx} + (\lambda_0+\lambda_r)\Big(\underline{v} - \pi + \psi + \tfrac{1}{\eta r}\ln(1-\eta r\psi)\Big)\right\},$$
$$\mathcal{L}^S\pi := \frac{1}{r}\max_{a,\psi}\left\{-\tfrac{\hat{\phi}}{2}a^2 + az + \tfrac{1}{2}\sigma^2\pi_{xx} + \lambda_0\Big(\underline{v} - \pi + \psi + \tfrac{1}{\eta r}\ln(1-\eta r\psi)\Big)\right\},$$

with the IC constraints as indicated. a^s denotes the optimal effort in safe region and a^R the optimal effort in risky region.

Risky region. The FOC for ψ gives $(\lambda_0+\lambda_r)(1-\frac{1}{1-\eta r\psi}) = 0$, hence $\psi^R = 0$. The FOC for a gives $a^R = \hat{\phi}^{-1}z$. The IC constraint $g\phi a + \lambda_r\psi \geq 0$ becomes $g\phi\,\hat{\phi}^{-1}z \geq 0$, which holds since $z \geq 0$. Substituting:

$$r\mathcal{L}^R\pi = \tfrac{1}{2}\hat{\phi}^{-1}z^2 + gz + \tfrac{1}{2}\sigma^2\pi_{xx} + (\lambda_0+\lambda_r)(\underline{v} - \pi). \tag{A.19}$$

Safe region. We first show the IC binds. Suppose toward a contradiction that the IC is slack: $g\phi a + \lambda_r\psi < 0$ with $a \geq 0$. This implies $\psi < -\mu a \leq 0$. But the function $\lambda_0(\psi + \frac{1}{\eta r}\ln(1-\eta r\psi))$ is increasing in ψ on $(-\infty, 0]$. So increasing ψ to $-\mu a$ (the binding IC value) strictly improves the objective without violating any constraint. Hence the IC must bind: $\psi = -\mu a$.

Substituting $\psi = -\mu a$ into $\mathcal{L}^S$ and differentiating with respect to a, the FOC is:

$$\hat{\phi}\,a - z + \frac{\lambda_0\eta r\mu^2 a}{1+\eta r\mu\,a} = 0. \tag{A.20}$$

Multiplying by $(1+\eta r\mu\,a)$:

$$\hat{\phi}\eta r\mu\;a^2 + B(z)\,a - z = 0, \qquad B(z) := \hat{\phi} + \lambda_0\eta r\mu^2 - \eta r\mu\,z. \tag{A.21}$$

By the quadratic formula, the positive root is:

$$a^S(z) = \frac{-B(z) + y(z)}{2\,\hat{\phi}\,\eta r\,\mu}, \qquad y(z) := \sqrt{B(z)^2 + 4\,\hat{\phi}\,\eta r\,\mu\,z}. \tag{A.22}$$

Equivalently, from the FOC (A.20):

$$a^S = \hat{\phi}^{-1}\left[z - \lambda_0\mu\left(1 - \frac{1}{1+\eta r\mu\,a^S}\right)\right]. \tag{A.23}$$

Since $\lambda_0\mu(1-\frac{1}{1+\eta r\mu a}) > 0$ for $a > 0$, we have $a^S(z) < \hat{\phi}^{-1}z = a^R(z)$ whenever $\lambda_0 > 0$ and $z > 0$.

Substituting into $\mathcal{L}^S$:

$$r\mathcal{L}^S\pi = -\tfrac{\hat{\phi}}{2}(a^S)^2 + a^S z + \tfrac{1}{2}\sigma^2\pi_{xx} + \lambda_0\Big(\underline{v} - \pi - \mu a^S + \tfrac{1}{\eta r}\ln(1+\eta r\mu\, a^S)\Big). \tag{A.24}$$

Define the gap function $D^{SB}(x) := \mathcal{L}^R\pi(x) - \mathcal{L}^S\pi(x)$, and define:

$$\Psi(a) := \lambda_0\left(\mu\, a - \frac{1}{\eta r}\ln(1+\eta r\mu\, a)\right) \geq 0,$$
$$G(z) := \tfrac{\hat{\phi}}{2}(a^S)^2 - a^S z + \Psi(a^S) \leq 0.$$

Taking (A.19) minus (A.24):

$$rD^{SB} = \tfrac{1}{2}\hat{\phi}^{-1}z^2 + gz + \lambda_r(\underline{v} - \pi) + G(z). \tag{A.25}$$

Note: the $\frac{1}{2}\sigma^2\pi_{xx}$ and $\lambda_0(\underline{v} - \pi)$ terms cancel. Differentiating (A.25) with respect to x using $dz/dx = \pi_{xx}$ and $d\pi/dx = z$:

$$rD^{SB}_x = (\hat{\phi}^{-1}z + g)\,\pi_{xx} - \lambda_r z + G'(z)\,\pi_{xx}.$$

