The Optimal Schedules of Incentives and Cash Flows*

Felix Zhiyu Feng
Univ. of Georgia

Robin Yifan Luo
Univ. of Washington

Mark M. Westerfield
Univ. of Washington

July, 2026

Abstract

We model delegated management when an agent can invest in durable asset quality or raise instantaneous cash flow. Both actions add value, but only aggregate output is observable, so incentivizing durable investment necessarily also incentivizes instantaneous effort. The flexibility to schedule incentives makes durable investment relatively cheaper to induce, and the optimal contract begins back-loaded. As promised future incentives accumulate, further back-loading becomes increasingly costly, and the contract can eventually become front-loaded. Negative cash-flow surprises raise promised future incentives. The resulting dynamics generate mean reversion in cash flows and other empirical implications for the duration of cash flows and incentives.

JEL Classification: D82, D86, G31, G34, M12, M52
Keywords: asset quality, short-termism, dynamic contracting, persistent private information, incentive schedules, cash-flow duration

*Feng can be reached at ffeng@uga.edu, Luo at yifanluo@uw.edu, and Westerfield at mwesterf@uw.edu.
We thank Harjoat Bhamra, Philip Bond, Simon Gervais, Ahmed Guecioueur, Fahad Khalil, Jacques Lawarrée, Doron Levit, George Nurisso, Stefano Pegoraro, Adriano Rampini, Xuyuanda Qi, Can Urgun, Jiang Xu, Hongda Zhong, seminar and conference participants at Duke University, University of Georgia, University of Washington, FIRS, MFA, Econometric Society NASM, the UNC-Chapel Hill Economic Theory Conference, and the Vienna Festival of Finance Theory for helpful comments and suggestions.

1 Introduction

Delegated management typically involves two types of activities: maintaining day-to-day performance and investing in future prospects. A CEO may boost near-term earnings by cutting costs or ramping up advertising, while also building internal processes, organizational know-how, or a customer base, all of which generate benefits over time. Importantly, both short-term effort and long-term investment generally add value. "Working harder" and "working smarter" each increase output and are valuable in their own right. However, their outcomes mix to make them indistinguishable: the market observes earnings, but cannot determine whether they reflect recent cost-cutting efforts or earlier efforts to build organizational know-how and a customer base. This leads to the fundamental questions we study in this paper: when short-term and long-term actions are both profitable but their outcomes cannot be separately observed, how can they be separately incentivized? Which action will the compensation contract favor?

A natural conjecture is that long-term investment is more difficult and expensive to induce. Instantaneous effort produces an immediate effect, creating a relatively direct link between the agent's action and observed output. Durable investment instead builds a hidden stock whose effects emerge gradually and may be difficult to distinguish from other changes in performance. When the agent is risk averse, providing incentives also requires exposing him to output risk, seemingly reinforcing the disadvantage of long-term investment.

Our paper reaches a different and surprising conclusion: at low and moderate levels, durable investment is cheaper to incentivize. Inducing more instantaneous effort today requires strengthening incentives tied to today's performance. In contrast, the impact of durable investment is spread over time, so it can be induced with incentives that are spread over time. Durable investment can be encouraged through a *flexible schedule* of future incentives; this flexibility advantage allows those incentives to be concentrated at the times and states in which they are most effective or least costly. Incentive schedules are front-loaded when they provide stronger incentives today than they promise in the future, and this tilts effort toward the instantaneous action. Incentive schedules are back-loaded when they provide stronger incentives in the future, and this tilts effort toward durable investment. Our analysis allows us to characterize these optimal, flexible schedules.

We study a dynamic principal-agent problem in which the agent can add value through instantaneous effort or by building durable asset quality. The productive effects of the two actions differ only in timing: instantaneous effort generates cash flows today, while investment in asset quality generates the same present value of cash flows but spreads them over time. The principal observes aggregate cash flows but not the agent's actions or asset quality, so both actions must be induced through the same schedule of pay-for-performance sensitivity (PPS) – the agent's exposure to output surprises. Current PPS governs instantaneous effort, while the utility-weighted average of promised future PPS governs durable investment. The principal chooses both the level and timing of PPS, making the contract relatively front- or back-loaded. Because the agent is risk averse, PPS is costly: exposing his continuation utility to output surprises requires additional compensation.

Our model generates three core results. First, the optimal contract begins back-loaded, with its stock of promised future incentives exceeding current PPS. To see why, start from a constant PPS schedule fixed at the level that maximizes the principal's payoff from instantaneous effort. This constant schedule delivers the same moderate level of current and promised future incentives, and it satisfies the adding up constraint for short-term and long-term incentives. Then, we make a small increase to promised future PPS without changing the present. Because this is a small increase from the myopic maximum, it has no first-order cost. However, it does have a first-order benefit because it raises durable investment today. The change reduces durable investment in the future, when incentives revert as the perturbation closes, but that effect is discounted. Thus, the value-maximizing contract begins back-loaded.

Second, the stock of promised future incentives is stochastic and negatively correlated with output: negative cash-flow surprises raise both promised future incentives and durable investment. However, the cost advantage of future incentives is not global. As promised future incentives accumulate, further back-loading requires concentrating additional PPS in states where incentives are already strong or at their maximum value – the flexibility advantage is gone. Thus, further back-loading becomes increasingly costly and leaves less scope to improve the timing of incentives. After a sufficiently prolonged sequence of negative cash-flow surprises, current PPS exceeds the stock of promised future incentives, and the contract becomes front-loaded. The negative correlation with output also helps discipline the agent. Failure to invest in asset quality produces persistent negative cash-flow surprises, whereas bad luck is transitory. Thus, concentrating future PPS after losses strengthens promised incentives precisely along the histories generated by underinvestment. Because PPS is costly, this state-contingent schedule is the cost-minimizing way to provide long-term incentives.

Third, the model links the dynamics of optimal incentives to empirical measures of duration and autocorrelation. The duration of incentives is negatively related to the stock of promised future incentives. After a loss, both current PPS and promised future incentives rise, but current PPS rises more. The contract becomes more front-loaded: PPS is higher today and expected to decline, giving it a shorter duration. Thus, poor performance optimally leads to front-loaded incentives, rather than front-loaded incentives causing poor performance. This prediction is consistent with Gopalan, Milbourn, Song, and Thakor (2014), who find that the duration of CEO compensation falls after poor performance.

We also show that the main results survive exogenous stochastic project termination. We consider an extension in which the project's cash flows may cease randomly, at which point all contractual promises, including future compensation and incentive provision, are extinguished. At first glance, termination would appear to raise the cost of long-term incentives since they rely on promises about future payoffs. Nevertheless, we show that all of our main results continue to hold. The reason is that the core mechanism of the optimal contract – the flexibility to distribute long-term incentives over time and to concentrate them in states in which they are more effective and less costly – does not depend on the contract's horizon. As long as termination does not occur at a predetermined date, the principal can still promise stronger incentives following below-expectation performance, conditional on the project's survival, and these promises retain their disciplining effect on the agent's actions.

The model's incentive dynamics provide a new perspective on the tension between delivering current earnings and making long-term investments documented by Graham, Harvey, and Rajgopal (2005). Because long-term cash flows are sequences of short-term cash flows, incentives for current and durable actions cannot be cleanly separated. The model predicts that negative cash-flow surprises are followed by higher current PPS and a more front-loaded incentive schedule. Thus, we explain why a value-maximizing contract may place greater weight on short-term actions following poor performance, even when durable investment remains productive. The same mechanism generates mean reversion in cash flows. Losses lead the contract to strengthen incentives and induce greater investment, raising expected subsequent cash flows. The optimal contract can make cash flows more predictable over finite horizons even though their instantaneous volatility is fixed.

Our paper belongs to the dynamic contracting literature with a multidimensional action space and persistent private information.[1] The studies closest to ours include Gryglewicz, Mayer, and Morellec (2020), Hackbarth, Rivera, and Wong (2022), Hoffmann and Pfeil (2021), He, Wei, Yu, and Gao (2017), and Marinovic and Varas (2019). Gryglewicz et al. (2020) and Hackbarth et al. (2022) consider agency models with separately observable short- and long-term tasks, allowing the tasks to be incentivized separately. These models examine how static and dynamic production or cost complementarities generate over- or underinvestment in the two tasks. Hoffmann and Pfeil (2021) study a dynamic multitasking model

[1] This literature can be divided into two broad groups. In the first, the agent has private information about exogenously evolving states or serially correlated shocks; examples include Fernandes and Phelan (2000); Kennan (2001); Zhang (2009); Kapička (2013); Pavan, Segal, and Toikka (2014); Bergemann and Strack (2015); Cisternas (2018); Bloedel, Krishna, and Leukhina (2025); Fu and Krishna (2019); Krasikov and Lamba (2021). In the second, the agent takes a private action with persistent effects; examples include Williams (2011, 2015); Sannikov (2014); Di Tella and Sannikov (2021); Bloedel, Krishna, and Strulovici (2023); Feng and Westerfield (2024). We add a short-term private action and study the resulting dynamics of compensation and project cash-flow duration. Relatedly, Gryglewicz, Mayer, and Morellec (2024) study loan origination and monitoring when the agent chooses a long-term action once at time zero and then takes a sequence of short-term actions. Because the long-term decision is made only at inception, their optimal contract is always front-loaded. In our model, the agent continuously allocates effort between short- and long-term actions, and the contract can switch between back- and front-loading with performance history.

in which the agent can divert current cash flow and undertake non-contractible intangible investment. The investment lowers current net cash flow but raises the probability of observable success that improves future profitability. The contract can use separate instruments for cash-flow diversion and investment: the agent's continuation utility moves continuously with reported cash flows and jumps when investment succeeds or fails. In contrast, in our model the principal observes only the cash flows jointly generated by the two actions and must use the same PPS schedule to incentivize both. We also abstract from default and from primitive production or cost complementarities. Our results instead arise from the intertemporal incentive constraint linking current PPS to the stock of promised future incentives.

Marinovic and Varas (2019) study a dynamic agency model with persistent private information in which the agent exerts productive instantaneous effort but can also manipulate the performance measure. Manipulation is their persistent action: it raises cash flows today, reverses over time, and is weakly value-destroying by assumption. Long-term incentives in their model serve to deter an action the principal never wants to induce. The dynamics of their contract are driven entirely by the CEO's known retirement date: with an infinite horizon, their optimal contract is constant and linear, and manipulation is irrelevant. Our model is fundamentally different because the persistent action creates value. This difference allows us to pose our central question: when short-term and long-term actions are equally productive and their outputs pool in a single observable cash flow – so that they cannot be separately incentivized – which one is cheaper to incentivize? This question does not arise in their setting because the persistent action is never worth inducing. Thus, the cost of long-term incentives in their model is a deterrence cost, rather than a cost to be compared with that of inducing short-term effort. The same distinction also changes the source of contract dynamics. Our alternation between back- and front-loaded incentive schedules arises in precisely the stationary, infinite-horizon environment in which their optimal contract is constant, and it is driven by the agent's performance history rather than by calendar time or an approaching deadline. Section 5.3 provides a more detailed comparison after presenting our main mechanism.

Meanwhile, He et al. (2017) considers a dynamic agency model in which the principal and agent learn about the asset's underlying quality. The agent can manipulate the principal's beliefs by exerting excess effort today, making the principal more optimistic before subsequent performance falls short of those expectations. Such manipulation generates negatively serially correlated cash-flow surprises: positive surprises are followed by disappointment. The principal optimally responds by strengthening incentives after success, thereby raising subsequent cash flows and firm value. This response makes temporary and permanent cash-flow shocks appear positively correlated. These predictions are the opposite of those generated by our model of building and revealing asset quality, highlighting the qualitative distinction between a learning model and a building model.

Our paper also speaks to the broad literature on managerial myopia and short-termism (e.g., Stein, 1988, 1989; Shleifer and Vishny, 1990; Froot, Perold, and Stein, 1992; Holmström and Tirole, 1993; Von Thadden, 1995; Bolton, Scheinkman, and Xiong, 2006; Edmans, Gabaix, Sadzik, and Sannikov, 2012; Zhu, 2018). These studies generally emphasize the value-destroying nature of short-termism and investigate its sources or remedies. Our model does not assume that short-term action is inferior to long-term action. Instead, we study how persistent private information affects the optimal allocation of effort between the two and show that agency frictions can alternately tilt this allocation toward short- or long-term actions relative to benchmark economies without persistent agency frictions. Our analysis also complements empirical work on cash-flow and asset duration, including Dechow, Sloan, and Soliman (2004), Lettau and Wachter (2007), Weber (2018), Gonçalves (2021), and Gormsen and Lazarus (2023). Whereas these studies primarily examine the relationship between duration and equity returns, we study the economic mechanism determining duration and how it responds to cash-flow shocks.

2 Model and Benchmark Economies

We study a principal-agent problem in which a principal (she) owns a real asset that is managed by the agent (he). Time is continuous, starting at $t = 0$, and the principal and agent share a discount rate $r > 0$. The agent is risk-averse and has hidden savings, while the principal is risk neutral with deep pockets. We begin with a baseline model in which the project and agent live forever, and then extend the model in Section 4 to give the project (and any contractual promises) a stochastic duration.

2.1 Assets, Technology, and Preferences

The agent takes two hidden actions that determine the cash flows from the asset he manages. The first, $q_t \geq 0$, builds durable asset quality. We interpret q_t as managerial effort that develops organizational or asset quality – for example, internal processes, organizational know-how, talent, culture, governance, or a customer base – rather than necessarily as observable investment expenditure. The resulting latent stock of quality, $Q_t \geq 0$, accumulates according to

$$dQ_t = (q_t - \rho Q_t)\, dt, \tag{1}$$

where quality depreciates at rate $\rho > 0$. The second action, $a_t \geq 0$, is hidden effort directed toward current cash generation. Cumulative cash flows Y_t evolve as

$$dY_t = (a_t + (r + \rho)Q_t)\, dt + \sigma dZ_t, \tag{2}$$

whose drift combines instantaneous effort with the cash flows generated by accumulated quality. The process Z_t is a Brownian cash-flow shock with volatility $\sigma > 0$.

We normalize the cash flows from Q_t by $(r + \rho)$ so that one unit of quality generates a present discounted value equal to one. Thus, quality is measured in dollars, and Q_t can be viewed as embodied capital. Under the standard transversality condition stated in

Appendix A.1, integration by parts shows that incremental units of a_t and q_t have equal present values:

$$\int_t^\infty e^{-r(s-t)} dY_s = Q_t + \int_t^\infty e^{-r(s-t)} \left[(a_s + q_s)\, ds + \sigma dZ_s \right]. \tag{3}$$

The additive technology rules out production complementarities, while the normalization gives a marginal unit of each action the same present-value contribution to output. The actions differ only in timing: a_t affects current cash flow, whereas q_t builds quality that generates cash flows over time. Thus, neither action has an intrinsic productivity advantage.

The agent consumes at rate c_t out of a savings account S_t and receives a cumulative compensation process M_t from the principal. We assume that the common discount rate r is also the risk-free return on private savings, so that S_t evolves as

$$dS_t = (rS_t - c_t)\, dt + dM_t. \tag{4}$$

The agent has CARA utility over consumption net of the costs of actions a_t and q_t,

$$u(c, a, q) = -\frac{1}{\gamma} \exp\left[-\gamma \left(c - \frac{1}{2}\theta(a^2 + q^2) \right) \right], \tag{5}$$

where $\gamma > 0$ is the coefficient of constant absolute risk aversion and $\theta > 0$ is the common action-cost parameter. The separable quadratic term is a consumption-equivalent effort or opportunity cost. It assigns the same cost to both actions and rules out cost complementarities.

Together, these restrictions are identification choices. Equal present-value productivity, identical separable effort costs, and the absence of production and cost complementarities place the two actions on equal terms except for the timing of their cash-flow effects. Consequently, differences in their optimal provision – and the resulting movement between back- and front-loaded incentive schedules – arise from the endogenous costs of providing incentives through a shared performance signal.

2.2 The Agency Problem

The principal observes aggregate cash flows and contractual transfers but not the agent's actions, consumption, savings, or asset quality. She recommends the actions to be taken and conditions the contract on reported values of the hidden states Q_t and S_t. Thus, the agent has hidden actions $\{c_t, a_t, q_t\}$ and persistent hidden information $\{Q_t, S_t\}$. We assume that Q_0 and S_0 are known; e.g., the asset's quality and the agent's wealth may begin at zero.

We use "hats" to denote the principal's recommended policies and the reports of the agent, and we use plain letters to denote the agent's true decisions and state variables. For example, $\hat{c}$ is the amount of consumption the principal requests, while c is the amount of consumption actually undertaken. Similarly, we write expectations conditioned on the true state of the project and the agent's actions (the agent's information set) as $\mathrm{E}_t[(\cdot)_s]$. Expectations conditioned on cash flows, the principal's recommendations, and the agent's reports (the principal's information set) are written as $\hat{\mathrm{E}}_t[(\cdot)_s]$.

Assumption 1 *The agent can commit not to quit before the project terminates exogenously.*[2] *His reports of the hidden state variables must be consistent with an admissible action path and a realization of the exogenous shocks.*

Finally, we assume that $\sigma^2\gamma r\theta \geq 1$. Economically, this condition ensures that granting the agent the entire output claim to implement maximum instantaneous effort yields a non-positive flow payoff to the principal. Technically, it guarantees the nonnegative quadratic term needed for the global incentive-compatibility bound when $\phi_t \leq 0$; see the proof of Proposition 2 in Appendix C.

[2]The baseline project has an infinite horizon; Section 4 introduces exogenous stochastic termination. The same "no-voluntary-quitting" assumption is commonly adopted in similar studies of persistent private information under CARA utility to improve tractability, including Williams (2015), He et al. (2017), and Marinovic and Varas (2019). Such a commitment may be implemented by requiring the agent to forfeit a bond or deferred compensation following voluntary departure. A technologically infeasible report may trigger the same penalty because it is publicly verifiable.

2.3 The Contract

The agent's objective function is his lifetime expected discounted utility:

$$W_t = \max \mathrm{E}_t \left[\int_t^{\infty} e^{-r(s-t)} u(c_s, a_s, q_s) ds \right]. \tag{6}$$

The principal's objective is the expected discounted value of net cash flows from the project:

$$V_t = \max \hat{\mathrm{E}}_t \left[\int_t^{\infty} e^{-r(s-t)} (dY_s - dM_s) \right]. \tag{7}$$

We make two useful simplifications to the contracting environment. The first is to set the agent's recommended level of savings at his certainty equivalent wealth:

$$\hat{S}_t = -\frac{1}{\gamma r} \ln\left(-\gamma r W_t\right). \tag{8}$$

This normalization selects a convenient representation of the agent's reported savings. Because the principal, the agent, and private markets use the same interest rate, changes in the timing of reported wealth can be offset through contractual borrowing or saving without changing consumption, incentives, or either party's value.[3]

The second simplification restricts the agent's exposure to cash-flow risk to lie between zero and one. This is analogous to the standard static restriction that the principal's and agent's effective ownership shares be nonnegative, thereby ruling out incentives to manipulate cash flows, as in Innes (1990). Formally, we require an adapted process $\beta_t \in [0, 1]$ such that the agent's cash-flow exposure in the compensation process can be written as $dM_t = (\cdots)dt + \beta_t dY_t$. Appendix A.3 provides a simple implementation rationale for this assumption.[4]

[3]The principal and agent are indifferent to the equilibrium level of reported savings during the life of the project. We allow this reported savings balance to be negative.

[4]A finite upper bound on β_t is also critical for analytical tractability. If arbitrarily large values of β_t were admissible, both the control set and the induced state space would become unbounded. The compactness, value bounds, and boundary arguments underlying our HJB analysis would then no longer apply, making it extremely difficult to establish existence, regularity, and other analytical properties of the value function without imposing alternative restrictions. An unbounded degree of pay-for-performance sensitivity is also unrealistic, as contracts in practice do not expose the agent to an arbitrarily large share of residual cash-flow

The principal maximizes her objective function subject to providing the agent with incentives to take the desired actions:

Definition 1 (Incentive Compatibility and Optimality) *A contract is incentive compatible if the contract induces the agent to accurately report the path of* $\{Q_t, S_t\}$ *and to choose the principal's recommended actions. A contract is optimal if it maximizes the principal's expected payoff among incentive-compatible contracts satisfying* $\beta_t \in [0,1]$ *and delivering the promised initial utility* W_0.

2.4 Benchmark Economies

These benchmarks isolate the source of the contract dynamics in the full model. We consider the first best (FB), agent ownership (AF), and an observable-quality second best (SB). In each benchmark, the two actions can be controlled or internalized separately, eliminating the intertemporal incentive constraint that links them in the full model. The resulting actions and incentive exposures are constant.

Proposition 1 (Benchmark Economies) *Let* $\beta_t \in [0,1]$ *denote the agent's exposure to cash-flow innovations; when the principal offers a contract,* β_t *is current PPS. In all three benchmark economies, consumption adjusts so that flow utility satisfies* $u(c_t, a_t, q_t) = rW_t$, *where* $W_t < 0$ *under CARA utility. Continuation utility evolves with cash-flow surprises according to*

$$dW_t = -\gamma r W_t \beta_t \left(dY_t - \mathrm{E}_t[dY_t]\right). \tag{9}$$

The first best has $\beta_t = \beta^{FB} = 0$*; agent ownership has* $\beta_t = \beta^{AF} = 1$*; and the observable-quality second best has* $\beta_t = \beta^{SB} = \frac{1}{1+\sigma^2\gamma r\theta}$.

---

risk. However, the precise upper bound of 1 is not necessary for either the analytical behavior of the model or our main results. The analysis can accommodate a higher upper bound as long as it is finite. We choose 1 because it is the most natural and practically relevant benchmark: it corresponds to giving the agent the entire residual cash-flow claim and admits the simple implementation rationale provided in Appendix A.3.

Instantaneous effort satisfies $a^{FB} = a^{AF} = 1/\theta$ and $a^{SB} = \beta^{SB}/\theta$, while durable effort is $q_t = 1/\theta$ in all three benchmarks.

The benchmarks differ in the agent's exposure to cash-flow risk. The agent is fully insured in FB, bears all cash-flow risk under agent ownership, and receives the myopically optimal PPS in the observable-quality second best. In FB the principal dictates both actions, under agent ownership the agent internalizes both, and in SB the principal controls durable effort and uses PPS only to induce instantaneous effort. Because the resulting exposures are constant, consumption exposure and both actions are constant. The state-dependent dynamics we study below only happen because the two hidden actions must be induced through the same PPS schedule.

3 The Optimal Contract

We now analyze the full model, with the agent controlling the management of the asset. The principal's inability to observe Q_t creates an asymmetry between a_t and q_t, and the principal offers the agent a time-varying and stochastic schedule of incentives.

3.1 The Agent's Problem

The agent's incentives and optimal actions are summarized by two endogenous state variables. The first is continuation utility W_t, which records the utility promised to the agent by the principal. The exposure to current output surprises is governed by current PPS β_t, which incentivizes instantaneous effort.

The second state variable is z_t, which captures the agent's sensitivity to asset quality Q_t. The agent has two ways to consume more in the future. One is to improve asset quality and receive a share of the resulting future output. The other is to accumulate private savings for later consumption. The variable z_t is the marginal rate of substitution between these two future gains: $\frac{\partial W_t}{\partial Q_t} / \frac{\partial W_t}{\partial S_t}$.[5] Improving asset quality is relatively more attractive – z_t is higher – when the agent's share of future output is higher. Thus, the principal encourages investment in asset quality by increasing the sensitivity of the agent's utility to future output rather than current output. We use ϕ_t to denote the exposure of z_t to output shocks.

The agent maximizes his lifetime utility in (6), subject to (1), (2), (4), and the contract offered by the principal. We use Hamiltonian methods to find necessary and sufficient conditions for the agent to report truthfully and to follow the principal's recommended actions:

Proposition 2 (Incentive Compatibility) *A contract is incentive compatible only if there are integrable, adapted stochastic processes z_t, β_t, and ϕ_t such that*

$$dW_t = -\sigma\gamma r W_t \beta_t d\hat{Z}_t \tag{10}$$

$$dz_t = (r+\rho)(z_t - \beta_t + \sigma^2\gamma r \beta_t \phi_t)dt + (r+\rho)\sigma\phi_t d\hat{Z}_t \tag{11}$$

$$rW_t = u(\hat{c}_t, \hat{a}_t, \hat{q}_t) \tag{12}$$

$$\hat{a}_t = \frac{1}{\theta}\beta_t \tag{13}$$

$$\hat{q}_t = \frac{1}{\theta}z_t, \tag{14}$$

with $\beta_t \in [0,1]$ and $z_t \in [0,1]$.

If $\phi_t \leq 0$, then those conditions are also sufficient.

The continuation value W_t is the discounted value of the agent's future expected utility in (6). Its law of motion follows from (10):

$$dW_t = -\gamma r W_t \beta_t \sigma d\hat{Z}_t = -\gamma r W_t \beta_t \left(dY_t - \hat{\mathrm{E}}_t[dY_t] \right), \tag{15}$$

where β_t is the agent's current pay-for-performance sensitivity (PPS) and $-\gamma r W_t > 0$ is a normalization. When cash flows dY_t exceed expectations, the agent receives additional continuation utility proportional to $-\gamma r W_t \beta_t$ and the surprise. The path of β_t describes how

[5]We use $\partial W_t / \partial S_t$ informally for the shadow value of S_t, which we derive formally in Appendix C. Similarly, $\partial W_t / \partial Q_t$ informally represents the shadow value of Q_t.

the sensitivity of the agent's continuation utility to cash-flow surprises evolves over time.

An important feature of the agent's continuation utility is that it is a martingale ($\mathrm{E}_t[W_s] = W_t$ or $\mathrm{E}_t[dW_t] = 0$), meaning that all changes are surprises. This result stems from the agent's ability to privately save at rate r, which is also the agent's discount rate. If the agent's marginal utility were to drift up, the agent could gain on average by saving and consuming later; if it were to drift down, he could gain by borrowing. The agent would undo any drift, so equilibrium marginal utility must be a martingale. Because the agent is CARA, utility is proportional to marginal utility and must also be a martingale. The agent consequently smooths his flow utility across time, with $u_t = rW_t$ in (12). Because the agent consumes out of savings and marginal utility does not change on average, we have $\partial W_t/\partial S_t \approx u_c(c_t, a_t, q_t) = -\gamma r W_t$.

To generate instantaneous cash flows, the agent repeatedly chooses the static action a_t, incurring a cost of effort equal to $\theta a_t^2/2$ units of consumption. The marginal utility cost is $-u_a = \theta a_t(-\gamma r W_t)$. In return, current PPS gives the agent a utility exposure $(-\gamma r W_t)\beta_t$ to cash flows in (15). Matching marginal cost with marginal benefit and canceling $(-\gamma r W_t)$ yields $a_t = \beta_t/\theta$ in (13). Thus, β_t, which captures the sensitivity of the agent's utility to immediate cash-flow surprises, also measures his incentive to increase cash flows through instantaneous effort.