The envelope theorem gives $G'(z) = -a^S(z)$. More explicitly, $G'(z) = \frac{da^S}{dz}[\hat{\phi}a^S - z + \Psi'(a^S)] - a^S$, and the FOC (A.20) ensures that $\hat{\phi}a^S - z + \Psi'(a^S) = 0$, so $G'(z) = -a^S(z)$. Therefore:

$$rD^{SB}_x = (\hat{\phi}^{-1}z + g - a^S(z))\,\pi_{xx} - \lambda_r z. \tag{A.26}$$

We show that along the locus $D^{SB}(x) = 0$, the derivative D^{SB}_x is strictly decreasing in z and hence in x, since $z = \pi_x$ is increasing by Lemma A.3. This ensures at most one crossing from $D^{SB} > 0$ to $D^{SB} < 0$. We define:

$$K_1(z) := -\lambda_r z, \qquad K_2(z) := \hat{\phi}^{-1}z + g - a^S(z), \qquad K_3(z) := \pi_{xx}(z),$$

so that $rD^{SB}_x = K_1 + K_2K_3$. To conclude we need $\frac{d}{dz}(K_1 + K_2K_3) < 0$ for all $z > 0$, i.e.,

$$\lambda_r > K_2'(z)\,K_3(z) + K_2(z)\,K_3'(z). \tag{A.27}$$

We bound each factor. We start with K_2 from implicit differentiation of (A.20):

$$\frac{da^S}{dz} = \frac{1}{\hat{\phi} + \frac{\lambda_0\eta r\mu^2}{(1+\eta r\mu a^S)^2}} \in \left[\frac{1}{\hat{\phi} + \lambda_0\eta r\mu^2},\ \hat{\phi}^{-1}\right].$$

So $K_2' = \hat{\phi}^{-1} - da^S/dz \in [0,\ \frac{\lambda_0\eta r\mu^2}{\hat{\phi}(\hat{\phi}+\lambda_0\eta r\mu^2)}]$, and K_2' is non-increasing , since da^S/dz in-

creases with a^S. Thus:

$$\sup_z K_2'(z) = K_2'(0) = \frac{\lambda_0 \eta r \mu^2}{\hat{\phi}(\hat{\phi} + \lambda_0 \eta r \mu^2)} = \frac{\lambda_0 \eta r g^2}{\kappa(\kappa \lambda_r^2 + \lambda_0 \eta r \phi g^2)}. \tag{A.28}$$

Since $K_2' \geq 0$, the function K_2 is non-decreasing. As $z \to \infty$, the FOC gives $a^S \to \hat{\phi}^{-1}(z - \lambda_0 \mu)$, so:

$$\sup_z K_2(z) = \lim_{z \to \infty} K_2(z) = g + \frac{\lambda_0 \mu}{\hat{\phi}} = g\left(1 + \frac{\lambda_0}{\kappa \lambda_r}\right). \tag{A.29}$$

For K_3 from the risky HJB:

$$\tfrac{1}{2}\sigma^2 K_3 = \Lambda \pi - \tfrac{1}{2}\hat{\phi}^{-1} z^2 - gz - (\lambda_0 + \lambda_r)\underline{v}.$$

Using $\pi \leq b$ and $\frac{1}{2}\hat{\phi}^{-1} z^2 + gz \geq 0$:

$$\sup_z K_3(z) \leq \frac{2}{\sigma^2}\Big[\Lambda b - (\lambda_0 + \lambda_r)\underline{v}\Big] =: \frac{2M}{\sigma^2}. \tag{A.30}$$

For K_3', using $D^{SB} = 0$ to eliminate π from the risky HJB:

$$\tfrac{1}{2}\sigma^2 K_3 = r\underline{v} + \frac{r + \lambda_0}{\lambda_r}\left(\tfrac{1}{2}\hat{\phi}^{-1} z^2 + gz\right) + \frac{\Lambda}{\lambda_r} G(z). \tag{A.31}$$

Differentiating and using $G'(z) = -a^S(z)$:

$$K_3'(z) = \frac{2}{\sigma^2 \lambda_r}\Big[(r + \lambda_0)(\hat{\phi}^{-1} z + g) - \Lambda\, a^S(z)\Big]. \tag{A.32}$$

At $z = 0$: $K_3'(0) = \frac{2(r+\lambda_0)g}{\sigma^2 \lambda_r} > 0$. As $z \to \infty$: $K_3'(z) \to -\infty$ (since $(r + \lambda_0)\hat{\phi}^{-1} z - \Lambda \hat{\phi}^{-1} z = -\lambda_r \hat{\phi}^{-1} z$). Thus $\sup_z K_3'(z)$ is finite. The second derivative is:

$$K_3''(z) = \frac{2}{\sigma^2 \lambda_r}\left[(r + \lambda_0)\hat{\phi}^{-1} - \frac{\Lambda}{\hat{\phi} + \frac{\lambda_0 \eta r \mu^2}{(1 + \eta r \mu a^S)^2}}\right].$$

At $z = 0$, where $a^S = 0$: $K_3''(0) = \frac{2}{\sigma^2 \lambda_r}\left[\frac{r+\lambda_0}{\hat{\phi}} - \frac{\Lambda}{\hat{\phi} + \lambda_0 \eta r \mu^2}\right]$. This is non-positive provided:

$$(r + \lambda_0)\lambda_0 \eta r \mu^2 \leq \lambda_r \hat{\phi}, \quad \text{i.e.,} \quad (r + \lambda_0)\lambda_0 \eta r \phi g^2 \leq \lambda_r^3 \kappa. \tag{A.33}$$

Since $\kappa \geq \phi \eta r \sigma^2$, a sufficient condition is $\sigma^2 \geq \frac{g^2 \lambda_0 (\lambda_0 + r)}{\lambda_r^3}$.