The agent's second technology uses durable effort q_t to increase asset quality, generating positive and persistent cash flows. Its marginal utility cost is analogous to that of instantaneous effort: $-u_q = \theta q_t(-\gamma r W_t) \approx \theta q_t \partial W_t/\partial S_t$. The approximation follows because the effort cost is measured in consumption units, so its marginal value is equivalent to a change in private savings.

The benefit of an additional unit of asset quality is spread over time. At date $s \geq t$, that unit raises the cash-flow rate by $(r + \rho)e^{-\rho(s-t)}$. Discounting utility at rate r and using (15), the marginal benefit to the agent's continuation utility is

$$\partial W_t / \partial Q_t \approx \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} \left(-(r+\rho)\gamma r W_s \beta_s \right) ds \right]. \tag{16}$$

The first-order condition for durable effort equates its marginal benefit and marginal cost: $\partial W_t/\partial Q_t = \theta q_t \, \partial W_t/\partial S_t$. Dividing (16) by $\partial W_t/\partial S_t$, we obtain

$$\frac{\partial W_t/\partial Q_t}{\partial W_t/\partial S_t} \approx \theta q_t = z_t = \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} (r+\rho) \frac{W_s}{W_t} \beta_s ds \right] \tag{17}$$

which is the integral definition of z_t, with the evolution given in (11). Alternatively, remembering that $u_c(c, a, q) = -\gamma r W$, we have

$$\frac{\partial W_t/\partial Q_t}{\partial W_t/\partial S_t} \approx \theta q_t = z_t = \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} (r+\rho) \frac{u_c(c_s, a_s, q_s)}{u_c(c_t, a_t, q_t)} \beta_s ds \right] \tag{18}$$

which is a familiar Euler equation. Thus, z_t is also the value of future cash-flow sensitivity $\beta_{s>t}$, evaluated at the ratio of marginal utilities. Together, these calculations show that the marginal rate of substitution for the agent between long-term investment and savings is also the *utility-weighted average of future PPS*, which we denote by z_t. Thus, the principal controls the agent's incentives with a *joint stock of promises*: W_t for future utility and z_t for future PPS.

The principal cannot fully separate incentives for the instantaneous and durable actions, but she can tilt their relative strength through the incentive schedule. Because both actions pool in the same cash flow – a_t and Q_t are not separately observable – the principal has only one performance signal and one incentive instrument, the path of PPS β_t. A constant PPS schedule is simple to implement: the integral expression for z_t in (17), together with the martingale property of W_t, implies that $z_t = \beta_t$, so the two actions receive equal incentives. More generally, the contract is back-loaded when $z_t > \beta_t$: current PPS is below the utility-weighted average of promised future PPS, tilting effort toward durable investment. The contract is front-loaded when $\beta_t > z_t$, tilting effort toward the instantaneous action. These definitions compare current PPS with the utility-weighted average of promised future PPS, and they require that future PPS is flexible but not monotonic.

The principal can also make promised future incentives state contingent. The stock z_t is not the simple discounted expected value of β: it is weighted by marginal utility (or continuation utility). Future PPS contributes more to today's long-term incentives when it is provided in states in which the agent has high marginal utility (or low utility).[6] The principal controls the exposure of the stock of promised future incentives to cash-flow shocks through ϕ_t; the diffusion coefficient of z_t is $(r + \rho)\sigma\phi_t$. When ϕ_t is negative, innovations in W_t and z_t are negatively correlated: negative cash-flow surprises reduce promised utility and raise the stock of promised future incentives, while positive surprises do the opposite. Because z_t is the utility-weighted average of future PPS, this state contingency concentrates promised future PPS in low-utility states, where it contributes more to durable effort.

Both current PPS and the stock of promised future incentives are bounded. The agent's current PPS satisfies $\beta_t \in [0, 1]$.[7] The same bound applies to $z_t \in [0, 1]$ by (17). This is a direct implication of the incentive adding-up constraint: a bound on current PPS imposes an equivalent bound on the stock of promised future incentives.

Combining promise keeping with the incentive-compatibility conditions (12), (13), and (14) yields the agent's consumption and compensation:

$$\hat{c}_t = -\frac{1}{\gamma}\ln(-\gamma r W_t) + \frac{1}{2\theta}\beta_t^2 + \frac{1}{2\theta}z_t^2. \tag{19}$$

In fact, consumption in the baseline model takes the same form as in the three benchmark economies (Proposition 1 and Appendix B). Using the agent's consumption (19), the law of motion for savings (4), and the desired level of savings (8), we can back out the agent's required compensation. He receives an initial payment $S_0 = -\frac{1}{\gamma r}\ln(-\gamma r W_0)$ and continuous

[6] Exponential utility is negative, so W_s/W_t is higher when W_s is lower (more negative).

[7] This restriction is introduced in Section 2.3 and given an auxiliary implementation rationale in Appendix A.3.

compensation

$$dM_t = \left[\frac{1}{2}\sigma^2\gamma r\beta_t^2 + \frac{1}{2\theta}\left(\beta_t^2 + z_t^2\right)\right] dt + \beta_t \sigma dZ_t. \tag{20}$$

3.2 The Principal's Problem

The principal maximizes the expected discounted value of cash flows (7) subject to the results of Proposition 2. Given the stock of promised future incentives z_t, the principal chooses current PPS β_t and the exposure ϕ_t of that stock to cash-flow shocks. We use dynamic programming to characterize the optimal policies:

Proposition 3 (The Principal's Value and Optimal Policies) *For $t > 0$, the principal's continuation value is*

$$V_t = V(Q_t, z_t) = Q_t + v(z_t), \tag{21}$$

where v is defined by

$$v(z_t) = \max_{(\beta,\phi)\,\mathrm{IC}} \mathrm{E}_t\left[\int_t^\infty e^{-r(s-t)}\left(\frac{1}{\theta}\beta_s + \frac{1}{\theta}z_s - \frac{1}{2\theta}\beta_s^2 - \frac{1}{2\theta}z_s^2 - \frac{\sigma^2\gamma r}{2}\beta_s^2\right) ds\right]. \tag{22}$$

The function v is a viscosity solution of the principal's HJB equation (D-4), subject to the boundary conditions (D-7) and (D-8). It is continuous on $[z^*, 1]$, strictly concave and twice continuously differentiable on $(z^*, 1)$, and strictly decreasing to the right of its maximizer z^*. The optimal initial stock of promised future incentives is z^*, and the principal's initial value is

$$V_0 = Q_0 + v(z^*) + \frac{1}{\gamma r}\ln(-\gamma r W_0). \tag{23}$$

The optimal policies $\beta(z)$ and $\phi(z)$ satisfy the following properties:

1. At z^*,

$$\beta(z^*) = \beta^* \equiv \frac{1}{1+\gamma\sigma^2 r\theta}, \qquad \phi(z^*) = 0, \quad \textit{and} \quad z^* > \beta^*.$$

2. $\beta(z) > 0$ and $\phi(z) < 0$ for all $z \in (z^*, 1)$.

3. At the right boundary, $\beta(1) = 1$ and $\phi(1) = 0$. Moreover, there exists $z^{**} < 1$ such that $\beta(z) = 1$ for all $z \in (z^{**}, 1]$.

4. Under additional parameter restrictions, $\beta(z) < 1$ for all $z < z^{**}$, and $\beta(z)$ crosses z from below exactly once in the interior domain. [8]

The principal's flow payoff in (22) consists of three components. The terms $\beta_t/\theta + z_t/\theta$ are the outputs generated by instantaneous and durable effort, respectively. Recall that $\hat{a}_t = \beta_t/\theta$ and $\hat{q}_t = z_t/\theta$ from Proposition 2, so the terms $-\beta_t^2/(2\theta) - z_t^2/(2\theta)$ are the corresponding effort costs. Finally, $-\sigma^2\gamma r\beta_t^2/2$ compensates the risk-averse agent for exposure to output risk. Critically, this risk-compensation term is a function of β not z. In other words, promised future incentives have a local marginal-cost advantage over current PPS:

Corollary 4 (Relative Marginal Costs of Incentives) *Denote the principal's flow payoff in (22) as $f(\beta, z)$. For any common incentive level $x \in (0, 1)$, we have*

$$\frac{\partial}{\partial z} f(\beta, z)\bigg|_{\beta=z=x} - \frac{\partial}{\partial \beta} f(\beta, z)\bigg|_{\beta=z=x} = \sigma^2\gamma r x > 0. \tag{24}$$

The corollary follows from differentiating (22).

This result means that even though the agent's two actions have identical productivity and effort costs, the marginal flow value of z is higher than the marginal flow value of β. This difference arises entirely from the $-\frac{\sigma^2\gamma r}{2}\beta_s^2$ term in (22) – the additional risk-compensation cost attached to current PPS. Intuitively, inducing greater short-term effort requires increasing current PPS, which immediately exposes the agent to additional risk. In contrast,

[8] In the online appendix, (OA.3-49) is sufficient for $\beta(z) < 1$ for all $z < z^{**}$, and adding (OA.3-70) is sufficient for single crossing.

inducing greater durable effort requires increasing future PPS, which has flexible timing – the flexibility advantage is that the principal places this promised future PPS in the states where it is least costly.[9]

There is an adding up constraint in incentives because z_t is the discounted, utility-weighted average of future β. But, if the principal could control β_t and z_t separately, she would choose

$$\beta_t = \arg\max_{\beta} \frac{\beta}{\theta} - \frac{\beta^2}{2\theta} - \frac{\sigma^2 \gamma r \beta^2}{2} = \frac{1}{1+\gamma r \sigma^2 \theta} = \beta^* \tag{25}$$

$$z_t = \arg\max_{z} \frac{z}{\theta} - \frac{z^2}{2\theta} = 1. \tag{26}$$

These constant policies coincide with the policies in the observable-quality second best of Subsection 2.4, but are not feasible. So, how should the principal choose β_t and z_t jointly? Constant policies are feasible – $\beta_t = z_t = \beta^*$ delivers the required promises at the myopically optimal level of β_t – but the principal can do better by increasing z_t by increasing the future $\beta_{s>t}$.

We now work through a small upward perturbation of $\beta_{s>t}$, moving future PPS away from its myopic optimum and increasing z_t today. At the benchmark of $\beta_t = z_t = \beta^*$, the direct first-order effect – foregone instantaneous-effort payoff and additional compensation for volatility – is exactly zero:

$$\frac{\partial}{\partial \beta_s}\left[\frac{\beta_s}{\theta} - \frac{\beta_s^2}{2\theta} - \frac{\sigma^2\gamma r}{2}\beta_s^2\right]\Bigg|_{\beta_s=\beta^*} = 0.$$

This means that any direct loss is second order. In contrast, the gain from increasing z_t is first order: each unit of z_t adds $\frac{1}{\theta}(1-z_t) > 0$ to the principal's flow payoff in (22). Since $z_t = \beta^* < 1$, this perturbation is possible, and because the perturbation adds value, we will

[9]This analysis highlights why we assume that the short- and long-term actions a and q generate identical present values, have identical cost functions, and exhibit no production or cost complementarities (as opposed to Gryglewicz et al., 2020 or Hackbarth et al., 2022). In this way, we isolate differences in incentives so that we can attribute differences in their usage to the endogenous costs of inducing the two actions. The same reasoning motivates the normalization used when the model is extended to incorporate termination in Section 4.

have $z^* > \beta^*$.

The optimal value of z^* is defined so that the marginal value of additional promised future incentives is zero, $v'(z^*) = 0$. The promise-keeping term drops out of the first-order condition for current PPS, which reduces to the myopic condition: $\beta(z^*) = \beta^*$ (Proposition 3). The contract begins with current PPS at its myopic optimum and a strictly larger stock of promised future incentives, $z^* > \beta(z^*)$, so it is back-loaded.

The optimal contract is not back-loaded forever: Proposition 3 shows that $\beta(z)$ crosses z from below, meaning that when z has moved upwards from its initial value, the contract becomes front-loaded with current PPS exceeding promised future incentives. The earlier perturbation argument fails because the cost advantage of durable incentives stems from their flexibility: when z_t is high, many near-future paths have $\beta_{s>t} = 1$ so that flexibility has been lost. If we attempt a small upward perturbation of $\beta_{s>t}$, it may not be possible except in the far future after PPS has declined. Then we see that our perturbation argument ignored discount rates: compensation costs are discounted at rate r, whereas the contribution of future PPS to z_t is discounted at rate $r+\rho$ because asset quality depreciates. This discounting difference means that perturbing $\beta_{s>t}$ in the near future and in the far future have different costs. Raising sufficiently distant PPS to affect z_t today is expensive when z_t is already large, so promised future incentives have a low marginal cost near the value-maximizing state but a high marginal cost when z_t is close to one.

Together, these arguments explain the contract's changing incentive schedule. It begins at z^*, where $\beta(z^*) < z^*$, and is back-loaded. As promised future incentives accumulate, the marginal cost of further back-loading rises. Near $z = 1$, the stock of promised future incentives becomes relatively more expensive than current PPS, $\beta(z) > z$, and the contract becomes front-loaded.

Our second main result is that $\phi_t \leq 0$ is optimal, with strict inequality on the interior $(z^*, 1)$. In other words, the stock of promised future incentives grows after negative cash-flow surprises. The intuition follows from (17): this stock is the *utility-weighted average* of future PPS. This means that future PPS is more effective in inducing durable effort today when it is provided in low-utility, high-marginal-utility states. Because the agent's promised utility moves with cash-flow surprises, these are the states that follow losses. Cash flows and the stock of promised future incentives are consequently negatively correlated. This result complements the intuition for our first main result. At low and moderate levels, long-term incentives are cheaper because of the flexibility advantage: future PPS can be scheduled optimally and concentrated in the states in which it is most effective. Greater short-term effort, in contrast, requires a higher level of current PPS, which must be delivered today and cannot be shifted toward more effective future states.

Figure 1 illustrates how the value function and optimal policies vary with the stock of promised future incentives z_t. Our numerical examples demonstrate $\beta'(z) > 0$, which is stricter than what we have proven.[10] Together with $\phi(z) < 0$ for $z \in (z^*, 1)$, this implies that both current PPS and the stock of promised future incentives rise after negative cash-flow surprises. The numerical solution also shows that current PPS rises more rapidly. The intuition of this result is analogous to the response of z to negative cash-flow shocks. If the agent underinvests in Q_t, the resulting shortfall produces persistent negative cash-flow surprises. In contrast, bad luck produces an isolated negative surprise that is as likely to be followed by a positive surprise as by another negative one. Because volatility is costly, it is cheaper to provide both current PPS and promised future PPS state-contingently. Raising both after negative surprises exposes an agent who generates repeated losses to higher PPS and a larger stock of promised future incentives more persistently than an agent who follows the principal's recommended actions but experiences low cash flows because of bad luck.

Figure 2 illustrates two consequences of promise keeping. In the front-loaded region, $\beta_t > z_t$ and $\phi_t < 0$, so both terms in the drift of z_t are negative and z_t is expected to decline. The drift is positive at lower, back-loaded states and negative at higher states, creating a tendency for the contract's loading to moderate over time. The right panel shows that the utility-weighted average of future PPS exceeds its unweighted counterpart, consistent with

[10]We find $\beta'(z) > 0$ in every parameterization we considered, but we have been unable to prove it. Nevertheless, parameter-only conditions (H1)–(H3) establish the uniqueness of the maximum-PPS region $[z^{**}, 1]$ and of the single-crossing point $\widehat{z}$. These conditions are sufficient but not necessary.

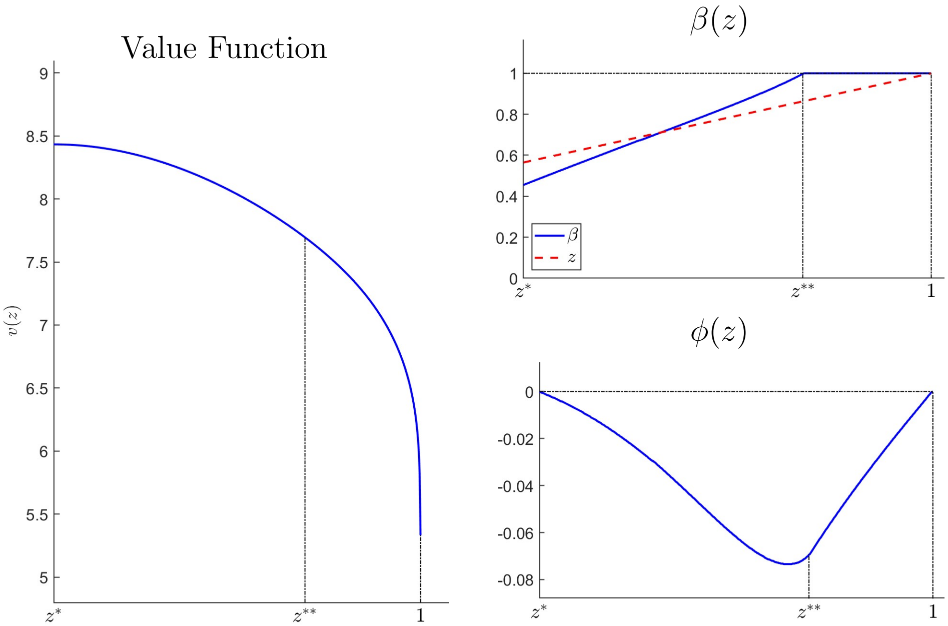

Figure 1: The Principal's Value and Policy Functions

The left panel plots the principal's value function $v(z)$ under the optimal contract. The top right panel plots current PPS $\beta(z)$ in blue and the stock of promised future incentives z in dashed red. The bottom right panel plots the exposure $\phi(z)$ of that stock to output shocks. The value z^* maximizes $v(z)$, and z^{**} is the left boundary of the region in which $\beta = 1$. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

future PPS being concentrated in low-utility states.[11]

Because $\phi_t < 0$, promised future incentives rise after negative cash-flow surprises and fall after positive surprises. Current PPS moves in the same direction and responds more strongly. The resulting changes in instantaneous and durable effort generate mean reversion in cash flows: positive histories are followed by lower-PPS, more back-loaded schedules, whereas adverse histories are followed by higher-PPS, more front-loaded schedules.

4 Risky Capital and Project Termination

We extend the model to allow the project to terminate. In general, termination makes quality accumulation risky because the project may end before an investment realizes its

[11] Although W_t is a martingale, it converges to zero almost surely, so its constant mean is sustained by rare extreme losses. In the baseline numerical solution, the modest wedge between the weighted and unweighted averages indicates that the incentive mechanism is not driven primarily by these tail histories. Section OA.4 of the Online Appendix provides the derivation and further discussion.

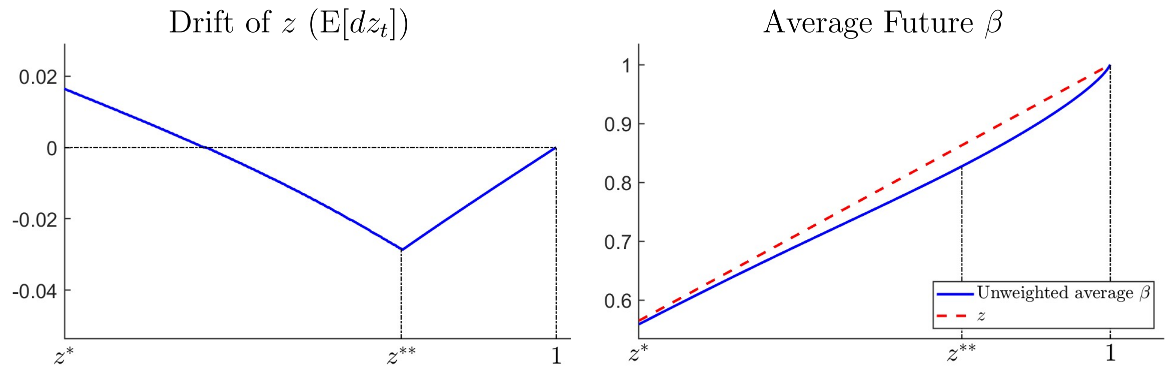

Figure 2: Weighted and Unweighted Averages of Future β and the Drift of z

The left panel plots the drift of z, $(r+\rho)(z_t - \beta_t + \sigma^2\gamma r \beta_t \phi_t)$. The right panel plots the utility-weighted and unweighted averages of future β. The former is $z = \hat{E}\left[\int_t^\infty e^{-(r+\rho)(s-t)}(r+\rho)(W_s/W_t)\beta_s ds\right]$, and the latter is $\hat{E}\left[\int_t^\infty e^{-(r+\rho)(s-t)}(r+\rho)\beta_s ds\right]$. The value z^* maximizes $v(z)$, and z^{**} is the left boundary of the region in which $\beta = 1$. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

full payoff. For a conservative robustness test, we assume that termination extinguishes both the project's future cash flows and its remaining contractual promises. This modeling choice creates a strong force *against* our main result that the flexibility to schedule long-term incentives can make them cheaper to provide. Durable effort generates value through future cash flows and is induced through a schedule of future incentives; termination may eliminate the former before they are realized and the latter before they are delivered. In contrast, current effort and current incentives are realized immediately.

Formally, we assume that the project is subject to an exogenous termination shock given by a Poisson process dN_t with intensity λ that hits at random time τ. This shock simply stops all cash flows, and the agent consumes out of savings from then on. We assume that the total quality of the asset, Q_t, accumulates as in the baseline model (1). However, cash flows are re-specified as

$$dY_t = (a_t + (r + \rho + \lambda)Q_t)\,dt + \sigma dZ_t, \tag{27}$$

for $t \in [0, \tau]$, with cash flows equal to zero thereafter. The normalization $(r + \rho + \lambda)Q_t$ preserves the first-best equivalence of the short- and long-term actions: one unit of quality continues to generate a present discounted value of cash flows equal to one,

$$\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}dY_s\right] = Q_t + \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}(a_s + q_s)ds\right]. \tag{28}$$

Thus, neither action has an intrinsic technological advantage, but long-term investment is risky because the project may end before all of its cash flows are realized.

To accommodate termination risk while maintaining tractability, we assume that public markets are complete enough for the agent to invest his hidden savings in securities contingent on public, exogenous shocks such as project termination.[12] The agent's savings evolve as

$$dS_t = (rS_t - c_t)\,dt + \nu_t(dN_t - \lambda dt) + dM_t, \tag{29}$$

where we added ν_t to reflect the agent's hidden risky portfolio.

Finally, the agent still maximizes lifetime expected utility, as in (6), and he consumes out of his savings after the project ends. The principal's objective function is still the expected discounted value of net cash flows from the project, which now ends:

$$V_t = \max \hat{\mathrm{E}}_t\left[\int_t^\tau e^{-r(s-t)}(dY_s - dM_s)\right]. \tag{30}$$

Termination makes investments in quality risky and permits hidden wealth transfers across survival and termination states. Nevertheless, the main results are unchanged:

Proposition 5 (Stochastic Project Termination) *Conditional on project survival, the incentive-compatibility conditions in Proposition 2 continue to hold, with the law of motion*

[12] The purpose of this assumption is mainly technical. Without it, the contracting problem would need additional state variables describing termination-contingent wealth or utility, and additional incentive-compatibility conditions governing how hidden saving affects consumption after termination. Consequently, the one-dimensional reduction of the contracting problem to z would generally fail. Economically, this assumption isolates the effect of a stochastic project horizon from this additional incomplete-markets problem.

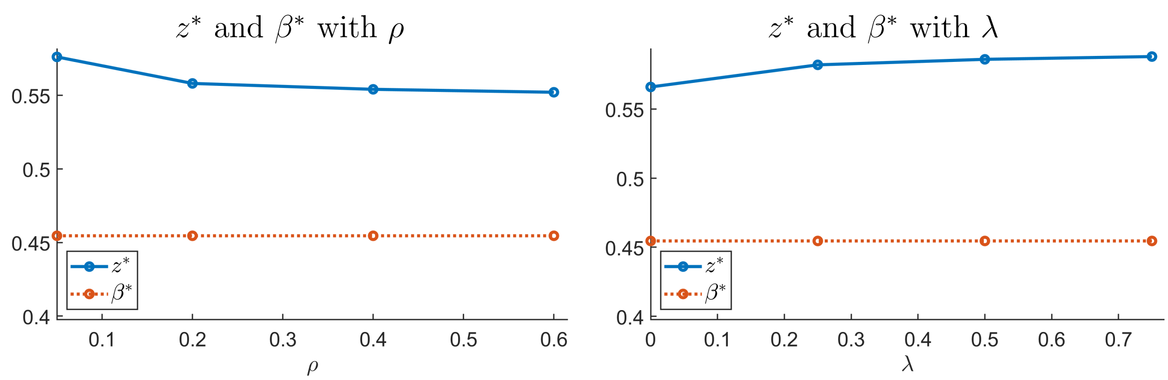

Figure 3: Comparative Statics of z^*

The figure plots the maximizing value z^* and current PPS β^* as ρ varies (left panel) and as λ varies (right panel). The baseline parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

for z_t replaced by

$$dz_t = (r + \rho + \lambda)(z_t - \beta_t + \sigma^2 \gamma r \beta_t \phi_t)dt + (r + \rho + \lambda)\sigma\phi_t d\hat{Z}_t - z_t dN_t. \tag{31}$$

The value-function and optimal-policy results in Proposition 3 also continue to hold after incorporating λ into the effective discounting and state dynamics. Specifically,

$$V_t = Q_t + \mathrm{E}_t \left[\int_t^\infty e^{-(r+\lambda)(s-t)} \left(\frac{1}{\theta}\beta_s + \frac{1}{\theta} z_s - \frac{1}{2\theta}\beta_s^2 - \frac{1}{2\theta} z_s^2 - \frac{\sigma^2 \gamma r}{2}\beta_s^2 \right) ds \right]. \tag{32}$$

Moreover, $z^* > \beta(z^*) = \frac{1}{1+\gamma\sigma^2 r\theta}$.

When the project terminates, all promises of future PPS are extinguished by assumption. The stock z_t consequently jumps to zero in (31). Before termination, it is the *stochastically terminating* utility-weighted average of future PPS:

$$z_t = \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho+\lambda)(s-t)} (r + \rho + \lambda) \frac{W_s}{W_t} \beta_s ds \right], \tag{33}$$

where β_s is the PPS that applies at date s conditional on the project surviving until then. The additional discounting at rate λ accounts for the probability that this promise is extinguished before it is delivered.

Proposition 5 establishes analytically that the core features of the optimal contract sur-

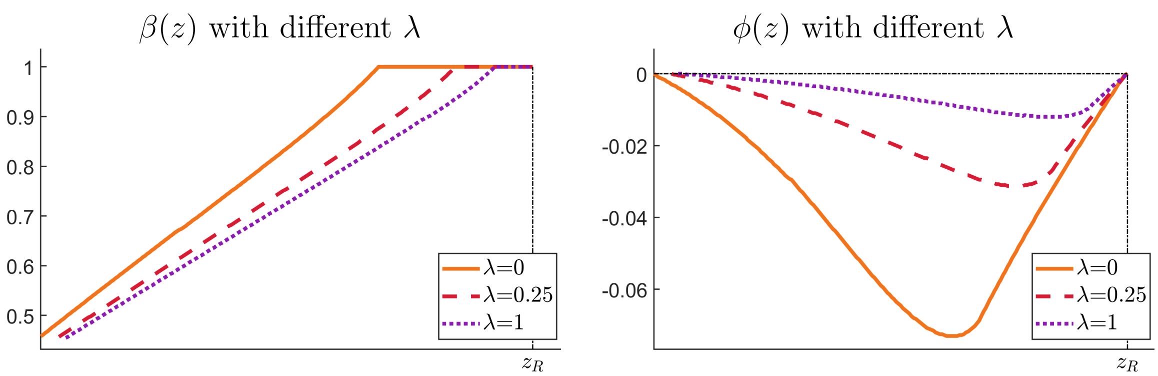

Figure 4: Comparative Statics of $\beta(z)$ and $\phi(z)$

The panels plot how $\beta(z)$ (left) and $\phi(z)$ (right) vary with λ. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, and $\theta = 1.5$.

vive termination: the contract begins back-loaded ($z^* > \beta^*$), cash flows and the stock of promised future incentives are negatively correlated ($\phi < 0$), and the contract has both front- and back-loaded regions. Figures 3 and 4 supplement these results with comparative statics for z^*, $\beta(z)$, and $\phi(z)$.

The comparative statics yield a perhaps surprising result: increasing λ raises z^*, making the initial contract more reliant on promised future incentives, while current PPS, $\beta(z^*) = 1/(1+\gamma\sigma^2 r\theta)$, remains constant. The explanation lies in relative discounting. In the baseline model, cash flows and utility are discounted at rate r, whereas depreciation causes investment in quality and its supporting future PPS to be discounted at rate $r + \rho$; see (17). This difference disadvantages long-term incentives relative to the current risk-compensation cost $-\sigma^2\gamma r\beta_t^2/2$. Consistent with this mechanism, the left panel of Figure 3 shows that increasing ρ lowers z^* and reduces initial back-loading.

Termination adds λ to both rates: cash flows are discounted at $r+\lambda$ in (32), while future PPS is discounted at $r+\rho+\lambda$ in (33). Although their absolute gap remains ρ, their relative difference is

$$\frac{r+\rho+\lambda}{r+\lambda} = 1 + \frac{\rho}{r+\lambda},$$

which declines with λ. As a result, promised future incentives are disadvantaged relatively

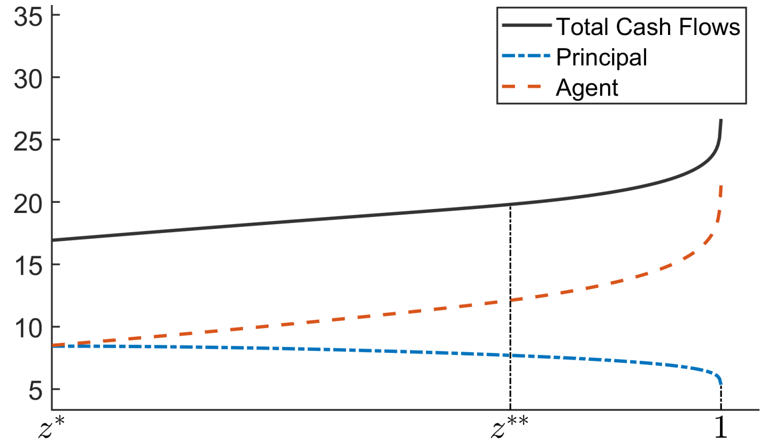

Figure 5: Value of Total Cash Flows, Principal's Cash Flows, and Agent's Compensation

The figure plots the expected value of total cash flows $\mathrm{E}_t[\int_t^{\tau} e^{-r(s-t)}dY_s]$, principal's cash flows $\mathrm{E}_t[\int_t^{\tau} e^{-r(s-t)}(dY_s - dM_s)]$, and agent's compensation $\mathrm{E}_t[\int_t^{\tau} e^{-r(s-t)}dM_s]$. The initial Q is set to be 0. z^* represents the point at which $v(z)$ is maximized, and z^{**} represents the left boundary of the region where $\beta = 1$. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

less as termination risk rises. The right panel of Figure 3 shows that the principal responds by increasing z^* and initial back-loading. This result contrasts sharply with the effect of a deterministic deadline in Marinovic and Varas (2019), as discussed in Section 5.3.

5 Discussion of Mechanisms and Results

In this section, we further describe our model's implications for cash flow dynamics and duration, and how those implications differ from those obtained by other contracting models.

5.1 Impulse Response and Autocorrelation

The model transforms transitory cash-flow innovations into persistent changes in incentives, effort, and expected cash flows. We examine these effects using three present values: the principal's net cash flows $\mathrm{E}_t[\int_t^{\tau} e^{-r(s-t)}(dY_s - dM_s)]$, the agent's compensation $\mathrm{E}_t[\int_t^{\tau} e^{-r(s-t)}dM_s]$, and total project cash flows $\mathrm{E}_t[\int_t^{\tau} e^{-r(s-t)}dY_s]$. Figure 5 plots these values. The first two add up to the third, and compensation becomes an increasing fraction of the total as the stock of promised future incentives z rises.

The key ingredient is that the principal responds to transitory shocks by changing promised future incentives, which gives those shocks durable effects. As z and current PPS rise, the agent devotes more effort to both durable and instantaneous actions, cash flows increase, and the agent's required compensation increases. At the same time, the principal's value declines because the marginal compensation cost exceeds the marginal cash-flow benefit. Because positive current cash-flow shocks reduce promised future incentives ($\phi \leq 0$), positive shocks today increase the principal's value but reduce the project's expected cash flows. Similarly, the agent's continuation utility increases ($dW/dZ > 0$), while his expected future compensation decreases.[13]

Critically, long-term incentives are cheaper than current PPS when the stock of promised future incentives is moderate, so optimal contracts begin back-loaded. The maximizing value z^* is the lower boundary of the continuation region, and the calibrated policy has $\beta(z) > \beta^*$ for $z > z^*$. Thus, a positive shock benefits both the principal and the agent in welfare terms (both $v(z)$ and W_t increase in response to positive cash-flow shocks) while the agent's compensation and effort both decline.

To illustrate the effect of temporary cash-flow shocks on future expected values, we plot impulse response functions in Figure 6. These valuation effects are durable, with impacts on expected cash flows well into the future. Indeed, the impact on cash-flow values can be greater in the medium term than in the short term, as illustrated by panels 6d through 6f in the second row of Figure 6. Following a positive cash-flow shock ($dZ_t > 0$), the stock of promised future incentives z and the expected value of future cash flows both decrease. Because a positive cash-flow innovation lowers subsequent expected cash flows, the model generates negative cash-flow autocorrelation over the corresponding horizons even though the underlying Brownian innovations are serially independent.

When z starts in the cap-binding region, the gap between the shocked and unshocked economies initially increases over time. A positive shock initially leaves optimal current PPS unchanged at $\beta = 1$ because the resulting z remains above z^{**}. The adjustment instead

[13]Recall here that the agent's compensation includes reimbursement for the private cost of effort.

occurs through a decline in the stock of promised future incentives. Current PPS begins to fall only after z crosses below z^{**}; see the $\beta(z)$ panel of Figure 1. Keeping current PPS high during this initial phase accelerates the decline in z through the adding-up relation between current and promised future incentives. This delayed response of current PPS causes the effect on expected cash flows to grow initially before dissipating.

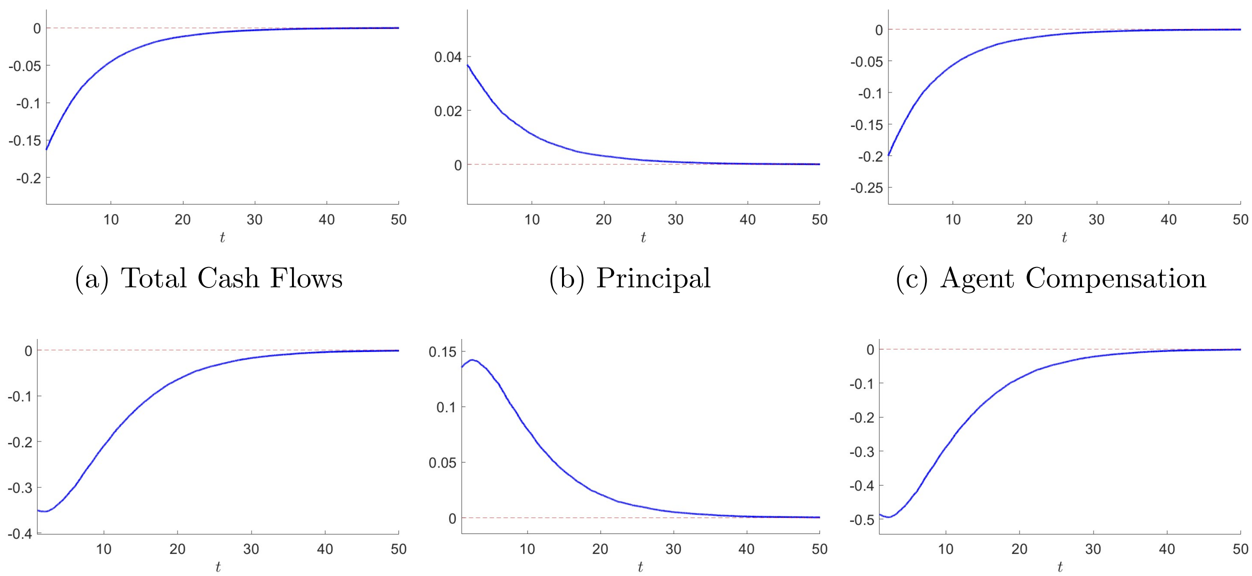

(a) Total Cash Flows (b) Principal (c) Agent Compensation

(d) Total Cash Flows (e) Principal (f) Agent Compensation

Figure 6: Impulse Response Functions

The figure plots differences between the shocked and unshocked expected present values of total cash flows, the principal's net cash flows, and the agent's compensation following a positive cash-flow innovation ($dZ_t > 0$). The x-axis is time, and the vertical axes report the corresponding differences in present value. We consider initial values $z_0 \in \{0.7, 0.95\}$, representing moderate (top row) and high (bottom row) existing incentives. The innovation changes the initial state by $\Delta z_0 = (r + \rho + \lambda)\sigma\phi(z_0)dZ_t < 0$. For each z_0, the figure averages 1,000 simulated responses. The initial Q is zero. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

5.2 Macaulay Duration

To measure the relative importance of long-term and short-term actions and incentives, we use Macaulay duration. For a nonnegative stochastic process X_t that is cut off at τ, with strictly positive present value, Macaulay duration is the cash-flow-weighted average time to payment:

$$D_t(X) = \frac{\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}(s-t)X_s ds\right]}{\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}X_s ds\right]}. \tag{34}$$

Intuitively, duration is the value-weighted average time at which X is received. The duration of a constant process, $X_t = x > 0$, is $1/(r+\lambda)$. We also apply the same ratio to the principal's net cash flows, which can be negative at some dates. In that application it is a generalized signed duration rather than literally an average time to payment.

For a nonnegative process, larger duration indicates greater back-loading. Thus, a long duration of PPS places relatively more weight on PPS delivered further in the future, whereas a short duration of project cash flows means that more of their value is realized sooner. Duration allows us to assess how far into the future PPS must extend to induce long-duration cash flows and whether the durations of PPS and cash flows are matched.

To simplify the calculation and interpretation, denote the expected discounted value of X by $F_t(X)$:

$$F_t(X) = \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}X_s ds\right]. \tag{35}$$

Fubini's theorem (Appendix A.2) then gives

$$D_t(X) = \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}\frac{F_s(X)}{F_t(X)}ds\right] = F_t\left(\frac{F_s(X)}{F_t(X)}\right). \tag{36}$$

Equation (36) expresses duration as the discounted integral of future remaining value relative to current value. Duration is high when a large fraction of the current value remains outstanding far into the future. For example, if X_t is total project cash flow, then $F_s(X)$ is the project's continuation value at s. Alternatively, X_t can measure incentives, allowing us to compare the timing of cash flows and incentives.

We plot the durations of current PPS β and the stock of promised future incentives z together with their values in Figure 7. Figure 2 shows that z is mean-reverting on average:

its drift is positive at low z and negative at high z. Consistent with this mean reversion, the durations of both β_t and z_t decline as z rises through most of the continuation region. However, near $z = 1$, both durations turn upward. At the absorbing endpoint, $\beta = z = 1$ remains constant, so both durations return to $1/(r+\lambda)$. The final panel compares z_t/β_t with $D(\beta_t)$ and shows that the ratio of promised future incentives to current PPS is not a good proxy for PPS duration: it does not capture either mean reversion or the endpoint dynamics implied by the adding-up constraint.

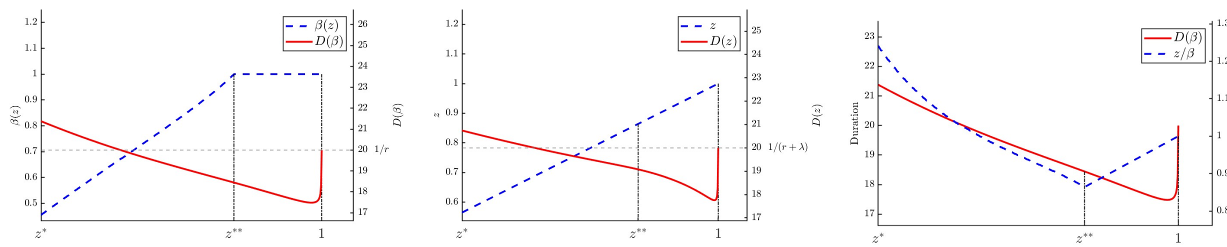

(a) β and Duration of β (b) z and Duration of z (c) Duration of β and z/β

Figure 7: Duration

The left panel of this figure plots β (left y-axis) and the duration of β (right y-axis). The dotted horizontal line is $1/r$ (right y-axis). The middle panel plots z (left y-axis) and the duration of z (right y-axis). The right panel plots the duration of β (left y-axis) and z/β (right y-axis). z^* represents the point at which $v(z)$ is maximized, and z^{**} represents the left boundary of the region where $\beta = 1$. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

Next, Figure 8 plots the durations of total project cash flows dY_t, the agent's compensation dM_t, and the principal's net cash flows $dY_t - dM_t$.[14] Total project cash flow reflects both instantaneous effort and the accumulated quality generated by durable effort. Its duration declines through most of the continuation region, as does the duration of compensation, before both turn upward near the absorbing endpoint. Compensation responds more sharply because its expected flow depends quadratically on β_t and z_t, making it especially sensitive to the decline in incentives from high interior states. The principal's net-cash-flow duration displays the opposite pattern in the simulation and rises with z. When z is high, mean

[14] Expected project cash flow and expected compensation are strictly positive. The principal's expected net cash flow can be negative near $z = 1$ because incentive and risk-compensation costs can exceed contemporaneous project cash flow. As a result, its reported duration should be interpreted as the generalized signed duration defined above.

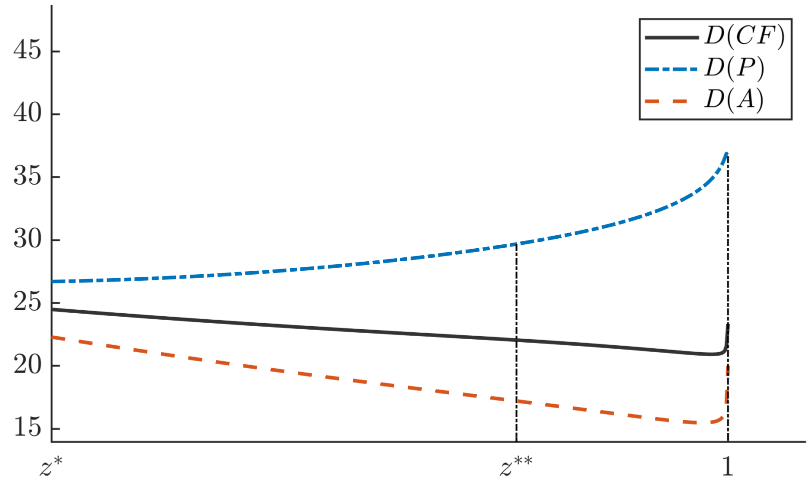

Figure 8: Duration of Total Cash Flows, Principal's Cash Flows, and Agent's Compensation

The figure plots the duration of total cash flows dY_t, principal's cash flows $dY_t - dM_t$, and agent's compensation dM_t. The initial Q is set to be 0. z^* represents the point at which $v(z)$ is maximized, and z^{**} represents the left boundary of the region where $\beta = 1$. Parameter values are $r = 0.05$, $\rho = 0.1$, $\gamma = 1$, $\sigma = 4$, $\theta = 1.5$, and $\lambda = 0$.

reversion lowers future incentive costs and raises the principal's continuation value, shifting a greater share of her net cash flows toward the future.

In the benchmark economies of Subsection 2.4, incentive policies and expected compensation flows are constant, so they have duration $1/r$ when $\lambda = 0$. The full model instead generates incentive and compensation durations both above and below $1/r$, so the optimal contract produces relatively short- or long-duration outcomes after different cash-flow histories. Actual project cash-flow duration also reflects the transition of asset quality from the common initialization $Q_0 = 0$; the comparisons across z in Figure 8 hold that initialization fixed and isolate how the optimal contract changes this duration across incentive states.

Over most of the continuation region, a positive shock lowers z and increases the durations of the agent's compensation, PPS, and total project cash flows, while decreasing the duration of the principal's share of cash flows (e.g., outside equity). These signs reverse locally near the upper boundary; at the absorbing endpoint, $\phi(1) = 0$, so cash-flow shocks do not change z. Over the main interior region, the compensation prediction is consistent with Gopalan, Milbourn, Song, and Thakor (2014), who find that executive-compensation duration is longer for firms with better recent stock performance. To our knowledge, the remaining duration predictions have not been tested directly and provide directions for future research.

5.3 Comparison of the Mechanism

The most theoretically similar paper to ours is Marinovic and Varas (2019) (hereafter MV). The two models share CARA utility with hidden savings, Brownian cash flows, a single output that pools two hidden actions, and a reduction to one scaled state variable for long-term incentives – MV's $z \equiv -p/W$ is the direct analogue of our z_t. This resemblance is substantive, but the persistent actions have different technologies and economic roles. Manipulation in MV's model raises cash flows today while subtracting a persistent stock later and is weakly negative-NPV under MV's Condition 1. Our durable action has no contemporaneous payoff and raises future cash flows. Reversing the value of the persistent action is not a relabeling: it invalidates MV's benchmark results and changes the long-term incentive state from a device that restrains a harmful action into one that induces productive investment.

The distinction is visible in the value functions: in MV's model, eliminating manipulation would weakly benefit the principal and make the long-term incentive state irrelevant (MV, Section IV). MV's contract uses long-term incentives to balance productive effort against manipulation, sometimes tolerating manipulation when prevention requires excessive post-retirement compensation. In our model, long-term incentives enter the principal's payoff directly through production, $\hat{q}_t = z_t/\theta$, so removing the durable action destroys value. The incentive state is determined by the marginal output of durable effort relative to the cost of compensating future volatility, producing a stationary interior optimum z^* with $v'(z^*) = 0$ and $z^* > \beta^*$. This structure also permits our central comparison of incentive costs. We give the two actions equal present values and identical effort costs and pool them in one output signal, so differences in their optimal provision arise from the agency problem. The actions in MV's model differ in value, cost, and timing by construction, so they do not provide the same comparison between the costs of current and future incentives.

In MV's model, a known retirement date makes the problem nonstationary: compensation delivered after retirement imposes risk without inducing further effort, so vesting accelerates and the target for long-term incentives declines toward zero. MV establishes that the drift of z_t is negative when manipulation is positive, while MV's numerical solutions show declining expected long-term incentives and a sensitivity to performance that becomes more negative as retirement approaches. In our stationary model, z_t drifts upward from z^*, ϕ is weakly negative everywhere and strictly negative away from z^*, and the contract moves between back- and front-loaded schedules. Section 4 further separates the mechanisms. Our stochastic termination extinguishes the project's cash flows and contractual promises together; retirement in MV's model ends productive effort while the project and potential post-retirement compensation continue. As a result, random project termination does not reproduce MV's known-retirement effect and instead increases initial reliance on promised future incentives.

These dynamics yield distinct empirical predictions: MV predicts that pay duration falls and manipulation rises with tenure, while we predict that pay duration falls while current PPS and promised future incentives rise after poor performance. In particular, Gopalan, Milbourn, Song, and Thakor (2014) document both margins: executive-compensation duration is lower for longer-tenured executives and for firms with worse recent performance. Our model, in turn, does not speak to the vesting and retirement effects that motivate MV's analysis; the two models thus explain complementary features of observed practices.

Our mechanism is likewise distinct from that of Gryglewicz et al. (2020) and Hackbarth et al. (2022), which feature optimally induced short- or long-termism. In those studies, the optimal mix of short- and long-term incentives varies with performance history because of complementarities either between shocks to the two performance signals or between the costs of the two actions. In our model, short- and long-term incentives are inevitably entangled and jointly designed, and the dynamic variation between more front-loaded and more back-loaded packages of incentives arises from how the costs of providing short- and long-term incentives endogenously vary with the observed output history.

Our mechanism also differs from those used in variable-effort extensions of DeMarzo and Sannikov (2006). In those models, the agency problem leads to financial constraints – the principal must liquidate the project after repeated bad news – so success reduces the agency problem and allows the principal to increase incentives. In turn, project values, incentives (when applicable), and the agent's expected compensation all increase after good news. Success also pushes liquidation further away, so cash-flow duration increases as incentives rise. In our model, the principal chooses current PPS and the stock of promised future incentives jointly. She prefers the value-maximizing, back-loaded initial schedule but must deliver the promises it creates. Good news increases the principal's value by reducing both current PPS and the stock of promised future incentives. Over most of the continuation region, cash-flow duration increases as those two incentive measures decline.

Our model of building asset quality also generates predictions opposite to those of learning models such as He et al. (2017) and Cisternas (2018). In those models, favorable output changes beliefs about project profitability and raises incentives, producing the analogue of $\phi \geq 0$. Cisternas (2018) interpret this response in part as an incentive-ratchet effect. Success relaxes the agency problem, so incentives and effort move toward their benchmark levels. In our model, a positive cash-flow innovation reduces the stock of promised future incentives and, over most of the continuation region, current PPS. The principal's value rises because future incentive costs decline, not because incentives and effort increase. The opposite signs of these incentive responses distinguish models of building asset quality from models of learning about asset quality.

6 Conclusion

In this paper, we study a setting in which a principal delegates management to an agent who can choose to enhance day-to-day performance or invest in long-term improvements. The optimal levels of instantaneous effort and durable quality investment are linked through incentives, without cost or production complementarities. The principal observes only cash flows and must use that single signal to incentivize both inputs. Current PPS adds up over time to induce investment in quality; conversely, promised future PPS induces repeated instantaneous effort as those promises come due. By controlling the level and timing of PPS, the principal can make the incentive schedule relatively front- or back-loaded, encouraging instantaneous effort or investment in quality.

We find that promised future incentives are initially cheaper to provide than current PPS – there is a flexibility advantage to generating incentives through promises about future PPS – so the optimal contract begins with a back-loaded PPS schedule. This cost advantage persists while the stock of promised future incentives remains low or moderate. Negative cash-flow surprises raise both promised future incentives and durable investment. However, once the stock of promised future incentives becomes sufficiently high, current PPS becomes relatively cheaper and the contract switches to a front-loaded schedule. These state-contingent dynamics generate mean reversion in cash flows and testable implications for the autocorrelation and durations of incentives and cash flows. The main results also survive exogenous stochastic project termination, even when termination extinguishes all contractual promises. Together, these results demonstrate that when productive short- and long-term actions share a common performance signal, the schedule of incentives is the central contracting instrument.

References

Bergemann, D. and P. Strack (2015). Dynamic revenue maximization: A continuous time approach. *Journal of Economic Theory 159*, 819–853.

Bloedel, A., R. V. Krishna, and B. Strulovici (2023). Persistent private information revisited. Working paper.

Bloedel, A. W., R. V. Krishna, and O. Leukhina (2025). Insurance and inequality with persistent private information. *Econometrica 93*(3), 821–857.

Bolton, P., J. Scheinkman, and W. Xiong (2006). Executive compensation and short-termist behaviour in speculative markets. *Review of Economic Studies 73*(3), 577–610.

Cisternas, G. (2018). Two-sided learning and the ratchet principle. *Review of Economic Studies 85*(1), 307–351.

Dechow, P. M., R. G. Sloan, and M. T. Soliman (2004). Implied equity duration: A new measure of equity risk. *Review of Accounting Studies 9*, 197–228.

DeMarzo, P. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance 61*, 2681 – 2724.

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

Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012). Dynamic ceo compensation. *Journal of Finance 67*(5), 1603–1647.

Feng, F. Z. and M. Westerfield (2024). Buying in and selling out: The dynamic returns to investing in expertise. Unpublished working paper, University of Washington.

Fernandes, A. and C. J. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory 91*(2), 223–247.

Froot, K. A., A. Perold, and J. C. Stein (1992). Shareholder trading practices and corporate investment horizons.

Fu, S. and R. V. Krishna (2019). Dynamic financial contracting with persistent private information. *RAND Journal of Economics 50*(2), 418–452.

Gonçalves, A. S. (2021). The short duration premium. *Journal of Financial Economics 141*(3), 919–945.

Gopalan, R., T. Milbourn, F. Song, and A. V. Thakor (2014). Duration of executive compensation. *Journal of Finance 69*(6), 2777–2817.

Gormsen, N. J. and E. Lazarus (2023). Duration-driven returns. *Journal of Finance 78*(3), 1393–1447.

Graham, J. R., C. R. Harvey, and S. Rajgopal (2005). The economic implications of corporate financial reporting. *Journal of Accounting and Economics 40*(1), 3–73.

Gryglewicz, S., S. Mayer, and E. Morellec (2020). Agency conflicts and short-versus long-termism in corporate policies. *Journal of Financial Economics 136*(3), 718–742.

Gryglewicz, S., S. Mayer, and E. Morellec (2024). The dynamics of loan sales and lender incentives. *Review of Financial Studies 37*(8), 2403–2460.

Hackbarth, D., A. Rivera, and T.-Y. Wong (2022). Optimal short-termism. *Management Science 68*(9), 6477–6505.

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

Hoffmann, F. and S. Pfeil (2021). Dynamic multitasking and managerial investment incentives. *Journal of Financial Economics 142*(2), 954–974.

Holmström, B. and J. Tirole (1993). Market liquidity and performance monitoring. *Journal of Political economy 101*(4), 678–709.

Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory 52*(1), 45–67.