Moreover, $K_3''(z) \leq K_3''(0)$ for $z > 0$; the negative term grows in magnitude because

da^S/dz increases. Hence, under (A.33), K_3' is non-increasing for all $z \geq 0$, and:

$$\sup_z [K_3'(z)]^+ = K_3'(0) = \frac{2(r+\lambda_0)g}{\sigma^2 \lambda_r}. \tag{A.34}$$

Since $K_2' \geq 0$, $K_3 \geq 0$ (at the relevant z), and $K_2 > 0$, we bound:

$$K_2' K_3 + K_2 K_3' \leq \sup K_2' \cdot \sup K_3 + \sup K_2 \cdot \sup [K_3']^+.$$

Substituting (A.28), (A.29), (A.30), and (A.34):

$$\lambda_r > \frac{\lambda_0 \eta r g^2}{\kappa(\kappa\lambda_r^2 + \lambda_0 \eta r \phi g^2)} \cdot \frac{2M}{\sigma^2} + g\left(1 + \frac{\lambda_0}{\kappa \lambda_r}\right) \cdot \frac{2(r+\lambda_0)g}{\sigma^2 \lambda_r}. \tag{A.35}$$

Rearranging:

$$\lambda_r^2 \sigma^2 > \frac{2\lambda_0 \eta r g^2 \lambda_r M}{\kappa(\kappa\lambda_r^2 + \lambda_0 \eta r \phi g^2)} + 2g^2(r+\lambda_0)\left(1 + \frac{\lambda_0}{\kappa\lambda_r}\right). \tag{A.36}$$

The LHS grows as λ_r^2 (or as σ^2), while the first RHS term is $O(1)$ in both λ_r and σ^2, and the second RHS term is $O(1)$. Therefore, the condition is satisfied for λ_r or σ^2 sufficiently large.

It remains to verify that the risky action is optimal near $\underline{x}$, so that there exists at least one crossing. At $\underline{x}$, the boundary conditions give $\pi(\underline{x}) = \underline{v}$ and $\pi_x(\underline{x}) = 0$ (smooth pasting), so $z = 0$, $a^S(0) = 0$, $G(0) = 0$, and $rD^{SB}(\underline{x}) = 0$. The derivative at the $\underline{x}$ is:

$$rD_x^{SB}(\underline{x}) = g\,\pi_{xx}(\underline{x}) > 0,$$

by Lemma A.3. So D^{SB} is increasing at $\underline{x}$, meaning the risky action is preferred immediately above $\underline{x}$. Combined with the uniqueness of the crossing, this establishes the cutoff structure: the risky action is optimal on $[\underline{x}, x_c^{SB}]$ and the safe action on $[x_c^{SB}, \bar{x}]$.

For the properties of the effort, in the risky region, $a^R(x) = \hat{\phi}^{-1}\pi_x(x)$, which is continuous and increasing since π is convex.

In the safe region, $a^S(x) = a^S(\pi_x(x))$, where $a^S(z)$ is a continuous, increasing function of z (from (A.22)), and $z = \pi_x(x)$ is increasing. So $a^S(x)$ is continuous and increasing on $(x_c^{SB}, \bar{x}]$.

At the cutoff x_c^{SB}, effort jumps down: from (A.23), $a^S(z) = \hat{\phi}^{-1}[z - \lambda_0\mu(1 - \frac{1}{1+\eta r \mu a^S})] < \hat{\phi}^{-1}z = a^R(z)$ for $\lambda_0 > 0$. This establishes the discontinuous, non-monotonic effort profile in part (ii).

Proof of part (iii) Assume towards a contradiction that the principal implements the risky action near $\bar{x}$. Then $\mathcal{L}^R\pi(x) \geq \mathcal{L}^S\pi(x)$. We construct a lower bound for $\mathcal{L}^S\pi$ by

evaluating the safe operator at the suboptimal risky controls $a = a^R = \hat{\phi}^{-1}z$ with the binding IC $\psi = -\mu\, a^R = -\frac{g}{\kappa\lambda_r}z$. The safe operator evaluated at these controls equals the risky operator minus:

$$gz + \lambda_r\Big(\underline{v} - \pi\Big) - \lambda_0\bigg(-\frac{g\kappa^{-1}}{\lambda_r}z + \frac{1}{\eta r}\ln\bigg(1 + \frac{g\eta r\kappa^{-1}}{\lambda_r}z\bigg)\bigg).$$

Since $\mathcal{L}^R\pi \geq \mathcal{L}^S\pi \geq$ this lower bound, the quantity above must be non-negative:

$$g\pi_x + \lambda_r(\underline{v} - \pi) \geq \lambda_0\bigg(-\frac{g\kappa^{-1}}{\lambda_r}\pi_x + \frac{1}{\eta r}\ln\bigg(1 + \frac{g\eta r\kappa^{-1}}{\lambda_r}\pi_x\bigg)\bigg). \tag{A.37}$$