Kapička, M. (2013). Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-order Approach. *Review of Economic Studies 80*(3), 1027–1054.

Kennan, J. (2001). Repeated bargaining with persistent private information. *Review of Economic Studies 68*(4), 719–755.

Krasikov, I. and R. Lamba (2021). A theory of dynamic contracting with financial constraints. *Journal of Economic Theory 193*, 105196.

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? a duration-based explanation of the value premium. *Journal of Finance 62*(1), 55–92.

Marinovic, I. and F. Varas (2019). CEO Horizon, Optimal Pay Duration, and the Escalation of ShortTermism. *Journal of Finance 74*, 2011 – 2053.

Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A myersonian approach. *Econometrica 82*(2), 601–653.

Sannikov, Y. (2014). Moral hazard and long-run incentives. Unpublished working paper, Princeton University.

Shleifer, A. and R. W. Vishny (1990). Equilibrium short horizons of investors and firms. *American Economic Review 80*(2), 148–153.

Stein, J. C. (1988). Takeover threats and managerial myopia. *Journal of Political Economy 96*(1), 61–80.

Stein, J. C. (1989). Efficient capital markets, inefficient firms: A model of myopic corporate behavior. *Quarterly Journal of Economics 104*(4), 655–669.

Von Thadden, E.-L. (1995). Long-term contracts, short-term investment and monitoring. *Review of Economic Studies 62*(4), 557–575.

Weber, M. (2018). Cash flow duration and the term structure of equity returns. *Journal of Financial Economics 128*(3), 486–503.

Williams, N. (2011). Persistent Private Information. *Econometrica 79*(4), 1233–1275.

Williams, N. (2015). A Solvable Continuous-time Dynamic Principal–agent Model. *Journal of Economic Theory 159*, 989–1015.

Zhang, Y. (2009). Dynamic Contracting with Persistent Shocks. *Journal of Economic Theory 144*(2), 635–675.

Zhu, J. Y. (2018). Myopic agency. *Review of Economic Studies 85*(2), 1352–1388.

Appendix

A Preliminaries

All proofs in this section of the appendix correspond to the model of Section 4 with $\lambda \geq 0$. We specialize the results to $\lambda = 0$ when we report them in Section 3.

A.1 Cash Flow Equivalence

First, we demonstrate that a_t and q_t are equally valuable for the principal, as in (3) and (28). Using integration by parts, we have

$$\begin{aligned}
\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}(r+\rho+\lambda)Q_s ds\right] &= \mathrm{E}_t\left[\int_t^\infty e^{-(r+\lambda)(s-t)}(r+\rho+\lambda)Q_s ds\right] \\
&= \mathrm{E}_t\left[\int_t^\infty -\frac{(r+\rho+\lambda)Q_s}{r+\lambda} de^{-(r+\lambda)(s-t)}\right] \\
&= \mathrm{E}_t\left[-\frac{(r+\rho+\lambda)Q_s}{r+\lambda} e^{-(r+\lambda)(s-t)}\bigg|_t^\infty\right] \\
&\quad - \mathrm{E}_t\left[\int_t^\infty -\frac{(r+\rho+\lambda)}{r+\lambda} e^{-(r+\lambda)(s-t)}\left(q_s - \rho Q_s\right) ds\right]
\end{aligned}$$

where we have assumed a transversality condition on Q_t that is valid in all three benchmark economies and the main model. Gathering terms, we have

$$\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}(r+\rho+\lambda)Q_s ds\right] = Q_t + \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)} q_s ds\right] \tag{A-1}$$

Adding the short-term action to both sides yields (28). To obtain (3), we repeat the calculations above with $\lambda = 0$ and without taking expectations.

A.2 Duration Expression

The duration of a cash flow $\{X\}_{s\geq t}$ starting from t is defined in (34), and the auxiliary function $F_t(X)$ is defined in (35). Using Fubini's theorem, we can rewrite the numerator of (34) as

$$\begin{aligned}
\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}(s-t)X_s ds\right] &= \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}\left(\int_t^s 1du\right) X_s ds\right] \\
&= \mathrm{E}_t\left[\int_t^\tau e^{-r(u-t)}\left(\int_u^\tau e^{-r(s-u)} X_s ds\right) du\right] \\
&= \mathrm{E}_t\left[\int_t^\tau e^{-r(u-t)} F_u(X) du\right]
\end{aligned} \tag{A-2}$$

Dividing by $F_t(X)$ gives $D_t(X) = F_t(F_s(X))/F_t(X)$, which is (36).

A.3 Cash Flow Shares

This subsection provides an auxiliary implementation rationale for the restriction $\beta_t \in [0,1]$ imposed in Section 2.3, analogous to the static argument in Innes (1990). Fix a contractual history and suppose that the agent can secretly destroy reported cash flow or supplement it dollar-for-dollar using private wealth. Let μ denote the resulting marginal change in reported cash flow. When $\mu > 0$, the agent incurs the private cost μ; when $\mu < 0$, destroying cash flow produces no private benefit.

At the fixed history, a marginal dollar of reported cash flow changes the agent's contractual payoff by β_t dollars. His net gain from manipulation is

$$\beta_t \mu - \mu^+, \qquad \mu^+ \equiv \max\{\mu, 0\}. \tag{A-3}$$

For $\mu < 0$, zero manipulation is optimal if and only if $\beta_t \geq 0$. For $\mu > 0$, it is optimal if and only if $\beta_t \leq 1$. Thus, a manipulation-free contract must have $\beta_t \in [0,1]$.

B Benchmark Economies (Proposition 1)

B.1 The First Best

Under the first best, the principal can direct the agent's choices to maximize her objective function, which is the expected discounted value of $dY_t - dM_t$. As in (8), we have $S_t = -\frac{1}{\gamma r}\ln(-\gamma r W_t)$. Let $V^{FB}(Q,W)$ denote the principal's continuation value after imposing this certainty-equivalent savings normalization. Thus,

$$dS_t = \left(-\frac{rW_t - u_t}{\gamma r W_t} + \frac{1}{2}\gamma r \sigma^2 \beta_t^2\right) dt + \beta_t \sigma dZ_t \tag{B-1}$$

$$dM_t = \left(-\frac{rW_t - u_t}{\gamma r W_t} + \frac{1}{2}\gamma r \sigma^2 \beta_t^2 + \frac{1}{\gamma}\ln(-\gamma r W_t) + c_t\right) dt + \beta_t \sigma dZ_t \tag{B-2}$$

where the first follows from Ito's lemma and the second from the law of motion for savings (4). The principal's problem is

$$V^{FB}(Q_t, W_t) = \max_{c,\, a,\, q,\, \beta} \mathrm{E}_t \Bigg[\int_t^\infty e^{-r(s-t)} \bigg(a_s + (r+\rho) Q_s \tag{B-3}$$

$$+ \frac{rW_s - u_s}{\gamma r W_s} - \frac{1}{2}\gamma r \sigma^2 \beta_s^2 - \frac{1}{\gamma}\ln(-\gamma r W_s) - c_s \bigg) ds \Bigg] \tag{B-4}$$

$$s.t. \quad dQ_t = (q_t - \rho Q_t)\, dt \tag{B-5}$$

$$dW_t = (rW_t - u(c_t, a_t, q_t))dt - \gamma r W_t \beta_t \sigma dZ_t \tag{B-6}$$

where W_t is the agent's continuation utility.

We then solve the principal's problem with the standard dynamic programming approach.

The HJB equation is

$$rV^{FB} = \max_{c,a,q,\beta} \Big\{ a + (r+\rho)Q + \frac{rW-u}{\gamma r W} - \frac{1}{2}\gamma r \sigma^2 \beta^2 - \frac{1}{\gamma}\ln(-\gamma r W) - c \qquad \text{(B-7)}$$
$$+ V_Q^{FB}(q-\rho Q) + V_W^{FB}(rW - u(c,a,q)) + \frac{1}{2}V_{WW}^{FB}\gamma^2 r^2 W^2 \beta^2 \sigma^2 \Big\}.$$

The solution is independent of promised utility:

$$V^{FB}(Q,W) = Q + \frac{1}{\theta r}. \qquad \text{(B-8)}$$

The associated policies are $\beta_t = 0$, $c_t = -\frac{1}{\gamma}\ln(-\gamma r W_t) + \frac{\theta}{2}(a_t^2 + q_t^2)$, $a_t = \frac{1}{\theta}$, and $q_t = \frac{1}{\theta}$. These policies yield $rW_t = u(c_t, a_t, q_t)$, $dS_t = dW_t = 0$, and $dM_t = \frac{\theta}{2}(a_t^2 + q_t^2)dt = \frac{1}{\theta}dt$. Equivalently, retaining the savings-accounting terms gives

$$V^{FB}(Q,S,W) = Q + S + \frac{1}{\gamma r}\ln(-\gamma r W) + \frac{1}{\theta r}. \qquad \text{(B-9)}$$

Under the normalization $S = -\frac{1}{\gamma r}\ln(-\gamma r W)$, the two middle terms cancel, yielding the reduced value above. A standard verification theorem confirms the sufficiency of the HJB solution.

B.2 The Agent Owns the Firm

When the agent owns the firm, cash flows are paid directly into the agent's savings ($dY_t = dM_t$). The agent's problem is

$$W(Q_t, S_t) = \max_{c,a,q} \mathrm{E}_t\left[\int_t^\infty e^{-r(s-t)} u(c,a,q)ds\right] \qquad \text{(B-10)}$$
$$s.t. \quad dS_t = (rS_t - c_t + a_t + (r+\rho)Q_t)\,dt + \sigma dZ_t \qquad \text{(B-11)}$$
$$dQ_t = (q_t - \rho Q_t)\,dt \qquad \text{(B-12)}$$

The HJB equation is

$$0 = \max_{c,a,q}\left\{u(c,a,q) - rW + W_Q(q-\rho Q) + W_S(rS - c + a + (r+\rho)Q) + \frac{1}{2}W_{SS}\sigma^2\right\} \qquad \text{(B-13)}$$

We conjecture and verify that the agent's lifetime utility is

$$W(Q_t, S_t) = -\frac{1}{\gamma r}\exp\left(-\gamma\left[rS_t + rQ_t + \frac{1}{\theta} - \frac{1}{2}\gamma r\sigma^2\right]\right) \qquad \text{(B-14)}$$

with the controls $\beta_t = 1$, $c_t = -\frac{1}{\gamma}\ln(-\gamma r W_t) + \frac{\theta}{2}(a_t^2 + q_t^2)$, $a_t = \frac{1}{\theta}$, and $q_t = \frac{1}{\theta}$. Further, $dW_t = -\gamma r \sigma W_t dZ_t$. We can rewrite (B-14) as $0 = \frac{1}{\gamma r}\ln(-\gamma r W_t) + S_t + Q_t + \frac{1}{\theta r} - \frac{1}{2}\gamma\sigma^2$. A standard verification theorem confirms the sufficiency of the HJB solutions.

B.3 The Observable-Quality Second Best

Here we solve the model with observable asset quality Q_t, i.e., the principal controls the quality improvement q_t while the agent needs to be incentivized to exert the unobservable instantaneous effort a_t. This is a special case of our full model proof in Sections C and D with $\Delta q = \Delta Q = \lambda = 0$.

We set up the Hamiltonian as

$$\mathcal{H}(\Delta c, \Delta a, \xi, \Delta S, p, j) = \xi u(\hat{c} + \Delta c, \hat{a} + \Delta a, q) + j^{\xi} \left(\frac{\Delta a}{\sigma} \right) \xi \\ + p^S \left(r \Delta S - \Delta c \right) \tag{B-15}$$

Note that we do not need the adjoint variable p^Q when the principal controls durable effort q, and the rest of the proof remains valid. Following the calculations in Sections C and D, we have $\hat{a}_t = \frac{\beta_t}{\theta}$ and $\hat{c}_t = -\frac{1}{\gamma} \ln(-\gamma r W_t) + \frac{\theta}{2}(\hat{a}_t^2 + q_t^2)$, where $\beta_t \in [0, 1]$ and $dW_t = -\sigma \gamma r W_t \beta_t d\hat{Z}_t$.

Proceeding from compensation to the agent (20), the principal's problem is

$$V(Q_t) = Q_t + \max_{\beta, q} \mathrm{E}_t \left[\int_t^\infty e^{-r(s-t)} \left(\frac{1}{\theta} \beta_s + q_s - \frac{1}{2\theta} \beta_s^2 - \frac{\theta}{2} q_s^2 - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right] \tag{B-16}$$

which can be maximized pointwise with respect to β and q to find $\beta_t = \frac{1}{1+\sigma^2 \gamma r \theta}$, $\hat{a}_t = \frac{1}{\theta} \frac{1}{1+\sigma^2 \gamma r \theta}$, $q_t = \frac{1}{\theta}$, and $\hat{c}_t = -\frac{1}{\gamma} \ln(-\gamma r W_t) + \frac{\theta}{2} \left(\hat{a}_t^2 + q_t^2 \right)$.

C The Agent's Problem

We continue with the "hat" notation described in Section 2.2. A variable or expectation with a hat, as in $\hat{c}_t$, represents the principal's recommendation, the agent's report, or an expectation conditioned on those variables and cash flows (i.e., the principal's information set). Plain letters, as in c_t, represent the agent's true decisions and expectations conditioned on the true state of the project (i.e., the agent's information set). We use Δ to denote differences; for example, $\Delta S_t = S_t - \hat{S}_t$ and $\Delta c_t = c_t - \hat{c}_t$.

Because the recommended actions and the agent's deviations must be consistent with the technology available (Assumption 1), we can take differences in (1) and (4) or (29) to obtain

$$d\Delta Q_t = (\Delta q_t - \rho \Delta Q_t) dt \tag{C-1}$$
$$d\Delta S_t = (r \Delta S_t - \Delta c_t) dt + \Delta \nu_t (dN_t - \lambda dt) \tag{C-2}$$

where N_t is the Poisson process with intensity λ and ν_t is the agent's investment in risky assets.

Because cash flows (Y) are observable, misreporting asset quality implies an offsetting misreport of the path of the Brownian motion. From (27), we have $dY_t = (a_t + (r + \rho + \lambda) Q_t) dt + \sigma dZ_t$ and $dY_t = (\hat{a}_t + (r + \rho + \lambda) \hat{Q}_t) dt + \sigma d\hat{Z}_t$, both on $t \in [0, \tau]$, so we also have

$$\left(\Delta a_t + (r + \rho + \lambda) \Delta Q_t \right) dt = - \sigma d\Delta Z_t. \tag{C-3}$$

We denote the Radon-Nikodym derivative between the probability measure induced by Z (the agent's information) and the measure induced by $\hat{Z}$ (the principal's information) with ξ_t, meaning $\mathrm{E}_t[(\cdot)_s] = \hat{\mathrm{E}}_t[\frac{\xi_s}{\xi_t}(\cdot)_s]$ for $s \geq t$ and

$$d\xi_t = \left(\frac{\Delta a_t + (r+\rho+\lambda)\Delta Q_t}{\sigma}\right)\xi_t d\hat{Z}_t. \tag{C-4}$$

for $t \in [0, \tau]$. For $t \geq \tau$, we have $d\xi_t = 0$ because the project has ended. Admissible deviations satisfy the standard measurability and integrability conditions; in particular, ξ and the adjoint stochastic integrals below are true martingales on every finite horizon. We also impose the usual costate transversality condition

$$\lim_{T\to\infty} \hat{\mathrm{E}}_t\left[e^{-r(T-t)}|p_T^{\xi}|\right] = 0, \tag{C-5}$$

which removes the terminal term when the ξ-costate equation is integrated forward.

We now prove Proposition 2 and the first part of Proposition 5 by finding the necessary conditions for incentive compatibility (Section C.1) and showing they are sufficient (Section C.2).

C.1 Required Incentives

The agent's goal is to maximize lifetime expected utility, given by (6). We rewrite the objective function as

$$\mathrm{E}_t\left[\int_t^{\infty} e^{-r(s-t)}u(c_s, a_s, q_s)ds\right] = \hat{\mathrm{E}}_t\left[\int_t^{\infty} e^{-r(s-t)}\frac{\xi_s}{\xi_t}u(\hat{c}_s + \Delta c_s, \hat{a}_s + \Delta a_s, \hat{q}_s + \Delta q_s)ds\right] \tag{C-6}$$

where we have used the change of measure ξ in (C-4) to go from E_t to $\hat{\mathrm{E}}_t$ and our Δ notation means $c_s = \hat{c}_s + \Delta c_s$. Then, the agent maximizes (C-6) subject to (C-1), (C-2), and (C-4).

We set up the Hamiltonian as

$$\begin{aligned}\mathcal{H}(\Delta c, \Delta a, \Delta q, \xi, \Delta Q, \Delta S, p, j, k) =& \xi u(\hat{c} + \Delta c, \hat{a} + \Delta a, \hat{q} + \Delta q) \\ &+ j^{\xi}\left(\frac{\Delta a + (r+\rho+\lambda)\Delta Q}{\sigma}\right)\xi \\ &+ p^Q(\Delta q - \rho\Delta Q) \\ &+ p^S(r\Delta S - \Delta c) + k^S \Delta\nu\lambda \end{aligned} \tag{C-7}$$

Taking first order conditions and evaluating along the optimal path ($\Delta c^* = \Delta a^* = \Delta q^* = 0$, $\xi^* = 1$, and $\Delta Q^* = \Delta S^* = 0$) gives

$$[\Delta c] \qquad -\gamma u(\hat{c}, \hat{a}, \hat{q}) = p^S \tag{C-8}$$

$$[\Delta a] \qquad \theta\gamma\hat{a}u(\hat{c}, \hat{a}, \hat{q}) = -j^{\xi}/\sigma \tag{C-9}$$

$$[\Delta q] \qquad \theta\gamma\hat{q}u(\hat{c}, \hat{a}, \hat{q}) = -p^Q \tag{C-10}$$

Optimality for $\Delta\nu$ requires $k^S = 0$.

The next step is to characterize the stochastic differential equations of the adjoint variables. Taking the derivative of $\mathcal{H}$ with respect to the state variables and evaluating along the optimal path, we obtain

$$dp_t^\xi = \Big(rp_t^\xi - u(\hat{c}_t, \hat{a}_t, \hat{q}_t)\Big)\, dt + j_t^\xi d\hat{Z}_t + k^\xi(dN_t - \lambda dt) \tag{C-11}$$

$$dp_t^Q = \left((r+\rho)p_t^Q - \frac{(r+\rho+\lambda)j_t^\xi}{\sigma}\right) dt + j_t^Q d\hat{Z}_t + k^Q(dN_t - \lambda dt) \tag{C-12}$$

$$dp_t^S = q_t^S d\hat{Z}_t + k^S(dN_t - \lambda dt). \tag{C-13}$$

Continuing, optimality for $\Delta\nu$ gives $k^S = 0$, so (C-13) makes p^S a local martingale. The standard integrability and adjoint transversality conditions imposed on admissible contracts make it a true martingale. The first-order condition (C-8), $p^S = -\gamma u(\hat{c}, \hat{a}, \hat{q})$, therefore makes the recommended flow utility a true martingale as well. The transversality condition for the ξ-costate and (C-11) give the conditional-present-value representation

$$p_t^\xi = \hat{\mathrm{E}}_t\left[\int_t^\infty e^{-r(s-t)} u(\hat{c}_s, \hat{a}_s, \hat{q}_s) ds\right], \tag{C-14}$$

and the martingale property of u evaluates it directly:

$$p_t^\xi = \int_t^\infty e^{-r(s-t)} \hat{\mathrm{E}}_t[u_s] ds = \frac{u_t}{r}. \tag{C-15}$$

Thus $rp_t^\xi = u_t$ and (C-8) gives $p_t^S = -\gamma r p_t^\xi$. Comparing the semimartingale decompositions of p^S and $-\gamma r p^\xi$ and recalling $k^S = 0$ gives $k^\xi = 0$. Finally, (C-14) is exactly the continuation-utility objective (6) along an incentive-compatible contract, so $p_t^\xi = W_t$. Consequently,

$$p_t^S = -\gamma r W_t, \qquad rW_t = u(\hat{c}_t, \hat{a}_t, \hat{q}_t). \tag{C-16}$$

We next observe that p_t^Q must jump to zero at τ because no additional incentives are possible, so $k^Q = -p^Q$. By redefining $j^\xi = -\sigma\gamma r W\beta$, and $j^Q = \sigma p^Q \eta$, we have:

$$dW_t = -\sigma\gamma r W_t \beta_t d\hat{Z}_t \tag{C-17}$$

$$dp_t^Q = (r+\rho+\lambda)(p_t^Q + \gamma r W_t \beta_t)dt + \sigma p_t^Q \eta_t d\hat{Z}_t - p^Q dN_t \tag{C-18}$$

$$\hat{a}_t = -\frac{j_t^\xi}{\sigma\theta\gamma u(\cdot)} = \frac{\sigma\gamma r W \beta_t}{\sigma\theta\gamma r W_t} = \frac{\beta_t}{\theta} \tag{C-19}$$

The restriction $\beta_t \in [0,1]$ is imposed as part of the contracting environment in Section 2.3. Its boundedness makes the stochastic exponential defining W_t a true martingale on every finite horizon, while ξ_t is a true martingale by the admissibility condition above. Integrating

(C-18), we have

$$p_t^Q = \hat{\mathrm{E}}_t \left[\int_t^\tau e^{-(r+\rho)(s-t)} (-(r+\rho+\lambda)\gamma r W_s \beta_s) ds \right] \tag{C-20}$$
$$= \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho+\lambda)(s-t)} (-(r+\rho+\lambda)\gamma r W_s \beta_s) ds \right],$$

which is the discounted value of future incentives. Since W_t is a martingale and $0 \le \beta_t \le 1$, we have $p_t^Q \in [0, -\gamma r W_t]$. We can reduce the two state variables p_t^Q and W_t to one:

$$z_t = \frac{p_t^Q}{-\gamma r W_t} \tag{C-21}$$
$$dz_t = (r+\rho+\lambda)(z_t - \beta_t + \sigma^2 \gamma r \beta_t \phi_t) dt + (r+\rho+\lambda)\sigma \phi_t d\hat{Z}_t - z_t dN_t \tag{C-22}$$

where the second equation follows from applying Ito's lemma to the first and by using $\phi_t = \frac{1}{r+\rho+\lambda} z_t(\gamma r \beta_t + \eta_t)$. From the bounds on p_t^Q, we also have $z_t \in [0,1]$ and (C-10) becomes

$$\hat{q}_t = \frac{z_t}{\theta} \tag{C-23}$$

Finally, combining $rW_t = u(\hat{c}_t, \hat{a}_t, \hat{q}_t)$ with (C-19) and (C-23) yields

$$\hat{c}_t = -\frac{1}{\gamma} \ln(-\gamma r W_t) + \frac{1}{2\theta}\beta_t^2 + \frac{1}{2\theta} z_t^2 \tag{C-24}$$

Optimality for $\Delta\nu_t$ leaves the agent indifferent over the recommended termination-contingent portfolio, so we normalize $\hat{\nu}_t = 0$. Combining (C-24) with the savings equation (29) (or (4) when $\lambda = 0$) and the certainty-equivalent normalization (8) then yields

$$dM_t = \left[\frac{1}{2}\sigma^2 \gamma r \beta_t^2 + \frac{1}{2\theta}\left(\beta_t^2 + z_t^2\right) \right] dt + \beta_t \sigma dZ_t. \tag{C-25}$$

C.2 A Bound on Gains from Deviation

We construct an upper bound for the payoff after deviation $(\Delta S, \Delta Q)$ of the form[15]:

$$\mathcal{J}(W, \Delta S, \Delta Q, z) = W \exp\left(-\gamma r \Delta S - \gamma r z \Delta Q\right), \tag{C-26}$$

where W is the equilibrium level of continuation utility. This upper bound establishes the sufficiency of the necessary conditions from Section C.1 because it is tight at $\Delta S = \Delta Q = 0$.

We include the standard infinite-horizon admissibility condition used in the verification below. For every admissible deviation, the change-of-measure density is a true martingale on finite horizons, the stochastic integrals below are locally square-integrable, and

$$\left\{ e^{-r(t\wedge\tau)} \mathcal{J}_{t\wedge\tau} : t \ge 0 \right\} \text{ is of class } D, \qquad \lim_{T\to\infty} \mathrm{E}\left[e^{-rT} |\mathcal{J}_T| \mathbf{1}_{\{T<\tau\}} \right] = 0. \tag{C-27}$$

[15]This is analogous to the bound in Sannikov (2014), He et al. (2017), or Marinovic and Varas (2019), extended to our setting.

For the baseline model without termination, take $\tau = \infty$, so the second condition is simply $\mathrm{E}[e^{-rT}|\mathcal{J}_T|] \to 0$. The same conditions are imposed conditionally after every stopping time, as required for sequential incentive compatibility. They rule out Ponzi-type deviations whose discounted continuation term fails to vanish, without placing a finite-horizon restriction on feasible actions.

Using (C-3) and (C-4) to change the probability measure, (10) and (31) become

$$\begin{aligned} dW_t =& -\gamma r W_t \beta_t \left(\Delta a + (r+\rho+\lambda)\Delta Q_t\right) dt - \sigma\gamma r W_t \beta_t dZ_t && \text{(C-28)} \\ dz_t =& (r+\rho+\lambda)\left(z_t + \sigma^2 \gamma r \beta_t \phi_t + \Delta a_t \phi_t + (r+\rho+\lambda)\Delta Q_t \phi_t - \beta_t\right) dt && \text{(C-29)} \\ &+ (r+\rho+\lambda)\sigma\phi_t dZ_t - z_t dN_t. \end{aligned}$$

We also have laws of motion for $d\Delta Q_t$ (C-1) and $d\Delta S_t$ (C-2).