Rearranging (moving the first λ_0 term to the LHS):

$$g\pi_x\bigg(1 + \frac{\lambda_0\kappa^{-1}}{\lambda_r}\bigg) + \lambda_r\underline{v} \geq \lambda_r\pi + \frac{\lambda_0}{\eta r}\ln\bigg(1 + \frac{g\eta r\kappa^{-1}}{\lambda_r}\pi_x\bigg) \geq \lambda_r\pi,$$

where the last inequality uses $\ln(1+t) > 0$ for $t > 0$. Hence:

$$g\pi_x\bigg(1 + \frac{\lambda_0\kappa^{-1}}{\lambda_r}\bigg) + \lambda_r\underline{v} \geq \lambda_r\pi.$$

From the risky HJB with $\pi_{xx} \geq 0$, $\pi_x \leq \frac{-g+\sqrt{g^2+2\hat{\phi}^{-1}[\Lambda\pi-(\lambda_0+\lambda_r)\underline{v}]}}{\hat{\phi}^{-1}}$. Substituting back:

$$g\cdot\hat{\phi}\bigg(-g + \sqrt{g^2 + 2\hat{\phi}^{-1}[\Lambda\pi - (\lambda_0+\lambda_r)\underline{v}]}\bigg)\bigg(1 + \frac{\lambda_0}{\kappa\lambda_r}\bigg) + \lambda_r\underline{v} \geq \lambda_r\pi.$$

As $\pi(x) \to b$ near $\bar{x}$, the LHS grows as $O(\sqrt{b})$ while the RHS grows as $O(b)$. For b sufficiently large, the inequality fails, yielding a contradiction. Therefore, the principal must implement the safe action near $\bar{x}$, and the cutoff x_c^{SB} is interior.

References

Bergemann, D. and U. Hege (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance 22*(6-8), 703–735.

Bergemann, D. and U. Hege (2005). The financing of innovation: Learning and stopping. *RAND Journal of Economics*, 719–752.

Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large risks, limited liability, and dynamic moral hazard. *Econometrica 78*(1), 73–118.

Bloedel, A., R. V. Krishna, and B. Strulovici (2023). Persistent private information revisited. Technical report, UCLA, FSU, and Northwestern University.

Branco, F., M. Sun, and J. M. Villas-Boas (2012). Optimal search for product information. *Management Science 58*(11), 2037–2056.

Bromberg-Silverstein, S., S. Moreno-Bromberg, and G. Roger (2021). Scale effects in dynamic contracting. *Mathematics and Financial Economics 15*, 431–472.

Bubeck, S. et al. (2023). Sparks of artificial general intelligence: Early experiments with gpt-4. Technical report. https://arxiv.org/abs/2303.12712.

Camboni, M. and T. Durandard (2024). Comparative statics for optimal stopping problems in nonstationary environments. Technical report, University of Wisconsin.

Cetemen, D., F. Z. Feng, and C. Urgun (2023). Renegotiation and dynamic inconsistency: contracting with non-exponential discounting. *Journal of Economic Theory 208*, 105606.

Cusumano, M. A. (1995). *The Software Business: Lessons from Microsoft and Other Leading Companies*. Free Press.

DeMarzo, P. and Y. Sannikov (2016). Learning, termination and payout policy in dynamic incentive contracts. *The Review of Economic Studies*.

DeMarzo, P. M., D. Livdan, and A. Tchistyi (2014). Risking other people's money: Gambling, limited liability, and optimal incentives. Technical Report 3149, Stanford GSB.

Di Tella, S. and Y. Sannikov (2021, 5). Optimal asset management contracts with hidden savings. *Econometrica 89*(3), 1099–1139.

Feng, F. Z., C. R. Taylor, M. M. Westerfield, and F. Zhang (2024). Setbacks, shutdowns, and overruns. *Econometrica 92*(3), 815–847.

Feng, F. Z., C. R. Taylor, M. M. Westerfield, and F. Zhang (2025). Setbacks big and small. *Available at SSRN 5880422*.

Fudenberg, D., B. Holmström, and P. Milgrom (1990). Short-term contracts and long-term agency relationships. *Journal of economic theory 51*(1), 1–31.

Georgiadis, G. (2015). Projects and team dynamics. *The Review of Economic Studies 82*(1), 187–218.

Georgiadis, G., S. A. Lippman, and C. S. Tang (2014). Project design with limited commitment and teams. *The RAND Journal of Economics 45*(3), 598–623.

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

Guo, Y. (2016). Dynamic delegation of experimentation. *American Economic Review 106*(8), 1969–2008.

Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *The Review of Economic Studies 83*(3), 1040–1091.

Hay, M., D. W. Thomas, et al. (2021). Clinical development success rates 2011–2020. *Biotechnology Innovation Organization*. https://go.bio.org/rs/490-EHZ-999/images/ClinicalDevelopmentSuccessRates2011_2020.pdf.

He, Z. (2011). A model of dynamic compensation and capital structure. *Journal of Financial Economics 100*(2), 351–366.