To show that $\mathcal{J}$ is an upper bound on the agent's payoff, it is sufficient to show that $u(\hat{c}+\Delta c, \hat{a}+\Delta a, \hat{q}+\Delta q) - r\mathcal{J} + E\left[d\mathcal{J}\right]\frac{1}{dt} \le 0$. Calculating from (C-26), we have

$$\begin{aligned} E\left[d\mathcal{J}\right]\frac{1}{dt} =& \mathcal{J}\Big[-\gamma r \beta_t \Delta a_t - \gamma r (r+\rho+\lambda)\beta_t \Delta Q_t - \gamma r (r\Delta S_t - \Delta c_t) - \gamma r z_t (\Delta q_t - \rho \Delta Q_t) \\ &- \gamma r \Delta Q_t (r+\rho+\lambda)\left(z_t + \sigma^2\gamma r \beta_t \phi_t + \Delta a_t \phi_t + (r+\rho+\lambda)\Delta Q_t \phi_t - \beta_t\right) \\ &+ \frac{1}{2}\gamma^2 r^2 (\Delta Q_t)^2 \sigma^2 (r+\rho+\lambda)^2 \phi_t^2 + \sigma^2\gamma^2 r^2 (r+\rho+\lambda)\beta_t\phi_t\Delta Q_t \\ &+ \lambda\left(\gamma r \Delta\nu + e^{-\gamma r\Delta\nu + \gamma r z \Delta Q} - 1\right)\Big] && \text{(C-30)} \end{aligned}$$

From (12) we have

$$\begin{aligned} u(\hat{c}+\Delta c, \hat{a}+\Delta a, \hat{q}+\Delta q) =& rW_t \exp\left(-\gamma\Delta c_t + \frac{\gamma\theta}{2}\Delta a_t^2 + \frac{\gamma\theta}{2}\Delta q_t^2 + \gamma\beta_t\Delta a_t + \gamma z_t \Delta q_t\right) \\ =& r\exp\Big(-\gamma\Delta c_t + \frac{\gamma\theta}{2}\Delta a_t^2 + \frac{\gamma\theta}{2}\Delta q_t^2 + \gamma\beta_t\Delta a_t + \gamma z_t\Delta q_t \\ &+ \gamma r\Delta S_t + \gamma r z_t \Delta Q_t\Big)\mathcal{J} && \text{(C-31)} \end{aligned}$$

Since $\mathcal{J} < 0$, we divide and reverse the inequality, so $u(\hat{c}+\Delta c, \hat{a}+\Delta a, \hat{q}+\Delta q) - r\mathcal{J} + E\left[d\mathcal{J}\right]\frac{1}{dt} \le 0$ if and only if

$$\begin{aligned} 0 \le& -\gamma r \beta_t \Delta a_t - \gamma r (r+\rho+\lambda)\beta_t \Delta Q_t - \gamma r (r\Delta S_t - \Delta c_t) - \gamma r z (\Delta q_t - \rho\Delta Q_t) \\ &- \gamma r \Delta Q_t \left(r+\rho+\lambda\right)\left(z_t + \sigma^2\gamma r\beta_t\phi_t + \Delta a_t\phi_t + (r+\rho+\lambda)\Delta Q_t\phi_t - \beta_t\right) \\ &+ \frac{1}{2}\gamma^2 r^2 (\Delta Q_t)^2\sigma^2 (r+\rho+\lambda)^2\phi_t^2 + \sigma^2\gamma^2 r^2 (r+\rho+\lambda)\beta_t\phi_t\Delta Q_t \\ &+ r\exp\left(-\gamma\Delta c_t + \frac{\gamma\theta}{2}\Delta a_t^2 + \frac{\gamma\theta}{2}\Delta q_t^2 + \gamma\beta_t\Delta a_t + \gamma z_t\Delta q_t + \gamma r\Delta S_t + \gamma r z_t \Delta Q_t\right) \\ &+ \lambda\left(\gamma r\Delta\nu + e^{-\gamma r\Delta\nu + \gamma r z\Delta Q} - 1\right) - r && \text{(C-32)} \end{aligned}$$

Using $e^x \geq 1 + x$, it is sufficient to show that

$$\begin{aligned}
0 \leq & -\gamma r\beta_t\Delta a_t - \gamma r(r+\rho+\lambda)\beta_t\Delta Q_t - \gamma r(r\Delta S_t - \Delta c_t) - \gamma r z(\Delta q_t - \rho\Delta Q_t) \\
& - \gamma r\Delta Q_t\,(r+\rho+\lambda)\left(z_t + \sigma^2\gamma r\beta_t\phi_t + \Delta a_t\phi_t + (r+\rho+\lambda)\Delta Q_t\phi_t - \beta_t\right) \\
& + \frac{1}{2}\gamma^2 r^2(\Delta Q_t)^2\sigma^2(r+\rho+\lambda)^2\phi_t^2 + (-\sigma\gamma r\beta_t)\,\sigma(r+\rho+\lambda)\phi(-\gamma r\Delta Q) \\
& + r\left(1 - \gamma\Delta c_t + \frac{\gamma\theta}{2}\Delta a_t^2 + \frac{\gamma\theta}{2}\Delta q_t^2 + \gamma\beta_t\Delta a_t + \gamma z_t\Delta q_t + \gamma r\Delta S_t + \gamma r z_t\Delta Q_t\right) \\
& + \lambda\,(\gamma r z_t\Delta Q_t) - r
\end{aligned} \tag{C-33}$$

Canceling terms, this is equivalent to

$$\begin{aligned}
0 \leq & -\gamma r(\Delta Q_t)^2(r+\rho+\lambda)^2\phi_t - \gamma r(r+\rho+\lambda)\Delta a_t\Delta Q_t\phi_t \\
& + \frac{1}{2}\gamma^2 r^2(\Delta Q_t)^2\sigma^2(r+\rho+\lambda)^2\phi_t^2 + r\frac{\gamma\theta}{2}(\Delta a_t)^2 + r\frac{\gamma\theta}{2}(\Delta q_t)^2
\end{aligned} \tag{C-34}$$

$$\begin{aligned}
= & -\gamma r(\Delta Q_t)^2(r+\rho+\lambda)^2\phi_t + \frac{1}{2}\left(\sigma\gamma r(r+\rho+\lambda)\Delta Q\phi - \frac{1}{\sigma}\Delta a\right)^2 \\
& + \frac{\theta\sigma^2\gamma r - 1}{2\sigma^2}(\Delta a_t)^2 + r\frac{\gamma\theta}{2}(\Delta q_t)^2
\end{aligned} \tag{C-35}$$

which is always met if $\phi_t \leq 0$ and $\theta\sigma^2\gamma r - 1 \geq 0$, the maintained primitive restriction in Section 2.2.

It remains to turn the generator inequality into the promised global bound. Let ϑ_n localize all Brownian and compensated-jump stochastic integrals in Itô's formula for $e^{-rt}\mathcal{J}_t$, and put

$$\sigma_{n,T} \equiv T \wedge \tau \wedge \vartheta_n. \tag{C-36}$$

The generator inequality above and optional sampling give

$$\mathrm{E}\left[\int_0^{\sigma_{n,T}} e^{-rt}u(c_t, a_t, q_t)dt + e^{-r\sigma_{n,T}}\mathcal{J}_{\sigma_{n,T}}\right] \leq \mathcal{J}_0. \tag{C-37}$$

The stopped stochastic integrals have zero expectation. For fixed T, the class-D condition in (C-27) permits $n \to \infty$ in the terminal term, while monotone convergence applied to $-u \geq 0$ handles the running-utility term. Hence (C-37) holds with $\sigma_{n,T}$ replaced by $T \wedge \tau$.

At termination, z jumps to zero and savings jump by $\Delta\nu$. Thereafter the agent only consumes from private savings. The CARA/risk-free-savings problem has continuation value

$$\mathcal{J}_\tau = W_\tau \exp(-\gamma r\Delta S_\tau), \tag{C-38}$$

which is exactly (C-26) evaluated after the termination jumps. Finally, decompose

$$e^{-r(T\wedge\tau)}\mathcal{J}_{T\wedge\tau} = e^{-r\tau}\mathcal{J}_\tau\mathbf{1}_{\{\tau\leq T\}} + e^{-rT}\mathcal{J}_T\mathbf{1}_{\{T<\tau\}}. \tag{C-39}$$

The class-D and transversality conditions in (C-27), together with monotone convergence for the negative flow utility, now permit $T \to \infty$ and yield

$$\mathrm{E}\left[\int_0^\tau e^{-rt}u(c_t, a_t, q_t)dt + e^{-r\tau}\mathcal{J}_\tau\right] \le \mathcal{J}_0, \tag{C-40}$$

with the terminal term omitted when $\tau = \infty$. The left side is the agent's entire lifetime utility, including optimal consumption from savings after project termination. At the recommended policy, all deviations are zero, so $\mathcal{J} = W$ and every inequality above is an equality. Thus no admissible deviation improves the agent's payoff, completing the sufficiency argument.

D The Principal's Problem

The principal's value is the discounted value of cash flows net of compensation to the agent. Using (C-25) for compensation, we have

$$V(Q_t, z_t) = \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}\left(a_s + (r+\rho+\lambda)Q_s - \frac{1}{2}\sigma^2\gamma r\beta_s^2 - \frac{1}{2\theta}\beta_s^2 - \frac{1}{2\theta}z_s^2\right)ds\right] \tag{D-1}$$

Next, we use (28), (C-19), and (C-23) to obtain

$$V(Q_t, z_t) = \; Q_t + \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}\left(\frac{1}{\theta}\beta_s + \frac{1}{\theta}z_s - \frac{1}{2\theta}\beta_s^2 - \frac{1}{2\theta}z_s^2 - \frac{\sigma^2\gamma r}{2}\beta_s^2\right)ds\right] \tag{D-2}$$

To find the principal's value, we analyze

$$v(z_t) = \max_{\beta,\phi}\mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}\left(\frac{1}{\theta}\beta_s + \frac{1}{\theta}z_s - \frac{1}{2\theta}\beta_s^2 - \frac{1}{2\theta}z_s^2 - \frac{\sigma^2\gamma r}{2}\beta_s^2\right)ds\right] \tag{D-3}$$

The HJB equation for $v(z)$ is

$$\begin{aligned}(r+\lambda)v(z) = \max_{\beta,\,\phi}\Bigg\{&\frac{1}{\theta}\beta + \frac{1}{\theta}z - \frac{1}{2\theta}\beta^2 - \frac{1}{2\theta}z^2 - \frac{\sigma^2\gamma r}{2}\beta^2 \\ &+ v'(z)(r+\rho+\lambda)\left(z - \beta + \sigma^2\gamma r\beta\phi\right) + \frac{1}{2}v''(z)(r+\rho+\lambda)^2\phi^2\sigma^2\Bigg\}\end{aligned} \tag{D-4}$$

To assess the HJB equation and boundary conditions, define the auxiliary value $v_{stop}(z)$ delivered by the stationary feasible feedback

$$\beta_s = z_s = z, \qquad \phi_s = 0, \qquad s < \tau. \tag{D-5}$$

Under this feedback, (31) gives $dz_s = 0$ before termination, so the incentive state is stopped at its initial value. Equivalently, the integral representation (C-20)–(C-21) makes a stationary constant incentive schedule satisfy $\beta = z$. Therefore,

$$v_{stop}(z) = \frac{2}{\theta}\frac{z}{r+\lambda} - \left(\frac{1}{\theta} + \frac{\sigma^2\gamma r}{2}\right)\frac{z^2}{r+\lambda} \tag{D-6}$$

The endpoint values follow from the state constraint, not merely from the bounds on β

and z. As shown formally in Lemma OA.3, a viable evolution started at $z = 0$ must have $(\beta, \phi) = (0, 0)$ almost everywhere, while one started at $z = 1$ must have $(\beta, \phi) = (1, 0)$ almost everywhere. Both endpoints are absorbing, and the boundary conditions are

$$v(0) = v_{stop}(0) = 0 \tag{D-7}$$

$$v(1) = v_{stop}(1) = \frac{1}{(r + \lambda)} \left(\frac{1}{\theta} - \frac{\sigma^2 \gamma r}{2} \right) \tag{D-8}$$

The formal analysis of the HJB system, including the remainder of the properties stated in Proposition 3, requires substantially longer and more technical proofs, which are presented in the Online Appendix. Section OA.1 formulates the relaxed, state-constrained control problem and establishes value bounds, endpoint absorption, concavity and continuity, the dynamic-programming and viscosity equations, strict concavity, and C^2 regularity on $(z^*, 1)$. Section OA.2 then proves smooth fit and the endpoint policy limits, constructs an invariant weak solution under the feedback controls, and verifies that the relaxed value is attained by an incentive-compatible stationary Markov contract. Finally, Section OA.3 derives additional properties of the optimal policies, including their endpoint expansions and primitive sufficient conditions for unique thresholds at which the incentive cap binds and the contract switches between back-loading and front-loading.

Online Appendix for

The Optimal Schedules of Incentives and Cash Flows

OA.1 Properties of the Value Function

This section establishes the properties of the principal's value function used in Propositions 3 and 5. We first formulate a relaxed control problem and establish basic value bounds, concatenation, and endpoint absorption (Lemmas OA.1–OA.3). The spread argument in Lemma OA.4 yields concavity; Lemmas OA.5 and OA.6 then establish continuity on the closed state space and the existence of an interior maximizer z^*. Lemmas OA.7 and OA.8 provide the dynamic-programming and viscosity characterizations. Finally, Lemmas OA.10, OA.13, and OA.16 establish strict concavity and C^2 regularity on the continuation region $(z^*, 1)$. Section OA.2 subsequently verifies that the relaxed value is attained by an incentive-compatible contract.

OA.1.1 The Relaxed Control Problem

We begin with the necessary conditions in Proposition 2, as extended to project termination in Proposition 5. A *relaxed admissible pair* is a pair of processes (β_t, ϕ_t), adapted to the filtration generated by $\hat{Z}$, such that $\beta_t \in [0, 1]$, $\phi_t \in \mathbb{R}$, both processes satisfy the integrability requirements of these propositions, and the pre-termination state satisfies

$$dz_t = (r + \rho + \lambda)\left(z_t - \beta_t + \sigma^2\gamma r\beta_t\phi_t\right) dt + (r + \rho + \lambda)\sigma\phi_t \, d\hat{Z}_t, \tag{OA.1-1}$$

and remains in $[0, 1]$ almost surely. A *restricted admissible pair* is a relaxed admissible pair that also satisfies $\phi_t \leq 0$. Taking expectations over the independent, exponentially distributed termination time in (D-3), the principal's payoff from either type of pair started at z_0 is

$$J(z_0; \beta, \phi) = \mathrm{E}\left[\int_0^\infty e^{-(r+\lambda)s} f(\beta_s, z_s)\, ds\right], \tag{OA.1-2}$$

$$f(\beta, z) \equiv \frac{1}{\theta}\beta + \frac{1}{\theta}z - \frac{1}{2\theta}\beta^2 - \frac{1}{2\theta}z^2 - \frac{\sigma^2\gamma r}{2}\beta^2, \tag{OA.1-3}$$

and we define

$$\bar{v}(z_0) \equiv \sup_{(\beta,\phi) \text{ relaxed admissible}} J(z_0; \beta, \phi), \tag{OA.1-4}$$

$$v^-(z_0) \equiv \sup_{(\beta,\phi) \text{ restricted admissible}} J(z_0; \beta, \phi). \tag{OA.1-5}$$

Finally, let $v^{\mathrm{IC}}(z_0)$ denote the principal's value over all incentive-compatible contracts, net of the additive quality-in-place term Q_t and of the initial payment S_0. By the compensation process (C-25) and the cash-flow equivalence (28), the payoff of *any* incentive-compatible contract – not only the optimum – takes the form (D-3), so $v^{\mathrm{IC}}(z_0)$ is the supremum of $J(z_0;\beta,\phi)$ over the pairs (β,ϕ) generated by incentive-compatible contracts.

Lemma OA.1 (Relaxation and value bounds) *(i) For every $z_0\in[0,1]$, both admissible sets are non-empty: the stationary pair $\beta_t\equiv z_0$, $\phi_t\equiv 0$ delivers $J=v_{stop}(z_0)$, so $\bar v(z_0)\geq v^-(z_0)\geq v_{stop}(z_0)$. (ii) The flow payoff satisfies $|f(\beta,z)|\leq B\equiv \frac{2}{\theta}+\frac{\sigma^2\gamma r}{2}$ on $[0,1]^2$, so both auxiliary values are well-defined and bounded in absolute value by $B/(r+\lambda)$. (iii) The three values satisfy*

$$v^-(z)\leq v^{\mathrm{IC}}(z)\leq \bar v(z). \tag{OA.1-6}$$

Proof. (i) Under $\beta_t\equiv z_0$ and $\phi_t\equiv 0$, the drift in (OA.1-1) is $(r+\rho+\lambda)(z_t-z_0)$ and the volatility is zero, so $z_t\equiv z_0\in[0,1]$ solves the law of motion. This pair is restricted admissible and delivers $f(z_0,z_0)/(r+\lambda)=v_{stop}(z_0)$. (ii) On $[0,1]^2$, $f\leq \frac{1}{\theta}\beta+\frac{1}{\theta}z\leq\frac{2}{\theta}$ and $f\geq -\frac{1}{2\theta}\beta^2-\frac{1}{2\theta}z^2-\frac{\sigma^2\gamma r}{2}\beta^2\geq -\frac{1}{\theta}-\frac{\sigma^2\gamma r}{2}\geq -B$; discounting gives the asserted bounds. (iii) Every incentive-compatible contract admits the processes in the necessary direction of Propositions 2 and 5, and therefore generates a relaxed admissible pair with the same payoff; hence $v^{\mathrm{IC}}\leq\bar v$. Conversely, every restricted admissible pair satisfies $\phi\leq 0$ and the remaining hypotheses in the sufficient direction of these propositions; together with the compensation process (C-25), it defines an incentive-compatible contract with payoff J. Hence $v^-\leq v^{\mathrm{IC}}$. ∎

Two further structural facts are used repeatedly below. The first is that the admissible classes are stable under concatenation at stopping times, and that conditioning at a stopping time bounds continuation payoffs by the relaxed value. The second is that the endpoints of the state space are absorbing: the state constraint alone forces the boundary controls and pins down the boundary values of $\bar v$.

Lemma OA.2 (Concatenation and conditioning) *(i) Let (β,ϕ) be a relaxed admissible pair, let τ be a stopping time, and let $\{A_i\}_{i\geq 1}\subset\mathcal{F}_\tau$ be a countable partition of $\{\tau<\infty\}$ such that $z_\tau=\zeta_i$ on A_i for constants $\zeta_i\in[0,1]$, with (β^i,ϕ^i) a relaxed admissible pair for the initial state ζ_i. The strategy that follows (β,ϕ) on $[0,\tau)$ and (β^i,ϕ^i) on A_i thereafter, driven by the shifted Brownian motion $(\hat Z_{\tau+t}-\hat Z_\tau)_{t\geq 0}$, is relaxed admissible. The same statement holds for the restricted class.*

(ii) For every relaxed admissible pair (β,ϕ) and every stopping time τ, with the convention $e^{-(r+\lambda)\tau}\equiv 0$ on $\{\tau=\infty\}$,

$$J(z_0;\beta,\phi)\leq \mathrm{E}\left[\int_0^\tau e^{-(r+\lambda)s}f(\beta_s,z_s)\,ds+e^{-(r+\lambda)\tau}\,\bar v(z_\tau)\right], \tag{OA.1-7}$$

whenever $\bar v(z_\tau)$ is a well-defined random variable (in particular whenever z_τ takes countably many values, and, once the continuity in Lemma OA.5 is available, always).

Proof. (i) Adaptedness, the pointwise control constraints, and the state constraint $z\in[0,1]$ hold separately on the two segments and are preserved by pasting, and the integrability requirements hold segment by segment. By the strong Markov property, $(\hat Z_{\tau+t}-\hat Z_\tau)_{t\geq 0}$ is a Brownian motion independent of $\mathcal{F}_\tau$, so each continuation pair remains admissible for its initial state, and the pasted state path solves (OA.1-1) under the pasted controls. If every ingredient satisfies $\phi \leq 0$, so does the pasted strategy.

(ii) Split the payoff at τ and take conditional expectations. By the strong Markov property, the restriction of (β, ϕ) to $[\tau, \infty)$ is, conditionally on $\mathcal{F}_\tau$, a relaxed admissible pair for the initial state z_τ, so its conditional payoff is at most $\bar{v}(z_\tau)$ on $\{\tau < \infty\}$; on $\{\tau = \infty\}$ the tail payoff vanishes by discounting and $|f| \leq B$. Substituting into the tower property yields (OA.1-7), all interchanges being justified by $|f| \leq B$ and $|\bar{v}| \leq B/(r + \lambda)$. ■

Lemma OA.3 (Absorption at the endpoints) *Let (β, ϕ) be a relaxed admissible pair. If $z_0 = 1$, then almost surely $z_t = 1$ for all $t \geq 0$ and $(\beta_t, \phi_t) = (1, 0)$ for a.e. t. If $z_0 = 0$, then almost surely $z_t = 0$ for all $t \geq 0$ and $(\beta_t, \phi_t) = (0, 0)$ for a.e. t. Consequently,*

$$\bar{v}(0) = \frac{f(0,0)}{r + \lambda} = 0, \qquad \bar{v}(1) = \frac{f(1,1)}{r + \lambda} = v_{stop}(1). \tag{OA.1-8}$$

Proof. Consider $z_0 = 1$; the floor is symmetric and is summarized at the end. Write $k \equiv r + \rho + \lambda$ and decompose $z_t = 1 + D_t + M_t$ with

$$D_t = \int_0^t k \left(z_s - \beta_s + \sigma^2 \gamma r \beta_s \phi_s \right) ds, \tag{OA.1-9}$$

$$M_t = \int_0^t k \sigma \phi_s \, d\hat{Z}_s, \tag{OA.1-10}$$

$$V_t \equiv \langle M \rangle_t = \int_0^t k^2 \sigma^2 \phi_s^2 \, ds. \tag{OA.1-11}$$

Intervals of zero quadratic variation. If $V_\epsilon = 0$, then $\phi = 0$ a.e. on $[0, \epsilon]$, so $dz_t = k(z_t - \beta_t)\, dt \geq k(z_t - 1)\, dt$ there, and Grönwall's inequality gives $z_t \geq 1 + (z_0 - 1)e^{kt} = 1$. With the state constraint $z_t \leq 1$, we conclude $z \equiv 1$ on $[0, \epsilon]$.

Exclusion of positive quadratic variation. Let $\vartheta \equiv \inf\{t : V_t > 0\}$ and suppose $\mathrm{P}(\vartheta < \infty) > 0$. By the preceding argument and path continuity, $z \equiv 1$ on $[0, \vartheta]$. Replace (z, M, V) and the controls by their shifts at ϑ, which form an admissible evolution from state 1 relative to the shifted filtration; after this reduction, on an event E of positive probability, $z_0 = 1$ and $V_t > 0$ for every $t > 0$.

Two estimates produce a contradiction on E. First, since $z, \beta \in [0, 1]$, writing $Y_t \equiv 1 - z_t \geq 0$,

$$D_t \;\geq\; -k \int_0^t Y_s \, ds - k\sigma^2 \gamma r \int_0^t |\phi_s| \, ds \;\geq\; -k \int_0^t Y_s \, ds - \sigma \gamma r \sqrt{t} \sqrt{V_t}, \tag{OA.1-12}$$

where the last step uses the Cauchy–Schwarz inequality $\int_0^t |\phi_s| \, ds \leq \sqrt{t} \left(\int_0^t \phi_s^2 \, ds \right)^{1/2} = \sqrt{t}\sqrt{V_t}/(k\sigma)$. Second, by the Dambis–Dubins–Schwarz theorem, $M_t = W_{V_t}$ for a standard Brownian motion W (on an enlargement of the probability space if $V_\infty < \infty$). The local law of the iterated logarithm at zero, together with the monotonicity of $u \mapsto u \log\log(1/u)$ near zero, gives almost surely both a bound for all sufficiently small u and a favorable sequence

$u_n \downarrow 0$:

$$\sup_{u' \le u} |W_{u'}| \le 2\sqrt{2u \log\log(1/u)}, \qquad W_{u_n} \ge \tfrac{1}{2}\sqrt{2u_n \log\log(1/u_n)}. \tag{OA.1-13}$$

On E, continuity of V with $V_0 = 0$ and $V_t > 0$ for $t > 0$ supplies times $t_n \downarrow 0$ with $V_{t_n} = u_n$, hence, since $M_s = W_{V_s}$ and V is nondecreasing,

$$M_{t_n} \ge \tfrac{1}{2}\sqrt{2V_{t_n} \log\log(1/V_{t_n})}, \quad \sup_{s \le t} |M_s| \le 2\sqrt{2V_t \log\log(1/V_t)} \;\text{ for small } t. \tag{OA.1-14}$$

From $Y_t = -D_t - M_t$, the bound (OA.1-12), and Grönwall's inequality applied to the increasing majorant,

$$Y_t \le \left(\sigma\gamma r\sqrt{t}\sqrt{V_t} + \sup_{s \le t} |M_s|\right) e^{kt}. \tag{OA.1-15}$$

Substituting this back into (OA.1-12) through $\int_0^t Y_s\, ds \le t \sup_{s \le t} Y_s$ and evaluating at t_n,

$$z_{t_n} - 1 = D_{t_n} + M_{t_n} \ge M_{t_n} - \sigma\gamma r\sqrt{t_n}\sqrt{V_{t_n}}\left(1 + kt_n e^{kt_n}\right) - 2kt_n e^{kt_n}\sqrt{2V_{t_n} \log\log(1/V_{t_n})}.$$

As $n \to \infty$, $\log\log(1/V_{t_n}) \to \infty$ makes the middle term $o\big(\sqrt{V_{t_n} \log\log(1/V_{t_n})}\big)$, and $t_n \to 0$ makes the last term o of the same order. By (OA.1-14), therefore, $z_{t_n} - 1 \ge \big(\frac{1}{2} - o(1)\big)\sqrt{2V_{t_n} \log\log(1/V_{t_n})} > 0$ for n large, on a subset of E of positive probability. This violates the state constraint. Hence $\mathrm{P}(\vartheta < \infty) = 0$.

Conclusion. $V \equiv 0$, so the zero-quadratic-variation argument gives $z \equiv 1$ on $[0, \infty)$ and $\phi = 0$ a.e. Then $dz = 0$ forces the drift $k(1 - \beta_t)$ to vanish a.e., so $\beta = 1$ a.e., and the unique payoff from $z_0 = 1$ is $f(1, 1)/(r + \lambda) = v_{stop}(1)$.

At the floor $z_0 = 0$, the same decomposition applies with $z_t = D_t + M_t \ge 0$. When $V_\epsilon = 0$, $dz = k(z - \beta)\, dt \le kz\, dt$ and Grönwall's inequality force $z \equiv 0$ on $[0, \epsilon]$. If quadratic variation becomes positive, the drift bound becomes $D_t \le k\int_0^t z_s\, ds + \sigma\gamma r\sqrt{t}\sqrt{V_t}$, the law of the iterated logarithm supplies times with $M_{t_n} \le -\frac{1}{2}\sqrt{2V_{t_n} \log\log(1/V_{t_n})}$, and the same computation gives $z_{t_n} < 0$ for n large, violating the state constraint. Hence $z \equiv 0$, $\phi = 0$ a.e., the vanishing drift $-k\beta_t$ gives $\beta = 0$ a.e., and the payoff is $f(0, 0)/(r + \lambda) = 0$. ■

OA.1.2 Concavity and Continuity

The key observation is that ϕ is an unbounded control: the relaxed problem can make the stock of promised future incentives z_t arbitrarily volatile. Setting $\beta_t = 0$ at the same time shuts down the bilinear drift term $\sigma^2\gamma r\beta_t\phi_t$ in (OA.1-1). In the limit, the relaxed problem replaces z_0 with an instantaneous mean-preserving spread at zero flow cost. Its bounded value must therefore be midpoint-concave.