He, Z., B. Wei, J. Yu, and F. Gao (2017). Optimal long-term contracting with learning. *The Review of Financial Studies 30*(6), 2006–2065.

Hellwig, M. F. and K. M. Schmidt (2002). Discrete-time approximations of the Holmström-Milgrom Brownian-motion model of intertemporal incentive provision. *Econometrica 70*(6), 2225–2264.

Holmström, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society*, 303–328.

Hopenhayn, H. A. and J. P. Nicolini (1997). Optimal unemployment insurance. *Journal of Political Economy 105*(2), 412–438.

Hörner, J. and L. Samuelson (2013). Incentives for experimenting agents. *The RAND Journal of Economics 44*(4), 632–663.

Jeanblanc, M., M. Yor, and M. Chesney (2009). *Mathematical methods for financial markets*. Springer Science & Business Media.

Kocherlakota, N. R. (2004). Figuring out the impact of hidden savings on optimal unemployment insurance. *Review of Economic Dynamics 7*(3), 541–554.

Kola, I. and J. Landis (2004). Can the pharmaceutical industry reduce attrition rates? *Nature Reviews Drug Discovery 3*, 711–716.

Li, R. and N. Williams (2025). Optimal contracts with hidden risk. *Review of Economic Dynamics*, 101306.

Madsen, E. (2022). Designing deadlines. *American Economic Review 112*(3), 963–997.

Manso, G. (2011). Motivating innovation. *Journal of Finance 66*, 1823–1860.

Marinovic, I. and F. Varas (2019). Ceo horizon, optimal pay duration, and the escalation of short-termism. *The Journal of Finance 74*(4), 2011–2053.

Mason, R. and J. Välimäki (2015). Getting it done: dynamic incentives to complete a project. *Journal of the European Economic Association 13*(1), 62–97.

McClellan, A. (2022). Experimentation and approval mechanisms. *Econometrica 90*(5), 2215–2247.

Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica 70*(2), 583–601.

Myerson, R. B. (2015). Moral hazard in high office and the dynamics of aristocracy. *Econometrica 83*(6), 2083–2126.

Pallmann, P. et al. (2018). Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC Medicine 16*, 29.

Rochet, J.-C. and G. Roger (2016). Risky utilities. *Economic Theory 62*, 361–382.

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *The Review of Economic Studies 75*(3), 957–984.

Shavell, S. and L. Weiss (1979). The optimal payment of unemployment insurance benefits. *Journal of Political Economy 87*(6), 1347–1362.

Strulovici, B. and M. Szydlowski (2015). On the smoothness of value functions and the existence of optimal strategies in diffusion models. *Journal of Economic Theory 159*, 1016–1055.

Teschl, G. (2012). *Ordinary Differential Equations and Dynamical Systems*, Volume 140. American Mathematical Society Providence.

Toxvaerd, F. (2006). Time of the essence. *Journal of Economic Theory 129*(1), 252–272.

Williams, N. (2015). A solvable continuous time dynamic principal–agent model. *Journal of Economic Theory 159*, 989–1015.

Wong, T.-Y. (2019). Dynamic agency and endogenous risk-taking. *Management Science 65*(9), 4032–4048.

Supplementary Appendix

S.1 Endogenous Project Scope

With endogenous scope, the principal's stopping payoff $\mathbf{1}_{\{X_\tau \geq \bar{x}\}} b + \mathbf{1}_{\{X_\tau < \bar{x}\}} \underline{v}$ is replaced with $g(X_\tau) := \max\{\underline{v}, b(X_\tau)\}$, where the function b is strictly increasing, concave, and continuously differentiable, with $b(0) < \underline{v} < \lim_{x\to\infty} b(x)$ and $b_x(x) \to 0$ as $x \to \infty$. Let $x_0 > 0$ be the progress level at which $b(x_0) = \underline{v}$. The stopping payoff enters the derivation in Section 4 only through the boundary conditions, so the project value is

$$\pi(x) = \max_{A,\tau} \mathbb{E}_x^A \left[e^{-r\tau} g(X_\tau) - \frac{\hat{\phi}}{2} \int_0^\tau e^{-rs} A_s^2 \, ds \right],$$

and the first-best value v is given by the same problem with ϕ in place of $\hat{\phi}$. These are one-dimensional control and stopping problems, so the value functions are continuously differentiable, and value matching and smooth pasting hold at every boundary point of the continuation region at which g is differentiable (see, e.g., Strulovici and Szydlowski, 2015).

Let $[\underline{x}, \bar{x}]$ be the component of the continuation region of the second-best problem that contains $X_0 = 0$. The value function π is nondecreasing, since a higher initial progress level yields a weakly higher payoff under any given strategy, and $\pi(0) > \underline{v}$ as the project is worth starting. Hence no point of $[0, x_0]$ is a stopping point, and $\underline{x} < 0 < x_0 < \bar{x}$, where $\bar{x} < \infty$ because $b_x(x) \to 0$. Near $\underline{x}$, $g = \underline{v}$, so value matching and smooth pasting read $\pi(\underline{x}) = \underline{v}$ and $\pi_x(\underline{x}) = 0$; near $\bar{x}$, $g = b$, so they read $\pi(\bar{x}) = b(\bar{x})$ and $\pi_x(\bar{x}) = b_x(\bar{x})$. On $[\underline{x}, \bar{x}]$, π satisfies the HJB equation (15) and hence the ODE (19). The same argument applies to the first-best problem and its continuation region $[\underline{x}^{FB}, \bar{x}^{FB}]$.

S.2 Flow Payoffs

Suppose that the principal receives a flow payoff $h(X_t)$ while the project is active, where h is increasing, strictly convex, and twice continuously differentiable, with $h(\bar{x}) \leq rb$ and $h(x) \leq r\underline{v}$ for x low enough, so that the principal prefers completing the project at $\bar{x}$, and abandoning it at low progress levels, to continuing indefinitely. The derivation in Section 4 goes through with the flow payoff added to the principal's objective, and the project value solves

$$r\pi(x) = h(x) + \frac{1}{2\hat{\phi}} \pi_x(x)^2 + \frac{1}{2}\sigma^2 \pi_{xx}(x) \tag{S.1}$$

on $[\underline{x}, \bar{x}]$ subject to value matching (16) and smooth pasting (17), with effort $a = \hat{\phi}^{-1}\pi_x$. Differentiating (S.1) gives $\hat{\phi}^{-1}\pi_x\pi_{xx} + \frac{1}{2}\sigma^2\pi_{xxx} = r\pi_x - h_x$, so by Itô's formula

$$d\left(e^{-rt}a(X_t)\right) = -e^{-rt}\hat{\phi}^{-1}h_x(X_t)\,dt + e^{-rt}a_x(X_t)\sigma\,dZ_t.$$

Discounted effort is thus a supermartingale up to τ: the marginal value of progress, and hence effort, grows in expectation at a rate less than r, the shortfall being the marginal flow payoff. The next result shows that effort can decrease in progress, but only near the target.

Proposition S.1 *There exists $x^c \in (\underline{x}, \bar{x}]$ such that effort is increasing on $[\underline{x}, x^c]$ and nonincreasing on $[x^c, \bar{x}]$.*

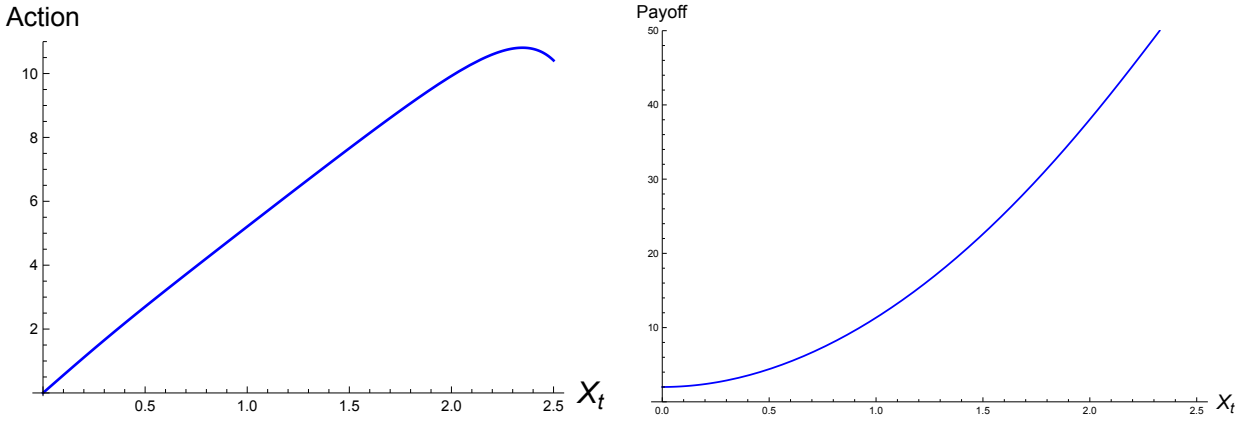


Figure S.1: Left: Value function. Right: Effort profiles. $(b, \bar{x}, \sigma, \underline{v}, r, \eta, \phi) = (30, 5, 1, 1, 2, 1/2, 1)$.

Proof of Proposition S.1

Let $m := a_x$ and $F := ra - \hat{\phi}^{-1}h_x$. Dividing the differentiated ODE by $\hat{\phi}$ gives

$$\frac{1}{2}\sigma^2 m_x + am = F, \qquad F_x = rm - \hat{\phi}^{-1}h_{xx} < rm.$$