Lemma OA.4 (Spread Lemma) *Let $z_0 \in (0, 1)$ and $\delta > 0$ satisfy $[z_0 - \delta, z_0 + \delta] \subset (0, 1)$. Then*

$$\bar{v}(z_0) \ge \frac{1}{2}\bar{v}(z_0 + \delta) + \frac{1}{2}\bar{v}(z_0 - \delta). \tag{OA.1-16}$$

In particular, $\bar{v}$ is midpoint-concave on $(0,1)$.

Proof. Write $k \equiv r+\rho+\lambda$ and fix $M>0$ and $\epsilon>0$. Consider the following strategy: set $\beta_t = 0$ and $\phi_t = -M$ for $t<\tau$, where

$$\tau \equiv \inf\{t \geq 0 : z_t \notin (z_0-\delta, z_0+\delta)\}, \tag{OA.1-17}$$

and at τ continue with an ϵ-optimal relaxed admissible pair for the initial condition z_τ. On $[0,\tau)$, (OA.1-1) reduces to

$$dz_t = kz_t\,dt - k\sigma M d\hat{Z}_t, \tag{OA.1-18}$$

a linear SDE with bounded drift ($|kz_t| \leq k$ before τ) and constant volatility. Because paths of z are continuous, $z_\tau \in \{z_0-\delta, z_0+\delta\}$; since z_τ takes only two values, the continuation requires only two ϵ-optimal relaxed pairs, which exist by the definition of the supremum, and concatenation at the stopping time preserves relaxed admissibility (Lemma OA.2(i)). During $[0,\tau)$ we have $\beta_t = 0 \in [0,1]$, $\phi_t = -M \in \mathbb{R}$, and $z_t \in (z_0-\delta, z_0+\delta) \subset (0,1)$ by the definition of τ. Moreover, the exit-time bound below gives $M^2\mathrm{E}[\tau]<\infty$ for every fixed M, so the stopped volatility is square-integrable. Thus the whole strategy is relaxed admissible.

Exit-time bound. We first show that $\mathrm{E}[\tau] \leq \frac{2\delta^2}{k^2\sigma^2M^2}$ for M large. Let $g(z) \equiv (z-z_0+\delta)(z_0+\delta-z)$, so that $g \geq 0$ on $[z_0-\delta, z_0+\delta]$, $g(z_0)=\delta^2$, $g(z_\tau)=0$, $|g'| \leq 2\delta$, and $g''=-2$. By Itô's formula applied on $[0,\tau\wedge T]$ and taking expectations (all integrands are bounded),

$$\begin{aligned} 0 \leq \mathrm{E}\left[g(z_{\tau\wedge T})\right] &= \delta^2 + \mathrm{E}\left[\int_0^{\tau\wedge T}\left(g'(z_s)kz_s - k^2\sigma^2M^2\right)ds\right] \\ &\leq \delta^2 + \left(2\delta k - k^2\sigma^2M^2\right)\mathrm{E}[\tau\wedge T]. \end{aligned} \tag{OA.1-19}$$

For $M^2 \geq 4\delta/(k\sigma^2)$ we have $2\delta k \leq \frac{1}{2}k^2\sigma^2M^2$, so $\mathrm{E}[\tau\wedge T] \leq \frac{2\delta^2}{k^2\sigma^2M^2}$; letting $T\to\infty$ by monotone convergence gives the claim, and in particular $\tau<\infty$ almost surely.

Exit probabilities. We next show that $\mathrm{P}(z_\tau = z_0+\delta) = \frac{1}{2}+O(M^{-2})$. The process $z_t - \int_0^t kz_s ds$ is a martingale on $[0,\tau]$, so by optional stopping (with dominated convergence, z being bounded and $\mathrm{E}[\tau]$ finite),

$$\mathrm{E}[z_\tau] = z_0 + \mathrm{E}\left[\int_0^\tau kz_s\,ds\right], \qquad 0 \leq \mathrm{E}\left[\int_0^\tau kz_s\,ds\right] \leq k\,\mathrm{E}[\tau]. \tag{OA.1-20}$$

Writing $p_M \equiv \mathrm{P}(z_\tau = z_0+\delta)$, we have $\mathrm{E}[z_\tau] = z_0+\delta(2p_M-1)$, and therefore

$$\left|p_M - \frac{1}{2}\right| \leq \frac{k\mathrm{E}[\tau]}{2\delta} \leq \frac{\delta}{k\sigma^2M^2}. \tag{OA.1-21}$$

Payoff accounting. The payoff earned before τ is bounded by $B\mathrm{E}[\tau]$ in absolute value (Lemma OA.1(ii)). The continuation payoff is at least $\mathrm{E}\left[e^{-(r+\lambda)\tau}\left(\bar{v}(z_\tau)-\epsilon\right)\right]$; using $|e^{-x}-1| \leq x$ and $|\bar{v}| \leq B/(r+\lambda)$,

$$\mathrm{E}\left[e^{-(r+\lambda)\tau}\left(\bar{v}(z_\tau)-\epsilon\right)\right] \geq \mathrm{E}\left[\bar{v}(z_\tau)\right] - \epsilon - \left(B+(r+\lambda)\epsilon\right)\mathrm{E}[\tau]. \tag{OA.1-22}$$

Combining, and using $\bar{v}(z_0) \geq J$ for this relaxed admissible strategy,

$$\bar{v}(z_0) \;\geq\; p_M\,\bar{v}(z_0+\delta) + (1-p_M)\,\bar{v}(z_0-\delta) - \epsilon - (2B + (r+\lambda)\epsilon)\,\mathrm{E}[\tau]. \tag{OA.1-23}$$

Letting $M \to \infty$ (so that $\mathrm{E}[\tau] \to 0$ and $p_M \to \frac{1}{2}$, with $\bar{v}$ bounded) and then $\epsilon \to 0$ yields (OA.1-16). ■

The construction uses the interior of the state space only through the requirement $[z_0 - \delta, z_0 + \delta] \subset (0,1)$; Lemma OA.6 treats the endpoints. The drift correction is absorbed into the $O(M^{-2})$ terms, so the symmetric spread suffices.

Lemma OA.4 shows that $\bar{v}$ is midpoint-concave on $(0,1)$, and Lemma OA.1(ii) shows that $\bar{v}$ is bounded there. These two facts alone upgrade to full concavity and continuity, by a classical result on Jensen-concave functions.

Lemma OA.5 (Concavity and interior continuity) *$\bar{v}$ is concave and continuous on $(0,1)$. Moreover, $\bar{v}$ is locally Lipschitz: on any compact subinterval $[a,b] \subset (0,1)$, with $d_{ab} \equiv \frac{1}{2}\min\{a,\,1-b\}$,*

$$|\bar{v}(y) - \bar{v}(x)| \leq \frac{2B}{(r+\lambda)d_{ab}}\,|y-x| \qquad \textit{for all } x,y \in [a,b]. \tag{OA.1-24}$$

Proof. *Concavity and continuity.* By the Bernstein–Doetsch theorem, a midpoint-concave function on an open interval that is bounded below on a subinterval of positive measure is concave and continuous. Lemma OA.4 gives midpoint concavity on $(0,1)$ and Lemma OA.1(ii) gives $\bar{v} \geq -B/(r+\lambda)$ throughout $(0,1)$, so the hypotheses hold and $\bar{v}$ is concave and continuous on $(0,1)$.

Local Lipschitz bound. Let $[a,b] \subset (0,1)$ and $x, y \in [a,b]$ with $x < y$. Choose auxiliary points $x' = x - d_{ab}$ and $y' = y + d_{ab}$, both in $(0,1)$ by the definition of d_{ab}. Concavity makes the slopes of $\bar{v}$ non-increasing, so

$$\frac{\bar{v}(x) - \bar{v}(x')}{x - x'} \;\geq\; \frac{\bar{v}(y) - \bar{v}(x)}{y - x} \;\geq\; \frac{\bar{v}(y') - \bar{v}(y)}{y' - y}. \tag{OA.1-25}$$

The outer difference quotients have denominator d_{ab} and numerators bounded in absolute value by $2\sup|\bar{v}| \leq 2B/(r+\lambda)$ (Lemma OA.1(ii)), so the middle quotient is bounded in absolute value by $2B/((r+\lambda)d_{ab})$, which is (OA.1-24). ■

The local Lipschitz bound (OA.1-24) supplies the modulus of continuity used in the dynamic-programming and viscosity arguments below. Deriving this modulus from concavity also avoids the state-constraint difficulty associated with coupling two solutions under a common control: their difference evolves as $z_t - z'_t = (z_0 - z'_0)e^{(r+\rho+\lambda)t}$ and may drive the lower path outside $[0,1]$.

The state constraint pins down the value at the endpoints (Lemma OA.3). A separate barrier argument rules out downward jumps there; this is needed both for a closed-interval concavity result and for existence of the optimal initial state.

Lemma OA.6 (Closed-interval continuity and an interior maximizer) *The relaxed value satisfies*

$$\bar{v}(0) = 0, \qquad \bar{v}(1) = v_{stop}(1) = \frac{1}{r+\lambda}\left(\frac{1}{\theta} - \frac{\sigma^2\gamma r}{2}\right). \tag{OA.1-26}$$

Moreover, $\bar{v}$ is concave and continuous on $[0,1]$. Consequently, the set of maximizers of $\bar{v}$ is nonempty and its largest element z^* lies in $(0,1)$.

Proof. Write $d \equiv r + \lambda$ and $k \equiv r + \rho + \lambda$.

Endpoint values. By Lemma OA.3, the endpoints are absorbing: from $z_0 = 1$ every relaxed admissible pair has $z \equiv 1$ with $(\beta, \phi) = (1, 0)$ a.e., and from $z_0 = 0$ it has $z \equiv 0$ with $(\beta, \phi) = (0, 0)$ a.e. Therefore

$$\bar{v}(0) = \frac{f(0,0)}{d} = 0, \qquad \bar{v}(1) = \frac{f(1,1)}{d} = v_{stop}(1). \tag{OA.1-27}$$

Concavity on the closed interval. Lemma OA.5 gives concavity on $(0,1)$, while Lemma OA.1(i) and continuity of v_{stop} give

$$\liminf_{z \downarrow 0} \bar{v}(z) \geq v_{stop}(0) = \bar{v}(0), \qquad \liminf_{z \uparrow 1} \bar{v}(z) \geq v_{stop}(1) = \bar{v}(1). \tag{OA.1-28}$$

Approximating either endpoint by interior points in the chord inequality therefore extends concavity to $[0,1]$. This conclusion alone allows downward jumps at the endpoints, which we now exclude.

Endpoint continuity. Because f is continuously differentiable on the compact square $[0,1]^2$, there is an $L < \infty$ such that

$$f(\beta, z) - f(0,0) \leq L(\beta + z), \qquad f(\beta, z) - f(1,1) \leq L(1 - \beta + 1 - z). \tag{OA.1-29}$$

Choose $\alpha \in (0,1)$ sufficiently small that

$$d > k\alpha + A_\alpha, \qquad A_\alpha \equiv \frac{\sigma^2 \gamma^2 r^2 \alpha}{2(1-\alpha)}. \tag{OA.1-30}$$

For $U_0(z) = Cz^\alpha$ and $z \in (0,1)$, direct maximization of the quadratic in ϕ gives

$$\sup_{\phi \in \mathbb{R}} \left\{ k\sigma^2 \gamma r \beta U_0'(z)\phi + \frac{1}{2} k^2 \sigma^2 U_0''(z)\phi^2 \right\} = A_\alpha C \beta^2 z^\alpha. \tag{OA.1-31}$$

Also,

$$k(z - \beta)U_0'(z) = kC\alpha z^\alpha - kC\alpha\beta z^{\alpha - 1}. \tag{OA.1-32}$$

For use in this proof, write

$$\mathcal{L}^{\beta,\phi} U(z) \equiv k(z - \beta + \sigma^2 \gamma r \beta \phi)U'(z) + \frac{1}{2} k^2 \sigma^2 \phi^2 U''(z). \tag{OA.1-33}$$

Using (OA.1-29)–(OA.1-31), choose $\varepsilon > 0$ small and then C large enough that, for all $z \in (0, \varepsilon)$ and $\beta \in [0,1]$,

$$f(\beta, z) + \sup_{\phi \in \mathbb{R}} \mathcal{L}^{\beta,\phi} U_0(z) \leq dU_0(z). \tag{OA.1-34}$$

Indeed, the difference between the left and right sides is bounded above by

$$\left(L - kC\alpha z^{\alpha-1}\right)\beta + Lz - C\left(d - k\alpha - A_\alpha\right)z^\alpha. \tag{OA.1-35}$$

The first term is nonpositive near zero, while (OA.1-30) and $z = o(z^\alpha)$ make the sum of the last two terms nonpositive. Increase C if necessary so that $U_0(\varepsilon) \geq B/d \geq \bar{v}(\varepsilon)$.

For any relaxed admissible pair starting from $z \in (0, \varepsilon)$, let τ be the first hitting time of $\{0, \varepsilon\}$. Localize twice: at τ_n, the first hitting time of $\{1/n, \varepsilon\}$ (with $1/n < z$), so that U_0' and U_0'' are bounded on the stopped region, and at $\eta_m \equiv \inf\{t : \int_0^t \phi_s^2\, ds \geq m\}$, so that the stochastic integral below is square-integrable. Writing $\varrho \equiv \tau_n \wedge \eta_m \wedge T$, Itô's formula applied to $e^{-dt}U_0(z_t)$ together with the pathwise inequality (OA.1-34) gives

$$\mathrm{E}\left[\int_0^\varrho e^{-ds} f(\beta_s, z_s)\, ds + e^{-d\varrho}\, U_0(z_\varrho)\right] \leq U_0(z), \tag{OA.1-36}$$

the stopped stochastic integral having zero expectation. Send $m \to \infty$, then $n \to \infty$ and $T \to \infty$: admissibility gives $\eta_m \uparrow \infty$ almost surely, the integrands are bounded ($|f| \leq B$ and $0 \leq U_0 \leq U_0(\varepsilon)$ on the stopped region), and path continuity gives $z_{\tau_n \wedge T} \to z_\tau$ on $\{\tau < \infty\}$, so dominated convergence yields

$$\mathrm{E}\left[\int_0^\tau e^{-ds} f(\beta_s, z_s)\, ds + e^{-d\tau}\, U_0(z_\tau)\mathbf{1}_{\{\tau<\infty\}}\right] \leq U_0(z). \tag{OA.1-37}$$

Since $U_0(0) = \bar{v}(0)$ and $U_0(\varepsilon) \geq \bar{v}(\varepsilon)$, combining this with the conditioning bound (OA.1-7) at τ (where z_τ takes two values) shows that the total payoff of the pair is at most $U_0(z)$. Taking the supremum over relaxed admissible pairs gives

$$v_{stop}(z) \leq \bar{v}(z) \leq Cz^\alpha, \qquad z \in [0, \varepsilon], \tag{OA.1-38}$$

and hence $\lim_{z\downarrow 0} \bar{v}(z) = \bar{v}(0)$.

The upper endpoint is symmetric. Put $x = 1 - z$ and

$$U_1(z) = \bar{v}(1) + C(1-z)^\alpha. \tag{OA.1-39}$$

Writing $y = 1 - \beta$, its drift term is

$$k(z-\beta)U_1'(z) = -kC\alpha x^{\alpha-1}y + kC\alpha x^\alpha, \tag{OA.1-40}$$

and maximizing over ϕ again gives at most $A_\alpha C x^\alpha$. Thus, after reducing ε and increasing C so that $U_1(1-\varepsilon) \geq B/d$, the analogue of (OA.1-34) holds on $(1-\varepsilon, 1)$. The same stopped verification yields

$$v_{stop}(z) \leq \bar{v}(z) \leq \bar{v}(1) + C(1-z)^\alpha, \qquad z \in [1-\varepsilon, 1], \tag{OA.1-41}$$

so $\lim_{z\uparrow 1} \bar{v}(z) = \bar{v}(1)$. Hence $\bar{v}$ is continuous on $[0, 1]$.

Interior maximizer. Continuity on the compact interval implies that $\bar{v}$ attains its maximum. The strictly concave quadratic v_{stop} has its maximizer $z^*_{stop} = 2\beta^*/(1+\beta^*) \in (0, 1)$

and

$$\bar{v}(z^*_{stop}) \geq v_{stop}(z^*_{stop}) > \max\{v_{stop}(0), v_{stop}(1)\} = \max\{\bar{v}(0), \bar{v}(1)\}. \tag{OA.1-42}$$

Thus no maximizer is an endpoint, and the largest maximizer z^* lies in $(0, 1)$. ■

Lemmas OA.4–OA.6 establish concavity and continuity of the relaxed value on the closed state space. We continue to denote this value by $\bar{v}$ until Section OA.2 proves that it equals both v^- and v^{IC}.

OA.1.3 Dynamic Programming and the Relaxed HJB

For $\psi \in C^2(0, 1)$ and finite controls $\beta \in [0, 1]$, $\phi \in \mathbb{R}$, define

$$\mathcal{L}^{\beta,\phi}\psi(z) \equiv k\left(z - \beta + \sigma^2\gamma r\beta\phi\right)\psi'(z) + \frac{1}{2}k^2\sigma^2\phi^2\psi''(z), \qquad k \equiv r + \rho + \lambda, \tag{OA.1-43}$$

and the extended-real Hamiltonian

$$\bar{H}(z, p, X) \equiv \sup_{\beta\in[0,1],\, \phi\in\mathbb{R}} \left\{ f(\beta, z) + k\left(z - \beta + \sigma^2\gamma r\beta\phi\right)p + \frac{1}{2}k^2\sigma^2\phi^2 X \right\}. \tag{OA.1-44}$$

If $X < 0$, the supremum over ϕ is finite. If $X > 0$, it is infinite. If $X = 0$, it is infinite whenever $p \neq 0$ and finite when $p = 0$. At the degenerate point $p = X = 0$ the pointwise Hamiltonian is finite but discontinuous; below $\bar{H}^*$ and $\bar{H}_*$ denote the upper and lower semicontinuous envelopes. On the open set $X < 0$ both envelopes equal $\bar{H}$, while $\bar{H}^*(z, 0, 0) = +\infty$ and $\bar{H}_*(z, 0, 0) = \max_{\beta\in[0,1]} f(\beta, z)$. Thus the equation

$$(r + \lambda)\bar{v}(z) - \bar{H}\left(z, \bar{v}'(z), \bar{v}''(z)\right) = 0 \tag{OA.1-45}$$

is understood in the viscosity-envelope sense with this extended-real Hamiltonian.

Lemma OA.7 (Local dynamic programming) *Let $I = (\ell, u) \Subset (0, 1)$, let $z_0 \in I$, and let*

$$\tau_I \equiv \inf\{t \geq 0 : z_t \notin I\}. \tag{OA.1-46}$$

Then

$$\bar{v}(z_0) = \sup_{(\beta,\phi)} \mathrm{E}\left[\int_0^{\tau_I} e^{-(r+\lambda)s} f(\beta_s, z_s)\, ds + e^{-(r+\lambda)\tau_I}\bar{v}(z_{\tau_I})\mathbf{1}_{\{\tau_I<\infty\}}\right], \tag{OA.1-47}$$

where the supremum is over relaxed admissible pre-exit controls on $[0, \tau_I)$. On $\{\tau_I < \infty\}$, path continuity gives $z_{\tau_I} \in \{\ell, u\}$.

Proof. The "$\leq$" inequality is the conditioning bound in Lemma OA.2(ii), applied at τ_I; the terminal state takes only the two values ℓ and u, and the tail on $\{\tau_I = \infty\}$ vanishes by discounting and boundedness of f.

For the reverse inequality, fix any admissible pre-exit control on $[0, \tau_I)$. Choose ϵ-optimal relaxed continuations from the two deterministic states ℓ and u, and concatenate the corresponding continuation on the events $\{z_{\tau_I} = \ell\}$ and $\{z_{\tau_I} = u\}$. Lemma OA.2(i) preserves admissibility and the state constraint. The resulting global relaxed strategy delivers the displayed payoff up to ϵ, and $\epsilon \downarrow 0$ gives the reverse inequality. Thus no measurable selector over a continuum of continuation states is needed for the local DPP used below. ■

Lemma OA.8 (Relaxed viscosity equation) *The continuous, concave function $\bar{v}$ is a viscosity solution of (OA.1-45) on $(0,1)$ in the envelope sense: upper tests use $\bar{H}^*$ and lower tests use $\bar{H}_*$. The boundary values are those in Lemma OA.6.*

Proof. Fix $z_0 \in (0,1)$.

Subsolution. Let $\psi \in C^2$ touch $\bar{v}$ from above at z_0, and put $p = \psi'(z_0)$, $X = \psi''(z_0)$. If $\bar{H}^*(z_0, p, X) = +\infty$, the subsolution inequality is automatic. Otherwise $X < 0$ and $\bar{H}^* = \bar{H}$ near (z_0, p, X). Suppose, toward a contradiction, that for some $\delta > 0$,

$$(r+\lambda)\bar{v}(z_0) \geq \bar{H}(z_0, p, X) + 4\delta. \tag{OA.1-48}$$

After adding a small multiple of $(z - z_0)^4$, assume the contact is strict on the boundary of a small interval $I \Subset (0,1)$ centered at z_0: for some $\chi > 0$,

$$\bar{v}(z) \leq \psi(z) - \chi, \qquad z \in \partial I. \tag{OA.1-49}$$

Since $X < 0$, the quadratic term in ϕ is coercive from above. Hence, after shrinking I if necessary,

$$f(\beta, z) + \mathcal{L}^{\beta,\phi}\psi(z) - (r+\lambda)\psi(z) \leq -\delta \tag{OA.1-50}$$

for every $z \in I$, every $\beta \in [0,1]$, and every finite $\phi \in \mathbb{R}$.

Use the local DPP (OA.1-47) on I with an ϵ-optimal pre-exit control, where $\epsilon < \frac{1}{2}\min\{\delta/(r+\lambda), \chi\}$. Apply Itô's formula to $e^{-(r+\lambda)t}\psi(z_t)$ up to τ_I, localizing at

$$\eta_n \equiv \inf\left\{t : \int_0^t \phi_s^2 ds \geq n\right\} \wedge n \tag{OA.1-51}$$

and then sending $n \to \infty$. The stopped stochastic integral has mean zero; boundedness of ψ handles the terminal term, while the generator term is bounded above by (OA.1-50), so Fatou's lemma gives the limiting inequality. Combining (OA.1-50) with the boundary gap gives

$$0 \leq \mathrm{E}\left[\int_0^{\tau_I} e^{-(r+\lambda)t}\left(f(\beta_t, z_t) + \mathcal{L}^{\beta_t,\phi_t}\psi(z_t) - (r+\lambda)\psi(z_t)\right)dt\right] \tag{OA.1-52}$$

$$+ \mathrm{E}\left[e^{-(r+\lambda)\tau_I}\left(\bar{v}(z_{\tau_I}) - \psi(z_{\tau_I})\right)\mathbf{1}_{\{\tau_I < \infty\}}\right] + \epsilon \tag{OA.1-53}$$

$$\leq -\delta\,\mathrm{E}\left[\int_0^{\tau_I} e^{-(r+\lambda)t}dt\right] - \chi\,\mathrm{E}\left[e^{-(r+\lambda)\tau_I}\mathbf{1}_{\{\tau_I<\infty\}}\right] + \epsilon \tag{OA.1-54}$$

$$\leq -\min\{\delta/(r+\lambda), \chi\} + \epsilon < 0, \tag{OA.1-55}$$

a contradiction. Therefore

$$(r+\lambda)\bar{v}(z_0) \leq \bar{H}^*\left(z_0, \psi'(z_0), \psi''(z_0)\right). \tag{OA.1-56}$$

Supersolution. Let $\psi \in C^2$ touch $\bar{v}$ from below at z_0. Concavity of $\bar{v}$ implies $\psi''(z_0) \le 0$. Fix any finite constant control (β, ϕ) with $\phi \neq 0$, use it until the first exit from a small interval about z_0, and then concatenate an ϵ-optimal continuation from the two exit states. The local DPP and Itô's formula, with $\epsilon = o(\mathrm{E}\int_0^{\tau_I} e^{-(r+\lambda)t}dt)$, yield

$$(r+\lambda)\bar{v}(z_0) \ge f(\beta, z_0) + \mathcal{L}^{\beta,\phi}\psi(z_0) \tag{OA.1-57}$$

for every $\beta \in [0,1]$ and every finite $\phi \neq 0$. The case $\phi = 0$ follows by continuity of the right side as $\phi \to 0$.

If $\psi''(z_0) < 0$, taking the supremum in (OA.1-57) gives the viscosity supersolution inequality because $\bar{H}_* = \bar{H}$ there and the quadratic in ϕ has a finite maximum. If $\psi''(z_0) = 0$ and $\psi'(z_0) \neq 0$, set $\beta = 1$ and send ϕ to the sign that makes $k\sigma^2\gamma r\psi'(z_0)\phi$ positive. The right side of (OA.1-57) then diverges, contradicting its finite left side. Hence no lower test with zero curvature and nonzero slope exists. Finally, if $\psi''(z_0) = \psi'(z_0) = 0$, the ϕ terms vanish and taking the supremum over β gives

$$(r+\lambda)\bar{v}(z_0) \ge \max_{\beta\in[0,1]} f(\beta, z_0) = \bar{H}_*(z_0, 0, 0). \tag{OA.1-58}$$

These cases exhaust all lower tests and prove the claim. ■

Retaining the Hamiltonian in extended-real form allows zero-curvature lower tests to be handled directly by the per-control DPP inequality.

OA.1.4 Strict Concavity

Recall from Lemma OA.6 that z^* is the largest maximizer of the continuous, concave relaxed value $\bar{v}$.

Lemma OA.9 (Strict decrease) *$\bar{v}$ is strictly decreasing on $(z^*, 1)$.*

Proof. For $z^* < x < y < 1$, maximality of z^* gives $\bar{v}(x) < \bar{v}(z^*)$. Concavity makes secant slopes non-increasing, so

$$\frac{\bar{v}(y) - \bar{v}(x)}{y - x} \le \frac{\bar{v}(x) - \bar{v}(z^*)}{x - z^*} < 0. \tag{OA.1-59}$$

Thus $\bar{v}(y) < \bar{v}(x)$. ■

The per-control supersolution inequality (OA.1-57) also provides a uniform curvature bound for every smooth lower test.

Lemma OA.10 (Quantitative curvature) *Let $K = [a, b] \Subset (z^*, 1)$. There are constants $0 < m_K \le M_K < \infty$ and $\eta_K > 0$ such that, if $\psi \in C^2$ touches $\bar{v}$ from below at some $z \in K$, then*

$$-M_K \le \psi'(z) \le -m_K < 0, \qquad \psi''(z) \le -\eta_K < 0. \tag{OA.1-60}$$

In particular, $\bar{v}$ has no affine segment in $(z^*, 1)$ and is strictly concave there.