At $\underline{x}$, $a = 0$ by smooth pasting, and $m(\underline{x}) = \frac{2}{\sigma^2\hat{\phi}}(r\underline{v} - h(\underline{x})) > 0$ by (S.1), since $h(\underline{x}) < r\underline{v}$ (otherwise the project would not be abandoned at $\underline{x}$). Let x^c be the first zero of m in $(\underline{x}, \bar{x}]$, or $x^c := \bar{x}$ if there is none. Then $m > 0$ on $(\underline{x}, x^c)$, and if $x^c < \bar{x}$, then $m_x(x^c) \leq 0$ and hence $F(x^c) \leq 0$. Let $x' := \sup\{x \leq \bar{x} : F \leq 0 \text{ on } [x^c, x]\}$. On $[x^c, x']$, $F \leq 0$, so

$$m(y) = \frac{2}{\sigma^2}\int_{x^c}^{y} e^{-\frac{2}{\sigma^2}\int_s^y a(u)\,du} F(s)\,ds \leq 0$$

by the linear ODE for m with $m(x^c) = 0$, so $F_x < 0$, and F is strictly decreasing. Hence $F < 0$ on $(x^c, x']$. If $x' < \bar{x}$, then $F < 0$ also on a right neighborhood of x': by continuity if $x' > x^c$, and because $F_x(x^c) = -\hat{\phi}^{-1} h_{xx}(x^c) < 0$ if $x' = x^c$. This contradicts the definition of x', so $x' = \bar{x}$ and $m \leq 0$ on $[x^c, \bar{x}]$.

S.3 Project Selection under Moral Hazard

In this section, we examine how moral hazard distorts the principal's choice among projects that differ in their volatility. This question arises naturally when the principal must select a technology or approach before contracting begins. In the first-best, higher volatility is unambiguously beneficial: it increases the option value embedded in the stopping problem without affecting the cost of effort. Under moral hazard, however, volatility also enters the effective cost of effort through $\kappa = 1 + \phi\eta r\sigma^2$, creating a tension between the option-value and incentive-cost channels.

We parameterize the project by its volatility σ and write $v^{FB}(x;\sigma)$ and $\pi(x;\sigma)$ for the first-best and second-best project values, respectively. Both solve their respective HJB equations:

$$r\, v^{FB}(x;\sigma) = \frac{1}{2\phi}\Big[v_x^{FB}(x;\sigma)\Big]^2 + \frac{1}{2}\sigma^2 v_{xx}^{FB}(x;\sigma), \tag{S.2}$$

$$r\, \pi(x;\sigma) = \frac{1}{2\hat{\phi}(\sigma)}\Big[\pi_x(x;\sigma)\Big]^2 + \frac{1}{2}\sigma^2 \pi_{xx}(x;\sigma), \tag{S.3}$$

where $\hat{\phi}(\sigma) = \kappa(\sigma)\phi$ with $\kappa(\sigma) = 1 + \phi\eta r\sigma^2$.

Proposition S.2 *Consider two projects that differ only in their volatility parameters, $\sigma_1 > \sigma_2$, with all other exogenous parameters $(b, \bar{x}, \phi, \eta, r, \underline{v})$ the same. The termination thresholds $\underline{x}^{FB}(\sigma)$ and $\underline{x}(\sigma)$ are determined endogenously.*

(i) The first-best project value is increasing in volatility: $v^{FB}(x;\sigma_1) > v^{FB}(x;\sigma_2)$ for all $x \in (\underline{x}^{FB}(\sigma_2), \bar{x})$. Moreover, $\underline{x}^{FB}(\sigma_1) < \underline{x}^{FB}(\sigma_2)$, so for $x \in (\underline{x}^{FB}(\sigma_1), \underline{x}^{FB}(\sigma_2)]$, the higher-volatility project is still active (with value above $\underline{v}$) while the lower-volatility project has been terminated (with value $\underline{v}$).

(ii) The second-best project value need not be monotone in volatility. There exist parameter configurations such that

$$v^{FB}(x_0;\sigma_1) > v^{FB}(x_0;\sigma_2) \quad \textit{but} \quad \pi(x_0;\sigma_1) < \pi(x_0;\sigma_2).$$

That is, the principal may prefer a lower-variance project even when the social planner would select the higher-variance one.

The proposition highlights a novel distortion introduced by moral hazard: the principal's project selection is biased toward lower-risk projects. This bias arises because the incentive-adjusted cost parameter $\hat{\phi}(\sigma)$ is increasing and convex in σ^2, so the marginal incentive cost of additional volatility grows without bound, whereas the marginal option value eventually diminishes.

Proof of Proposition S.2

Proof Part (i): Let $\sigma_1 > \sigma_2$, and write $v_i(x) := v^{FB}(x;\sigma_i)$. The comparison is the same one used in the proof of Proposition 2(iii). In the continuation region, strict convexity gives $v_{1,xx} > 0$, and therefore

$$rv_1 - \frac{1}{2\phi}v_{1,x}^2 - \frac{\sigma_2^2}{2}v_{1,xx} = \frac{\sigma_1^2 - \sigma_2^2}{2}v_{1,xx} > 0.$$

Thus the higher-volatility value is a strict supersolution of the lower-volatility continuation equation, while both problems have the same stopping payoff and completion payoff. The comparison principle for the associated variational inequalities yields

$$v^{FB}(x;\sigma_1) \geq v^{FB}(x;\sigma_2)$$

for every x, with strict inequality wherever the lower-volatility project continues. It also yields

$$\underline{x}^{FB}(\sigma_1) < \underline{x}^{FB}(\sigma_2).$$

Hence, for $x \in (\underline{x}^{FB}(\sigma_1), \underline{x}^{FB}(\sigma_2)]$, the higher-volatility project remains active and has value strictly above $\underline{v}$, while the lower-volatility project has already been terminated and has value $\underline{v}$.