Proof. Choose a compact interval $K' \Subset (z^*, 1)$ containing K in its interior. The local Lipschitz bound in Lemma OA.5 supplies $M_K < \infty$. If a smooth function ψ touches a concave function from below at an interior point, comparison of left and right difference quotients shows that the concave function is differentiable there and $\psi'(z) = \bar{v}'(z)$. Concavity and Lemma OA.9 then give, uniformly for $z \in K$,

$$\psi'(z) = \bar{v}'(z) \le \frac{\bar{v}(a) - \bar{v}(z^*)}{a - z^*} \equiv -m_K < 0. \tag{OA.1-61}$$

Together with local Lipschitz continuity this proves the slope bounds in (OA.1-60).

The supersolution argument in Lemma OA.8 rules out $\psi''(z) = 0$ when $\psi'(z) \neq 0$, so write $X \equiv \psi''(z) < 0$ and $p \equiv \psi'(z)$. In the per-control supersolution inequality (OA.1-57), set $\beta = 1$ and maximize its right side over finite $\phi \in \mathbb{R}$. Direct calculation gives

$$\sup_{\phi \in \mathbb{R}} \left\{ k\sigma^2 \gamma r p \phi + \frac{1}{2} k^2 \sigma^2 X \phi^2 \right\} = \frac{\sigma^2 \gamma^2 r^2 p^2}{-2X}. \tag{OA.1-62}$$

Consequently,

$$\frac{\sigma^2 \gamma^2 r^2 p^2}{-2X} \le (r + \lambda)\bar{v}(z) - f(1, z) - k(z - 1)p. \tag{OA.1-63}$$

The right side is bounded above on K by

$$C_K \equiv 2B + kM_K < \infty, \tag{OA.1-64}$$

using Lemma OA.1(ii), $|f| \le B$, $|z - 1| \le 1$, and $|p| \le M_K$. Since $p^2 \ge m_K^2$, (OA.1-63) implies

$$-X \ge \frac{\sigma^2 \gamma^2 r^2 m_K^2}{2C_K} \equiv \eta_K > 0. \tag{OA.1-65}$$

This proves (OA.1-60). An affine segment would provide an affine lower test with zero curvature at each of its interior points, contradicting the result. For a concave function, absence of affine segments is equivalent to strict concavity. ■

Corollary OA.11 (Sign of the relaxed volatility choice) *At any point $z \in (z^*, 1)$ at which the relaxed HJB holds classically with $\bar{v}''(z) < 0$, its maximizing controls satisfy $\beta(z) > 0$ and $\phi(z) < 0$.*

Proof. For fixed β, the first-order condition in ϕ gives

$$\phi(z) = -\frac{\gamma r \beta(z) \bar{v}'(z)}{k \bar{v}''(z)}. \tag{OA.1-66}$$

After maximizing over ϕ, the derivative of the Hamiltonian with respect to β at $\beta = 0$ is

$$\frac{1}{\theta} - k\bar{v}'(z) > 0, \tag{OA.1-67}$$

so the maximizing β is strictly positive. Lemma OA.9 gives $\bar{v}'(z) < 0$, and $\bar{v}''(z) < 0$ by hypothesis; hence (OA.1-66) is strictly negative. ■

Lemma OA.10 excludes flat curvature in the viscosity subjets but does not exclude singular distributional curvature, because a concave kink has no smooth lower test at the kink. We first obtain the classical HJB and the policy formulas at the almost-everywhere Alexandrov points.

OA.1.5 Interior Classical Regularity

The available regularity theorem for uniformly nondegenerate one-dimensional control problems cannot be invoked here: the control ϕ is unbounded and the diffusion coefficient vanishes when $\phi = 0$. We therefore retain the regularity supplied directly by concavity and the viscosity equation.

Let $D^2\bar{v}$ denote the distributional second derivative of $\bar{v}$ on $(z^*, 1)$. Concavity makes it a nonpositive Radon measure, with Lebesgue decomposition

$$D^2\bar{v} = \bar{v}''_{ac}(z)\,dz - \nu_s, \qquad \nu_s \geq 0, \tag{OA.1-68}$$

where ν_s is singular with respect to Lebesgue measure. Its atoms represent ordinary concave kinks; it may also contain a singular continuous component.

Lemma OA.12 (Almost-everywhere classical structure) *There is a Borel set $R \subset (z^*, 1)$ of full Lebesgue measure such that $\bar{v}$ has a second-order Alexandrov expansion at every $z \in R$. For each compact $K \Subset (z^*, 1)$, the constants in Lemma OA.10 can be chosen so that, for almost every $z \in K$,*

$$-M_K \leq \bar{v}'(z) \leq -m_K < 0, \qquad \bar{v}''_{ac}(z) \leq -\eta_K < 0. \tag{OA.1-69}$$

At every $z \in R$, the relaxed HJB holds classically and has unique maximizers $\beta(z) > 0$ and $\phi(z) < 0$. These maximizers are Borel measurable on R and satisfy

$$\phi(z) = -\frac{\gamma r \beta(z) \bar{v}'(z)}{k \bar{v}''_{ac}(z)}. \tag{OA.1-70}$$

Proof. A concave function on an interval is differentiable and twice Alexandrov-differentiable almost everywhere. Thus, outside a null set, there are finite numbers p and X such that

$$\bar{v}(z+h) = \bar{v}(z) + ph + \frac{1}{2}Xh^2 + o(h^2). \tag{OA.1-71}$$

At such a point, for every $\epsilon > 0$, the quadratic polynomials with slope p and curvatures $X + \epsilon$ and $X - \epsilon$ are respectively local upper and lower test functions after restricting to a sufficiently small neighborhood. Lemma OA.10, applied to the lower test and followed by $\epsilon \downarrow 0$, yields the bounds in (OA.1-69); the density of the absolutely continuous part of $D^2\bar{v}$ equals this Alexandrov second derivative almost everywhere.

For small ϵ, both $X - \epsilon$ and $X + \epsilon$ are strictly negative. On this branch the envelopes in Lemma OA.8 coincide with $\bar{H}$, so the viscosity subsolution and supersolution inequalities give

$$\bar{H}(z,p,X+\epsilon) \geq (r+\lambda)\bar{v}(z) \geq \bar{H}(z,p,X-\epsilon). \tag{OA.1-72}$$

The Hamiltonian is finite and continuous on $X<0$, so letting $\epsilon \downarrow 0$ proves

$$(r+\lambda)\bar{v}(z) = \bar{H}(z,p,X) \tag{OA.1-73}$$

at every $z \in R$.

For $p<0$ and $X<0$, maximizing first over $\phi \in \mathbb{R}$ reduces the Hamiltonian to

$$\bar{H}(z,p,X) = \max_{\beta\in[0,1]} \left\{ f(\beta,z) + k(z-\beta)p - \frac{\sigma^2\gamma^2 r^2}{2}\beta^2 \frac{p^2}{X} \right\}. \tag{OA.1-74}$$

Put

$$A(z) \equiv \frac{1}{\theta} - kp > 0, \tag{OA.1-75}$$

$$\Gamma(z) \equiv \frac{1}{\theta} + \sigma^2\gamma r + \sigma^2\gamma^2 r^2 \frac{p^2}{X}. \tag{OA.1-76}$$

The part depending on β is $A(z)\beta - \frac{1}{2}\Gamma(z)\beta^2$. Its unique maximizer is

$$\beta(z) = \begin{cases} \min\{1, A(z)/\Gamma(z)\}, & \Gamma(z) > 0, \\ 1, & \Gamma(z) \leq 0. \end{cases} \tag{OA.1-77}$$

In particular, $\beta(z) > 0$. The strictly concave quadratic in ϕ then has the unique maximizer (OA.1-70); since $p<0$ and $X<0$, it is strictly negative. Measurability follows from measurability of the almost-everywhere derivatives and the displayed formulas. ■

Lemma OA.12 proves neither that $\nu_s = 0$ nor that the almost-everywhere selectors extend continuously across its support. Lemma OA.13 removes the atomic part of ν_s, and the subsequent Hamiltonian-gap argument removes the singular-continuous part.

A kink of the concave $\bar{v}$ (a jump of $\bar{v}'$) is precisely an atom of ν_s. Kinks are invisible to lower-test curvature estimates because no smooth lower test exists at a kink. Upper tests with arbitrarily negative curvature do exist, however, and their combination with the almost-everywhere classical equation excludes kinks. Define the *deterministic Hamiltonian* – the $\phi = 0$ branch of (OA.1-44) –

$$H_0(z,p) \equiv \max_{\beta\in[0,1]} \left\{ f(\beta,z) + k(z-\beta)p \right\}, \tag{OA.1-78}$$

which is jointly continuous and, as a pointwise maximum of functions affine in p, convex in p. Because f is strictly concave in β, the maximizer is unique,

$$\beta_0(p;z) = \max\left\{0,\ \min\left\{1,\ \frac{1/\theta - kp}{1/\theta + \sigma^2\gamma r}\right\}\right\}, \tag{OA.1-79}$$

strictly positive for $p \leq 0$, strictly decreasing in p wherever it is interior, and by Danskin's theorem $\partial H_0/\partial p = k\,(z - \beta_0(p;z))$.

Lemma OA.13 (No interior kinks) $\bar{v} \in C^1(z^*, 1)$. *Consequently, the singular measure* ν_s *in* (OA.1-68) *has no atoms in* $(z^*, 1)$.

Proof. Suppose $\bar{v}$ has a kink at $z_1 \in (z^*, 1)$: $\bar{v}'_-(z_1) > \bar{v}'_+(z_1)$. Both one-sided derivatives are strictly negative: by concavity and Lemma OA.9, $\bar{v}'_-(z_1) \leq \frac{\bar{v}(z_1)-\bar{v}(z^*)}{z_1-z^*} < 0$.

(i) Upper tests at the kink. Fix $p \in (\bar{v}'_+(z_1), \bar{v}'_-(z_1))$ and $N > 0$, and set $\psi(z) = \bar{v}(z_1) + p(z-z_1) - \frac{N}{2}(z-z_1)^2$. For $h > 0$ the subgradient inequality gives $\bar{v}(z_1+h) \leq \bar{v}(z_1) + \bar{v}'_+(z_1)h$, and $ph - \frac{N}{2}h^2 \geq \bar{v}'_+(z_1)h$ holds for $h \leq 2(p - \bar{v}'_+(z_1))/N$; symmetrically for $h < 0$ using $\bar{v}'_-(z_1)$. So ψ is a C^2 upper test at z_1, and the subsolution property (Lemma OA.8) gives $(r+\lambda)\bar{v}(z_1) \leq \bar{H}(z_1, p, -N)$. Maximizing the ϕ-quadratic at $X = -N$,

$$\bar{H}(z_1, p, -N) = \max_{\beta \in [0,1]} \left\{ f(\beta, z_1) + k(z_1 - \beta)p + \frac{\sigma^2\gamma^2 r^2 \beta^2 p^2}{2N} \right\} \leq H_0(z_1, p) + \frac{\sigma^2\gamma^2 r^2 p^2}{2N}.$$

Letting $N \to \infty$, and then extending to the closed interval by continuity of H_0 in p,

$$(r+\lambda)\bar{v}(z_1) \leq H_0(z_1, p) \qquad \text{for all } p \in [\bar{v}'_+(z_1), \bar{v}'_-(z_1)]. \tag{OA.1-80}$$

(ii) Limits from either side. At every $z \in R$ (Lemma OA.12), the classical equation (OA.1-73) holds, and since $\phi = 0$ is feasible in the supremum, $(r+\lambda)\bar{v}(z) = \bar{H}(z, \bar{v}'(z), \bar{v}''_{ac}(z)) \geq H_0(z, \bar{v}'(z))$. Because R has full measure, we may take $z_n \uparrow z_1$ with $z_n \in R$; then $\bar{v}'(z_n) \to \bar{v}'_-(z_1)$ (one-sided limits of the derivative of a concave function), and continuity of $\bar{v}$ and H_0 give $(r+\lambda)\bar{v}(z_1) \geq H_0(z_1, \bar{v}'_-(z_1))$. The sequence $z_n \downarrow z_1$ gives $(r+\lambda)\bar{v}(z_1) \geq H_0(z_1, \bar{v}'_+(z_1))$.

(iii) The squeeze. Combining (i) and (ii),

$$\max\left\{H_0(z_1, \bar{v}'_-),\ H_0(z_1, \bar{v}'_+)\right\} \ \leq\ (r+\lambda)\bar{v}(z_1) \ \leq\ \min_{p \in [\bar{v}'_+, \bar{v}'_-]} H_0(z_1, p). \tag{OA.1-81}$$

Since $H_0(z_1, \cdot)$ is convex, its maximum over the interval $[\bar{v}'_+, \bar{v}'_-]$ is attained at an endpoint, so the display forces $H_0(z_1, \cdot)$ to be *constant* on the interval. Hence $\partial H_0/\partial p = k(z_1 - \beta_0(p; z_1)) = 0$, i.e., $\beta_0(p; z_1) \equiv z_1$ on a nondegenerate interval of p. By (OA.1-79), β_0 is strictly decreasing in p wherever it is interior and equal to $1 \neq z_1$ at the corner (as $z_1 < 1$); it cannot be constant at the interior value z_1 on a nondegenerate interval. Contradiction. Hence $\bar{v}'$ has no jumps in $(z^*, 1)$; being monotone without jumps, it is continuous. ■

The exclusion follows from the model's strict effort costs: a kink would force the deterministic Hamiltonian to be affine across the entire jump interval of slopes, which the β^2-terms in f rule out. Lemma OA.13 leaves only a possible singular-continuous component of ν_s. We eliminate that component using the exact gap between the full Hamiltonian and its deterministic branch.

Throughout this argument write

$$d \equiv r + \lambda, \qquad k \equiv r + \rho + \lambda, \qquad C \equiv \frac{1}{\theta} + \sigma^2\gamma r, \qquad s' \equiv \frac{\sigma^2\gamma^2 r^2}{2}, \tag{OA.1-82}$$

and

$$h(z, p, \beta) \equiv f(\beta, z) + k(z - \beta)p, \qquad H_0(z, p) = \max_{\beta \in [0,1]} h(z, p, \beta). \tag{OA.1-83}$$

Lemma OA.14 (Exact Hamiltonian-gap inversion) *For $p < 0$ and $q \geq 0$, define*

$$G(z, p, q) \equiv \max_{\beta \in [0,1]} \{h(z, p, \beta) + s'\beta^2 q\} - H_0(z, p). \tag{OA.1-84}$$

Then $G(z, p, 0) = 0$, and $q \mapsto G(z, p, q)$ is continuous and strictly increasing from $[0, \infty)$ onto $[0, \infty)$. Denote its inverse by

$$q = Q(z, p, \Delta), \qquad \Delta \geq 0. \tag{OA.1-85}$$

The function Q is continuous, locally Lipschitz for $p < 0$, and $Q(z, p, 0) = 0$. Moreover, for $X < 0$,

$$\bar{H}(z, p, X) - H_0(z, p) = G\left(z, p, -\frac{p^2}{X}\right). \tag{OA.1-86}$$

Consequently, whenever $\Delta \equiv d\,w - H_0(z, p) > 0$, the equation

$$d\,w = \bar{H}(z, p, X) \tag{OA.1-87}$$

has the unique solution

$$X = \Phi(z, w, p) \equiv -\frac{p^2}{Q\big(z, p, d\,w - H_0(z, p)\big)} < 0. \tag{OA.1-88}$$

The map Φ is continuous on $\{p < 0,\ d\,w > H_0(z, p)\}$.

Proof. Put $A(p) \equiv 1/\theta - kp > 0$. Apart from terms independent of β, the maximand in (OA.1-84) is

$$A(p)\beta - \frac{1}{2}\big(C - 2s'q\big)\beta^2. \tag{OA.1-89}$$

It has the unique maximizer

$$\beta_q(p) = \begin{cases} \min\{1, A(p)/(C - 2s'q)\}, & C - 2s'q > 0, \\ 1, & C - 2s'q \leq 0. \end{cases} \tag{OA.1-90}$$

In particular, $\beta_q(p) \geq \beta_0(p) > 0$. The envelope theorem gives

$$\frac{\partial G}{\partial q}(z, p, q) = s'\beta_q(p)^2 > 0 \tag{OA.1-91}$$

on each smooth branch, with matching one-sided derivatives when the cap begins to bind. Hence G is strictly increasing; and since the choice $\beta = 1$ makes it diverge linearly in q, its range is $[0, \infty)$. The inverse is continuous. On compact subsets with $p < 0$, $\beta_0(p)$ is bounded away from zero, so $\partial G/\partial q$ is bounded away from zero; the explicit piecewise-quadratic formula then gives local Lipschitz continuity of Q.

Maximizing the original Hamiltonian first over ϕ gives

$$\bar{H}(z,p,X) = \max_{\beta\in[0,1]} \left\{ h(z,p,\beta) - s'\beta^2 \frac{p^2}{X} \right\}, \qquad X<0. \tag{OA.1-92}$$

This is (OA.1-86); inversion gives (OA.1-88) and the asserted continuity. ■

We apply the lemma to the value itself. Since $\bar{v} \in C^1(z^*,1)$, write

$$p(z) \equiv \bar{v}'(z) < 0, \qquad \Delta(z) \equiv d\bar{v}(z) - H_0\big(z,p(z)\big). \tag{OA.1-93}$$

Lemma OA.15 (Strict interior Hamiltonian gap) *For every $z \in (z^*,1)$,*

$$\Delta(z) > 0. \tag{OA.1-94}$$

Consequently, for each compact $K \Subset (z^*,1)$ there is a $\delta_K > 0$ such that $\Delta \geq \delta_K$ on K.

Proof. At almost every point in the set R of Lemma OA.12, put $X = p'_{ac} = \bar{v}''_{ac} < 0$. Equations (OA.1-73) and (OA.1-86) give

$$\Delta(z) = G\left(z, p(z), -\frac{p(z)^2}{p'_{ac}(z)}\right) > 0 \qquad \text{a.e.} \tag{OA.1-95}$$

The function Δ is continuous, so $\Delta \geq 0$ everywhere.

Suppose, toward a contradiction, that $\Delta(z_0) = 0$ for some $z_0 \in (z^*,1)$. Let

$$\beta_0(p) = \min\left\{1, \frac{1/\theta - kp}{C}\right\}, \qquad y(z) \equiv z - \beta_0(p(z)). \tag{OA.1-96}$$

Because p is continuous and nonincreasing, it is locally of bounded variation. The BV chain rule and Danskin's formula $H_{0p} = k(z - \beta_0(p))$ give, in the sense of distributional first derivatives,

$$D\Delta = \left[-\rho p - \frac{1-z}{\theta}\right] dz - k\, y\, Dp, \qquad Dp = p'_{ac} dz - \nu_s. \tag{OA.1-97}$$

By continuity of Q and $Q(z,p,0) = 0$, (OA.1-95) implies that $p'_{ac} = -p^2/Q(z,p,\Delta)$ is arbitrarily negative, uniformly almost everywhere in a sufficiently small neighborhood of z_0. If $y(z_0) > 0$, shrink the neighborhood so that y is bounded away from zero. The absolutely continuous density in (OA.1-97) is then strictly positive and its singular part is $ky\nu_s \geq 0$; hence Δ is strictly increasing there. This contradicts $\Delta \geq 0$ and $\Delta(z_0) = 0$ by looking immediately to the left. If $y(z_0) < 0$, the same argument makes Δ strictly decreasing and gives the symmetric contradiction to the right. Therefore

$$y(z_0) = 0. \tag{OA.1-98}$$

Since $z_0 < 1$, (OA.1-98) puts the deterministic maximizer strictly inside $(0,1)$ near z_0. On that neighborhood define

$$E(z) \equiv \Delta(z) + \frac{C}{2}\big(z - \beta_0(p(z))\big)^2. \tag{OA.1-99}$$

Completing the square in the deterministic Hamiltonian cancels all dependence on p and gives the exact identity

$$E(z) = d\bar{v}(z) - \frac{2z}{\theta} + \left(\frac{1}{\theta} + \frac{\sigma^2\gamma r}{2}\right) z^2. \tag{OA.1-100}$$

Thus $E \geq 0$ and $E(z_0) = 0$. Its distributional second derivative is

$$D^2E = dD^2\bar{v} + \left(\frac{2}{\theta} + \sigma^2\gamma r\right) dz = \left[dp'_{ac} + \frac{2}{\theta} + \sigma^2\gamma r\right] dz - d\nu_s. \tag{OA.1-101}$$

Shrink the neighborhood again. Since p stays bounded away from zero, Δ is uniformly small, and $Q(z,p,\Delta) \to 0$ uniformly as $\Delta \to 0$, the identity $p'_{ac} = -p^2/Q(z,p,\Delta)$ makes the absolutely continuous density in (OA.1-101) strictly negative. Its singular part is also nonpositive. Hence E is strictly concave on a neighborhood of z_0, which is incompatible with the interior minimum $E(z_0) = 0 \leq E(z)$. This final contradiction proves (OA.1-94). Compact separation follows from continuity. ■

Lemma OA.16 (Interior classical regularity) *The relaxed value satisfies*

$$\bar{v} \in C^2(z^*, 1), \qquad \bar{v}''(z) = \Phi\big(z, \bar{v}(z), \bar{v}'(z)\big) < 0. \tag{OA.1-102}$$

The relaxed HJB holds classically at every $z \in (z^*, 1)$; its maximizing controls are unique and continuous there, and

$$\beta(z) > 0, \qquad \phi(z) = -\frac{\gamma r \beta(z)\bar{v}'(z)}{k\bar{v}''(z)} < 0. \tag{OA.1-103}$$

In particular, the singular measure ν_s in (OA.1-68) vanishes.

Proof. Fix $K \Subset (z^*, 1)$. Lemma OA.15 keeps $(z, \bar{v}(z), \bar{v}'(z))$ in a compact subset of the domain of Φ. Hence

$$F(z) \equiv \Phi\big(z, \bar{v}(z), \bar{v}'(z)\big) \tag{OA.1-104}$$

is continuous on K. Because $\bar{v}$ is C^1, every upper or lower test at z has slope $\bar{v}'(z)$. The strict gap keeps the relevant curvatures on the finite $X < 0$ branch, where the envelopes in Lemma OA.8 coincide with $\bar{H}$. The monotonicity of $X \mapsto \bar{H}(z, \bar{v}'(z), X)$ and the viscosity inequalities therefore say exactly that $\bar{v}$ is a viscosity solution on K of

$$-u'' + F(z) = 0. \tag{OA.1-105}$$

Let $W \in C^2(K)$ satisfy $W'' = F$. Then $\bar{v} - W$ is a viscosity solution of $-u'' = 0$, and hence is both convex and concave, so it is affine. Therefore $\bar{v} \in C^2(K)$ and $\bar{v}'' = F$. Since K was arbitrary, the conclusion holds on $(z^*, 1)$. The policy conclusions follow from the unique maximizers in Lemma OA.14; their continuity follows from the displayed formulas and $\bar{v} \in C^2$. Finally, a C^2 function has no singular distributional curvature, so $\nu_s = 0$. ■

Lemma OA.16 completes the value-function analysis: the relaxed value is continuous and concave on $[0, 1]$, strictly concave and C^2 on $(z^*, 1)$, and admits unique continuous maximizing controls throughout that region. Section OA.2 establishes the endpoint behavior and closed-loop existence required for exact verification.

OA.2 Verification

This section proves exact attainment of the principal's value in Propositions 3 and 5. Lemma OA.17 gives a closed-loop verification criterion. Lemmas OA.18 and OA.19 establish smooth fit and the Hamiltonian identity at the lower endpoint; Corollary OA.21 and Lemma OA.22 identify the associated policy limits. Lemmas OA.23–OA.25 establish the upper-endpoint limits. Theorem OA.26 then constructs an invariant closed-loop diffusion and proves that the relaxed value is attained by an incentive-compatible Markov contract.

OA.2.1 Verification Criterion

Let the continuous interior maximizers in Lemma OA.16 be extended to $[z^*, 1]$, and define

$$b(z) \equiv k\left(z - \beta(z) + \sigma^2\gamma r\beta(z)\phi(z)\right), \qquad a(z) \equiv k\sigma\phi(z). \tag{OA.2-1}$$

Lemma OA.17 (Closed-loop verification criterion) *Suppose that the feedback extends to $[z^*, 1]$ with $\phi \le 0$ and is restricted admissible. Suppose also that, for every $z_0 \in [z^*, 1]$, the equation*

$$dz_t = b(z_t)dt + a(z_t)d\hat{Z}_t, \qquad z_0 \in [z^*, 1], \tag{OA.2-2}$$

admits a weak solution remaining in $[z^*, 1]$. If

$$\bar{v}'_+(z^*) = 0, \;\; (\beta(z^*), \phi(z^*)) = (\beta^*, 0), \;\; z^* > \beta^*, \;\; (r+\lambda)\bar{v}(z^*) = f(\beta^*, z^*), \tag{OA.2-3}$$

and the upper endpoint is either unattained from below in finite time or absorbing with $(\beta(1), \phi(1)) = (1, 0)$ and $(r+\lambda)\bar{v}(1) = f(1, 1)$, then the feedback attains

$$J(z_0; \beta(\cdot), \phi(\cdot)) = \bar{v}(z_0), \qquad z_0 \in [z^*, 1]. \tag{OA.2-4}$$

Consequently,

$$v^-(z) = v^{\mathrm{IC}}(z) = \bar{v}(z) = \max_{(\tilde{\beta}, \tilde{\phi}) \text{ relaxed admissible}} J(z; \tilde{\beta}, \tilde{\phi}), \qquad z \in [z^*, 1]. \tag{OA.2-5}$$

Proof. Extend $\bar{v}$ to the left of z^* by the constant $\bar{v}(z^*)$. Smooth fit makes this extension continuously differentiable at z^*, with locally absolutely continuous derivative below every level $1 - \varepsilon$. The generalized Itô formula therefore applies without a local-time term at z^*. On $(z^*, 1)$, Lemma OA.16 gives the classical HJB with the chosen feedback; at z^* the assumed Hamiltonian identity gives the same equality. Stopping below the upper endpoint and localizing the stochastic integral yields

$$\bar{v}(z_0) = \mathrm{E}\left[\int_0^T e^{-(r+\lambda)t} f(\beta(z_t), z_t)dt + e^{-(r+\lambda)T}\bar{v}(z_T)\right] \tag{OA.2-6}$$

after removing the localizations below 1. If the process reaches 1, the absorbing identity $(r+\lambda)\bar{v}(1) = f(1,1)$ supplies the continuation contribution; otherwise the stopping times converge to infinity. Boundedness of f and $\bar{v}$ permits $T \to \infty$, proving (OA.2-4). Since the feedback is restricted admissible, Propositions 2 and 5 make it incentive compatible. The sandwich (OA.1-6) then yields (OA.2-5). ∎

OA.2.2 Lower-Endpoint Behavior

Because z^* lies in the interior of the primitive state space $[0,1]$, the viscosity inequalities are available on both sides of the maximizer. Their combination with the deterministic Hamiltonian yields smooth fit and identifies the limiting feedback at the lower endpoint.