Part (ii): In the second-best HJB equation (S.3), an increase in σ has two effects: (a) it raises the diffusion coefficient $\frac{1}{2}\sigma^2$, which increases the option value of continuation, and (b) it raises the effective cost $\hat{\phi}(\sigma) = (1 + \phi\eta r\sigma^2)\phi$, which lowers the optimal effort and reduces the value of the project. Effect (b) is absent in the first-best.

To see that the net effect can be negative, consider the numerical example from the text: set all parameters to 1 except $b = 10$, with starting progress $X_0 = -0.473$. For $\sigma = 0.6$, we have $\hat{\phi} = 1.36$ and $\pi(X_0) \approx 3.8878$. For $\sigma = 0.8$, we have $\hat{\phi} = 1.64$ and $\pi(X_0) \approx 3.5677$. Thus $\pi(X_0;0.6) > \pi(X_0;0.8)$. Since the first-best value is increasing in σ (part (i)), we have $v^{FB}(X_0;0.8) > v^{FB}(X_0;0.6)$, establishing the divergence.

Lemma S.1 *Consider the progress process* (1) *with* $X_0 = 0$ *and any effort process* $A = (A_t)_{t\geq 0}$ *progressively measurable with* $A_t \geq 0$. *Let* $\tau^A := \inf\{t \geq 0 : X_t \notin (\underline{x}, \bar{x})\}$ *denote the two-sided*

exit time. Define the reflected Brownian motion at the lower boundary $\underline{x}$,

$$R_t := \underline{x} + \sigma Z_t + \ell_t, \qquad \ell_t := \sup_{0 \le s \le t} (-\sigma Z_s)^+,$$

and its upper hitting time $\tau^{\text{ref}} := \inf\{t \ge 0 : R_t = \bar{x}\}$. Then

$$\tau^A \le \tau^{\text{ref}} \quad \textit{a.s.}$$

In particular, $\tau^A < \infty$ almost surely for every such A, and $\mathbb{E}[\tau^A] \le \mathbb{E}[\tau^{\text{ref}}] < \infty$.

Proof Shift the state space so that the lower boundary is at 0. Let $H := \bar{x} - \underline{x}$ and $x_0 := -\underline{x} \in (0, H)$, and define

$$Y_t^A := X_t - \underline{x} = x_0 + \int_0^t A_s \, ds + \sigma Z_t, \qquad Y_t^0 := x_0 + \sigma Z_t.$$

Then $\tau^A = \inf\{t \ge 0 : Y_t^A \notin (0, H)\}$.

Apply the one-sided Skorokhod reflection map on $[0, \infty)$: for a continuous path f with $f(0) \ge 0$, set $\Gamma(f)(t) := f(t) + \ell_f(t)$ where $\ell_f(t) := \sup_{0 \le s \le t}(-f(s))^+$. Let $R_t^0 := \Gamma(Y^0)(t)$. By construction, R^0 has the same law as the reflected process R shifted by $\underline{x}$, and $\tau^{\text{ref}} = \inf\{t \ge 0 : R_t^0 = H\}$.

We now establish the pathwise domination $\tau^A \le \tau^{\text{ref}}$. Fix ω and work pathwise. Write $\mathcal{A}_t := \int_0^t A_s \, ds$ for the cumulative drift. For any $t < \tau^A(\omega)$, we have $Y_s^A(\omega) \in (0, H)$ for all $s \le t$. In particular,

$$Y_s^0(\omega) = Y_s^A(\omega) - \mathcal{A}_s(\omega) > -\mathcal{A}_s(\omega),$$

so $(-Y_s^0(\omega))^+ < \mathcal{A}_s(\omega) \le \mathcal{A}_t(\omega)$ for each $s \le t$, where the last inequality uses $A \ge 0$. Taking suprema over $s \le t$,

$$\ell_{Y^0}(t; \omega) = \sup_{0 \le s \le t} (-Y_s^0(\omega))^+ \le \mathcal{A}_t(\omega).$$

Therefore, for all $t < \tau^A(\omega)$,

$$R_t^0(\omega) = Y_t^0(\omega) + \ell_{Y^0}(t; \omega) \le Y_t^0(\omega) + \mathcal{A}_t(\omega) = Y_t^A(\omega) < H.$$

This shows $R_t^0(\omega)$ cannot reach H before $\tau^A(\omega)$, i.e., $\tau^{\text{ref}}(\omega) \ge \tau^A(\omega)$. Since ω was arbitrary, we have $\tau^A \le \tau^{\text{ref}}$ almost surely.

Finally, $\tau^{\text{ref}} < \infty$ a.s. because reflected Brownian motion on $[0, \infty)$ hits any fixed level almost surely. Moreover, $\mathbb{E}[\tau^{\text{ref}}] < \infty$: the process R^0 has the same law as $|x_0 + \sigma Z_t|$, so

τ^{ref} has the same distribution as the first hitting time of $\{\pm H\}$ by $x_0 + \sigma Z_t$, whose mean is finite. Thus $\tau^A \leq \tau^{\text{ref}}$ yields $\tau^A < \infty$ a.s. and $\mathbb{E}[\tau^A] \leq \mathbb{E}[\tau^{\text{ref}}] < \infty$.