Lemma OA.18 (No corner at the optimal initial state) *$\bar{v}$ is differentiable at z^* with*

$$\bar{v}'(z^*) = 0. \tag{OA.2-7}$$

Proof. By concavity and the local Lipschitz property (Lemma OA.5), the one-sided derivatives

$$p_r \equiv \bar{v}'_+(z^*) = \lim_{z \downarrow z^*} \bar{v}'(z), \qquad p_\ell \equiv \bar{v}'_-(z^*) = \lim_{z \uparrow z^*} \bar{v}'(z) \tag{OA.2-8}$$

exist and are finite, the limits being taken along points of differentiability. Because z^* is a maximizer, $p_r \le 0 \le p_\ell$, the superdifferential of $\bar{v}$ at z^* is $[p_r, p_\ell]$, and the subgradient inequalities

$$\bar{v}(z^* + h) \le \bar{v}(z^*) + p_r h \;\; (h \ge 0), \qquad \bar{v}(z^* + h) \le \bar{v}(z^*) + p_\ell h \;\; (h \le 0) \tag{OA.2-9}$$

hold.

Intermediate corner slopes. Suppose $p_r < p_\ell$, and fix $s \in (p_r, p_\ell)$ and $N > 0$. By (OA.2-9), exactly as in part (i) of the proof of Lemma OA.13,

$$\psi_N(z) = \bar{v}(z^*) + s(z - z^*) - \frac{N}{2}(z - z^*)^2 \tag{OA.2-10}$$

is a C^2 upper test at z^* on the neighborhood $|z - z^*| \le 2\min\{s - p_r,\, p_\ell - s\}/N$, and the subsolution property together with

$$\bar{H}(z^*, s, -N) \le H_0(z^*, s) + \frac{\sigma^2\gamma^2 r^2 s^2}{2N} \tag{OA.2-11}$$

yields, after $N \to \infty$,

$$(r+\lambda)\bar{v}(z^*) \le H_0(z^*, s), \qquad s \in (p_r, p_\ell). \tag{OA.2-12}$$

Right-side bound. On $(z^*, 1)$ the HJB holds classically (Lemma OA.16) and $\phi = 0$ is feasible in the supremum, so $(r+\lambda)\bar{v}(z) \ge H_0(z, \bar{v}'(z))$. Letting $z \downarrow z^*$, with $\bar{v}$ and H_0 continuous and $\bar{v}'(z) \to p_r$,

$$(r+\lambda)\bar{v}(z^*) \ge H_0(z^*, p_r). \tag{OA.2-13}$$

Left-side bound. On $(0, z^*)$, the derivative $\bar{v}'$ exists a.e. and is nonincreasing, hence is itself differentiable a.e. At any z where $\bar{v}'(z)$ and $(\bar{v}')'(z) \equiv X$ both exist, integrating the relation $\bar{v}'(z+t) - \bar{v}'(z) = Xt + o(t)$ gives the second-order expansion $\bar{v}(z+h) = \bar{v}(z) + \bar{v}'(z)h + \frac{1}{2}Xh^2 + o(h^2)$, so for every $\epsilon > 0$ the function $y \mapsto \bar{v}(z) + \bar{v}'(z)(y-z) + \frac{1}{2}(X-\epsilon)(y-z)^2$ is a C^2 lower test at z. The per-control supersolution inequality (OA.1-57) with $\phi = 0$, under which the curvature term drops out, then gives $(r+\lambda)\bar{v}(z) \geq f(\beta, z) + k(z-\beta)\bar{v}'(z)$ for every $\beta \in [0,1]$, that is,

$$(r+\lambda)\bar{v}(z) \geq H_0(z, \bar{v}'(z)) \qquad \text{for a.e. } z \in (0, z^*). \tag{OA.2-14}$$

Letting $z \uparrow z^*$ along such points, with $\bar{v}'(z) \to p_\ell$,

$$(r+\lambda)\bar{v}(z^*) \geq H_0(z^*, p_\ell). \tag{OA.2-15}$$

Squeeze argument. Suppose $p_r < p_\ell$. Since $H_0(z^*, \cdot)$ is convex by (OA.1-78), its maximum over $[p_r, p_\ell]$ is attained at an endpoint, so (OA.2-13) and (OA.2-15) give, for every $s \in (p_r, p_\ell)$,

$$H_0(z^*, s) \leq \max\{H_0(z^*, p_r),\ H_0(z^*, p_\ell)\} \leq (r+\lambda)\bar{v}(z^*) \leq H_0(z^*, s), \tag{OA.2-16}$$

the last inequality being (OA.2-12). Hence $H_0(z^*, \cdot)$ is constant on $[p_r, p_\ell]$. This is impossible: by Danskin's theorem, $\partial H_0/\partial p = k(z^* - \beta_0(p; z^*))$, where the unique maximizer $\beta_0(p; z^*) \in [0,1]$ equals 1 for $p \leq -\sigma^2\gamma r/k$, equals 0 for $p \geq 1/(k\theta)$, and equals $(1/\theta - kp)/(1/\theta + \sigma^2\gamma r)$, strictly decreasing, in between. Since $z^* \in (0,1)$, the derivative is $k(z^*-1) < 0$ on the first branch, $kz^* > 0$ on the last, and strictly increasing on the middle branch; it therefore vanishes at no more than one point, and $H_0(z^*, \cdot)$ is constant on no nondegenerate interval. The contradiction proves $p_r = p_\ell$; since $p_r \leq 0 \leq p_\ell$, both equal zero, and $\bar{v}$ is differentiable at z^* with $\bar{v}'(z^*) = 0$. ■

The preceding argument does not rely on the degenerate constant upper test at $(p, X) = (0, 0)$, which yields only the automatic subsolution inequality for the upper semicontinuous envelope of the Hamiltonian. The following lemma instead uses the curvature identity to establish the endpoint Hamiltonian equality.

Lemma OA.19 (Vanishing endpoint Hamiltonian gap) *The endpoint Hamiltonian gap vanishes:*

$$\Delta_* \equiv \lim_{z \downarrow z^*} \Delta(z) = (r+\lambda)\bar{v}(z^*) - H_0(z^*, 0) = 0. \tag{OA.2-17}$$

Proof. By Lemma OA.18, $p(z) \equiv \bar{v}'(z) \to 0$ as $z \downarrow z^*$. Hence the displayed limit defining Δ_* exists by continuity of $\bar{v}$ and H_0. The right-side lower bound (OA.2-13), with $p_r = 0$, gives $\Delta_* \geq 0$.

Suppose, toward a contradiction, that $\Delta_* > 0$. On $(z^*, 1)$, the exact gap identity gives

$$\Delta(z) = G(z, p(z), q(z)), \qquad q(z) = -\frac{p(z)^2}{\bar{v}''(z)}. \tag{OA.2-18}$$

Since G extends continuously to $p = 0$, is strictly increasing in q, satisfies $G(z^*, 0, 0) = 0$, and diverges as $q \to \infty$, it follows that

$$q(z) \to q_* > 0, \qquad G(z^*, 0, q_*) = \Delta_*. \tag{OA.2-19}$$

Put $P(z) \equiv -p(z) > 0$. Since $\bar{v} \in C^2(z^*, 1)$ and $\bar{v}'' < 0$,

$$P'(z) = -\bar{v}''(z) = \frac{P(z)^2}{q(z)}. \tag{OA.2-20}$$

Fix $z_0 \in (z^*, 1)$ close enough to z^* that $q(z) \geq q_*/2$ on $(z^*, z_0]$. Then for $z \in (z^*, z_0]$,

$$\frac{1}{P(z)} = \frac{1}{P(z_0)} + \int_z^{z_0} \frac{1}{q(s)}\, ds \leq \frac{1}{P(z_0)} + \frac{2(z_0 - z^*)}{q_*} < \infty. \tag{OA.2-21}$$

But smooth fit gives $P(z) \to 0$ as $z \downarrow z^*$, so $1/P(z) \to \infty$, a contradiction. Therefore $\Delta_* = 0$. ■

Lemma OA.20 (Stopped benchmark) *Let*

$$\beta^* = \frac{1}{1 + \gamma\sigma^2 r\theta}. \tag{OA.2-22}$$

The stopped deterministic benchmark satisfies

$$v'_{stop}(\beta^*) = \frac{\gamma\sigma^2 r}{(r+\lambda)(1+\gamma\sigma^2 r\theta)} > 0, \tag{OA.2-23}$$

and its unique maximizer is

$$z^*_{stop} = \frac{2}{2 + \gamma\sigma^2 r\theta} = \frac{2\beta^*}{1+\beta^*} > \beta^*. \tag{OA.2-24}$$

Proof. The stopped benchmark is

$$v_{stop}(z) = \frac{1}{r+\lambda}\left[\frac{2z}{\theta} - \left(\frac{1}{\theta} + \frac{\sigma^2\gamma r}{2}\right) z^2\right], \tag{OA.2-25}$$

and therefore

$$v'_{stop}(z) = \frac{1}{r+\lambda}\left[\frac{2}{\theta} - \left(\frac{2}{\theta} + \sigma^2\gamma r\right) z\right]. \tag{OA.2-26}$$

Substitution of $\beta^* = 1/(1+\gamma\sigma^2 r\theta)$ gives (OA.2-23). The derivative vanishes at $z^*_{stop} = 2/(2+\gamma\sigma^2 r\theta) = 2\beta^*/(1+\beta^*)$, which is strictly larger than β^* because $\beta^* < 1$. ■

Corollary OA.21 (Lower-endpoint policy and initial back-loading) *As $z \downarrow z^*$,*

$$q(z) \to 0, \qquad \beta(z) \to \beta^* = \frac{1}{1+\theta\sigma^2\gamma r}, \tag{OA.2-27}$$

and the cap is slack near z^*. Moreover, the optimal initial state is strictly back-loaded: $z^* > \beta^*$.

Proof. Write $k \equiv r + \rho + \lambda$ and $s' \equiv \sigma^2\gamma^2 r^2/2$. On $(z^*, 1)$ the classical HJB and the exact gap identity (OA.1-86) give $\Delta(z) = G(z, p(z), q(z))$ with $q = -p^2/\bar{v}'' > 0$. From the envelope derivative $\partial G/\partial q = s'\beta_q(p)^2$ in the proof of Lemma OA.14, with $\beta_0(p) \leq \beta_q(p) \leq 1$,

$$s'\beta_0(p)^2\, q \;\leq\; G(z, p, q) \;\leq\; s' q. \tag{OA.2-28}$$

Since $p(z) \to 0$ (Lemma OA.18), $\beta_0(p(z)) \to \beta_0(0) = \beta^* > 0$, so for z near z^* the left inequality gives $q(z) \leq 2\Delta(z)/(s'\beta^{*2}) \to 2\Delta_*/(s'\beta^{*2}) = 0$. Consequently $A(z) = 1/\theta - kp(z) \to 1/\theta$ and $\Gamma(z) = 1/\theta + \sigma^2\gamma r - 2s'q(z) \to 1/\theta + \sigma^2\gamma r$, so the relaxed policy formula (OA.1-77) gives $\beta(z) \to \beta^* < 1$, with the cap slack near z^*.

For the strict ordering, Lemma OA.19 gives $(r+\lambda)\bar{v}(z^*) = H_0(z^*, 0) = f(\beta^*, z^*)$. Suppose $z^* \leq \beta^*$. Since $f_z = (1-z)/\theta > 0$ on $[0, 1)$, we would have $f(\beta^*, z^*) \leq f(\beta^*, \beta^*) = (r+\lambda)v_{stop}(\beta^*)$. But v_{stop} is strictly concave with unique maximizer $z^*_{stop} = 2\beta^*/(1+\beta^*) > \beta^*$ by (OA.2-24), so, using $\bar{v}(z^*) = \max \bar{v} \geq \bar{v}(z^*_{stop}) \geq v_{stop}(z^*_{stop})$ (Lemma OA.1),

$$(r+\lambda)\bar{v}(z^*) \geq (r+\lambda)v_{stop}(z^*_{stop}) > (r+\lambda)v_{stop}(\beta^*) \geq f(\beta^*, z^*) = (r+\lambda)\bar{v}(z^*),$$

a contradiction. Hence $z^* > \beta^*$. ■

Lemma OA.22 (Vanishing diffusion at the lower endpoint) *$\phi(z) \to 0$ as $z \downarrow z^*$. Consequently the closed-loop coefficients extend continuously to the lower endpoint with*

$$a(z^*) = 0, \qquad b(z^*) = k(z^* - \beta^*) > 0. \tag{OA.2-29}$$

Proof. Throughout, $z \in (z^*, z^* + \eta)$ for a small $\eta > 0$, $h \equiv z - z^*$, $P(z) \equiv -\bar{v}'(z) > 0$, and $P' = -\bar{v}'' = |\bar{v}''| > 0$; recall $\Delta > 0$ on $(z^*, 1)$ (Lemma OA.15), $P(z) \to 0$ (Lemma OA.18), and $\Delta(z) \to 0$ as $z \downarrow z^*$ (Lemma OA.19). Write $s' \equiv \sigma^2\gamma^2 r^2/2$ and define

$$W(z) \equiv \frac{\Delta(z)}{P(z)^2} > 0, \qquad c_0 \equiv \frac{1 - z^*}{\theta} > 0, \qquad y^* \equiv z^* - \beta^* > 0. \tag{OA.2-30}$$

Curvature–gap comparison. Evaluating (OA.2-28) at $q = P^2/P'$ and using $\beta_0(p(z))^2 \geq \beta^{*2}/2$ for η small: from $\Delta \leq s'P^2/P'$ we get $P' \leq s'P^2/\Delta$, and from $\Delta \geq s'\beta_0^2 P^2/P'$ we get $P' \geq (s'\beta^{*2}/2)P^2/\Delta$. That is,

$$\frac{c_1}{W(z)} \leq P'(z) \leq \frac{c_2}{W(z)}, \qquad c_1 \equiv \frac{s'\beta^{*2}}{2}, \quad c_2 \equiv s'. \tag{OA.2-31}$$

Derivative of the gap. On $(z^*, 1)$, $\bar{v} \in C^2$ and H_0 is continuously differentiable with $\partial_z H_0 = (1-z)/\theta + kp$ and $\partial_p H_0 = k(z - \beta_0(p; z))$ (Danskin's theorem), so, using $(r+\lambda) - k = -\rho$,

$$\Delta'(z) = \rho P(z) - \frac{1-z}{\theta} + k\, y(z)\, P'(z), \qquad y(z) \equiv z - \beta_0(p(z); z). \tag{OA.2-32}$$

By Corollary OA.21, $y(z) \to y^* > 0$. Shrink η so that on $(z^*, z^* + \eta)$: $y(z) \leq 2y^*$, $(1-z)/\theta \geq c_0/2$, and $\rho P(z) \leq c_0/8$.

Linear lower bound for the slope. Integrating (OA.2-32) from z^* (where Δ vanishes by Lemma OA.19) and using $\int_{z^*}^{z} yP'\,d\zeta \le 2y^*P(z)$ and $\int_{z^*}^{z} \rho P\,d\zeta \le (c_0/8)h$,

$$0 < \Delta(z) \le \frac{c_0}{8}h - \frac{c_0}{2}h + 2ky^*P(z), \quad \text{hence, } P(z) \ge c_P h,\ c_P \equiv \frac{3c_0}{16ky^*}. \tag{OA.2-33}$$

Boundedness of W. Let $\bar{W} \equiv \max\{8ky^*c_2/c_0, 1\}$. Wherever $W(z) \ge \bar{W}$ and $z \in (z^*, z^*+\eta)$, (OA.2-31) gives $kyP' \le 2ky^*c_2/\bar{W} \le c_0/4$, so

$$W' = \frac{1}{P^2}\left[\rho P - \frac{1-z}{\theta} + kyP'\right] - \frac{2WP'}{P} \tag{OA.2-34}$$

$$\le \frac{1}{P^2}\left[\frac{c_0}{8} - \frac{c_0}{2} + \frac{c_0}{4}\right] = -\frac{c_0}{8P^2} < 0. \tag{OA.2-35}$$

Suppose, toward a contradiction, that $\limsup_{z\downarrow z^*} W(z) > 2\bar{W}$. Pick $z_1 \in (z^*, z^*+\eta)$ with $W(z_1) > 2\bar{W}$ and $h_1 \equiv z_1 - z^*$ as small as desired, and let $z_2 \equiv \sup\big(\{z \in (z^*, z_1) : W(z) \le 2\bar{W}\} \cup \{z^*\}\big)$. If $z_2 > z^*$, then $W(z_2) \le 2\bar{W}$ by continuity, while on (z_2, z_1) we have $W > 2\bar{W} \ge \bar{W}$, so W is strictly decreasing there by (OA.2-35) and $W(z_2) \ge W(z_1) > 2\bar{W}$ – a contradiction. Hence $z_2 = z^*$ and $W > 2\bar{W}$ on all of (z^*, z_1). On that interval, (OA.2-31) gives $P' \le c_2/(2\bar{W}) \equiv \Lambda$, hence $P(z) \le \Lambda h$, and then (OA.2-35) gives $W'(z) \le -c_0/(8\Lambda^2 h^2)$. Integrating from z to z_1,

$$W(z) \ge \frac{c_0}{8\Lambda^2}\left(\frac{1}{h} - \frac{1}{h_1}\right) \ge \frac{c_0}{16\Lambda^2}\cdot\frac{1}{h}, \qquad h \le h_1/2. \tag{OA.2-36}$$

By (OA.2-31) again, $P'(z) \le (16c_2\Lambda^2/c_0)\,h$, and integrating, $P(z) \le (8c_2\Lambda^2/c_0)\,h^2$ for $h \le h_1/2$. Choosing h_1 small enough, this contradicts (OA.2-33) at $h = h_1/2$. Therefore $\limsup_{z\downarrow z^*} W(z) \le 2\bar{W}$; in particular $W \le 3\bar{W}$ on $(z^*, z^*+\eta')$ for some $\eta' > 0$.

Conclusion. The curvature–gap comparison and the bound on W imply $|\bar{v}''(z)| = P'(z) \ge c_1/(3\bar{W}) > 0$ near z^*, so

$$|\phi(z)| = \frac{\gamma r\,\beta(z)\,P(z)}{k\,|\bar{v}''(z)|} \le \frac{3\gamma r\bar{W}}{k\,c_1}\,P(z) \longrightarrow 0. \tag{OA.2-37}$$

Since also $\beta(z) \to \beta^*$ (Corollary OA.21), the coefficients $a(z) = k\sigma\phi(z)$ and $b(z) = k(z - \beta(z) + \sigma^2\gamma r\beta(z)\phi(z))$ extend continuously with $a(z^*) = 0$ and $b(z^*) = k(z^* - \beta^*) > 0$. ■

OA.2.3 Upper-Endpoint Behavior

Lemma OA.23 (Upper-end marginal cost) *The marginal value of relaxing the stock of promised future incentives diverges at the upper endpoint:*

$$\lim_{z\uparrow 1} \bar{v}'(z) = -\infty. \tag{OA.2-38}$$

Proof. Since $\bar{v}$ is concave and C^2 on $(z^*, 1)$, $p(z) \equiv \bar{v}'(z)$ is nonincreasing. Lemma OA.9

gives $p(z) < 0$ on $(z^*, 1)$, so the one-sided limit

$$p_1 \equiv \lim_{z \uparrow 1} p(z) \tag{OA.2-39}$$

exists in $[-\infty, 0)$. Suppose, toward a contradiction, that $p_1 > -\infty$. Then $p_1 < 0$. Put $d \equiv r + \lambda$, $k \equiv r + \rho + \lambda$, and $c \equiv -p_1 > 0$. By Lemma OA.3,

$$d\bar{v}(1) = f(1, 1). \tag{OA.2-40}$$

For every $z \in (z^*, 1)$, the relaxed HJB evaluated at the feasible control $(\beta, \phi) = (1, 0)$ gives

$$d\bar{v}(z) \geq f(1, z) + k(z - 1)p(z). \tag{OA.2-41}$$

Subtracting the endpoint identity and writing $x = 1 - z$,

$$d\big(\bar{v}(z) - \bar{v}(1)\big) \geq f(1, z) - f(1, 1) + k(z - 1)p(z). \tag{OA.2-42}$$

Moreover,

$$f(1, z) - f(1, 1) = -\frac{x^2}{2\theta}, \tag{OA.2-43}$$

and finite convergence of $p(z)$ implies

$$\frac{\bar{v}(z) - \bar{v}(1)}{x} = \frac{1}{x}\int_z^1 -p(s)\, ds \to c, \qquad \frac{(z - 1)p(z)}{x} = -p(z) \to c. \tag{OA.2-44}$$

Dividing (OA.2-42) by x and sending $z \uparrow 1$ yields $dc \geq kc$, contradicting $c > 0$ and $k - d = \rho > 0$. Therefore $p_1 = -\infty$. ■

Lemma OA.24 (Upper cap region) *There exists $z^{**} < 1$ such that*

$$\beta(z) = 1, \qquad z \in (z^{**}, 1). \tag{OA.2-45}$$

Proof. Let

$$q(z) \equiv -\frac{\bar{v}'(z)^2}{\bar{v}''(z)} > 0, \tag{OA.2-46}$$

and define

$$A(z) \equiv \frac{1}{\theta} - k\bar{v}'(z), \qquad \Gamma(z) \equiv \frac{1}{\theta} + \sigma^2\gamma r - \sigma^2\gamma^2 r^2 q(z). \tag{OA.2-47}$$

The difference between the linear and quadratic coefficients in the reduced β-problem is

$$A(z) - \Gamma(z) = -k\bar{v}'(z) - \sigma^2\gamma r + \sigma^2\gamma^2 r^2 q(z). \tag{OA.2-48}$$

By Lemma OA.23, $-k\bar{v}'(z) \to \infty$ as $z \uparrow 1$, while $q(z) > 0$. Hence $A(z) - \Gamma(z) > 0$ for all z sufficiently close to 1. If $\Gamma(z) \leq 0$, the reduced Hamiltonian is increasing in β on $[0, 1]$. If

$\Gamma(z) > 0$, then $A(z)/\Gamma(z) > 1$. In either case the unique maximizer in (OA.1-77) is $\beta(z) = 1$. The result follows. ■

Lemma OA.25 (Upper-end policy limits) *Set $\beta(1) = 1$ and $\phi(1) = 0$. Then*

$$\lim_{z\uparrow 1} q(z) = 0, \qquad \lim_{z\uparrow 1} \phi(z) = 0, \tag{OA.2-49}$$

where $q(z) \equiv -\bar{v}'(z)^2/\bar{v}''(z)$. Consequently,

$$\lim_{z\uparrow 1} b(z) = 0, \qquad \lim_{z\uparrow 1} a(z) = 0. \tag{OA.2-50}$$

Thus the feedback extends continuously to the absorbing upper endpoint.

Proof. By Lemma OA.24, there is $\bar{z} < 1$ such that $\beta(z) = 1$ on $(\bar{z}, 1)$. Write $d \equiv r + \lambda$, $s' \equiv \sigma^2\gamma^2 r^2/2$, $p(z) \equiv \bar{v}'(z)$, and $x \equiv 1 - z$. On this cap-binding region, the classical relaxed HJB and the optimized value of the ϕ-quadratic give

$$d\bar{v}(z) = f(1, z) + k(z-1)p(z) + s'q(z). \tag{OA.2-51}$$

Subtracting $d\bar{v}(1) = f(1, 1)$ and using $f(1, z) - f(1, 1) = -(1-z)^2/(2\theta)$ gives

$$s'q(z) = d\big(\bar{v}(z) - \bar{v}(1)\big) - kx\big(-p(z)\big) + \frac{x^2}{2\theta}. \tag{OA.2-52}$$

Since $q(z) \geq 0$,

$$0 \leq s'q(z) \leq d\big(\bar{v}(z) - \bar{v}(1)\big) + \frac{x^2}{2\theta}. \tag{OA.2-53}$$

Endpoint continuity from Lemma OA.6 therefore implies $q(z) \to 0$. On the same region $\beta = 1$, and the first-order condition for ϕ can be written as

$$\phi(z) = -\frac{\gamma r}{k}\frac{q(z)}{-p(z)}. \tag{OA.2-54}$$

Lemma OA.23 gives $-p(z) \to \infty$, so $\phi(z) \to 0$. Finally,

$$b(z) = k\left(z - 1 + \sigma^2\gamma r\phi(z)\right) \to 0, \qquad a(z) = k\sigma\phi(z) \to 0. \tag{OA.2-55}$$

The definitions $\beta(1) = 1$ and $\phi(1) = 0$ give the asserted absorbing extension. ■

OA.2.4 Exact Attainment

Theorem OA.26 (Exact attainment) *Extend the interior feedback controls of Lemma OA.16 to $[z^*, 1]$ by*

$$(\beta, \phi)(z^*) = (\beta^*, 0), \qquad (\beta, \phi)(1) = (1, 0), \tag{OA.2-56}$$

and let $b(z) \equiv k(z - \beta(z) + \sigma^2\gamma r\beta(z)\phi(z))$ *and* $a(z) \equiv k\sigma\phi(z)$*. Then for every* $z_0 \in [z^*, 1]$ *the closed-loop equation*

$$dz_t = b(z_t)\,dt + a(z_t)\,d\hat{Z}_t \tag{OA.2-57}$$

has a weak solution that remains in $[z^*, 1]$*, the induced feedback contract is restricted admissible, and it attains the relaxed value:*

$$J(z_0; \beta(\cdot), \phi(\cdot)) = \bar{v}(z_0), \qquad z_0 \in [z^*, 1]. \tag{OA.2-58}$$

Consequently,

$$v^-(z) = v^{\mathrm{IC}}(z) = \bar{v}(z) = \max_{(\tilde{\beta},\tilde{\phi}) \text{ relaxed admissible}} J(z; \tilde{\beta}, \tilde{\phi}), \qquad z \in [z^*, 1]. \tag{OA.2-59}$$

The suprema over the relaxed, restricted, and incentive-compatible classes are therefore attained. In particular, the optimal incentive-compatible contract is a stationary Markov contract in the single state z*, initialized at* z^**.*

Proof. *Continuity, boundedness, and admissibility of the extension.* On $(z^*, 1)$ the feedback controls are unique and continuous with $\beta \in (0, 1]$ and $\phi < 0$ (Lemma OA.16). At the lower endpoint, $\beta(z) \to \beta^*$ and $\phi(z) \to 0$ (Corollary OA.21 and Lemma OA.22); at the upper endpoint, $\beta(z) = 1$ near 1 and $\phi(z) \to 0$ (Lemmas OA.24 and OA.25). Hence the extended controls are continuous on $[z^*, 1]$, bounded, and satisfy $\phi \leq 0$ throughout, and a and b are continuous and bounded with

$$a(z^*) = 0, \quad b(z^*) = k(z^* - \beta^*) > 0, \qquad a(1) = 0, \quad b(1) = 0. \tag{OA.2-60}$$

Existence and invariance. Extend (a, b) to $\mathbb{R}$ by their endpoint values, $(a, b) \equiv (0, b(z^*))$ on $(-\infty, z^*]$ and $(a, b) \equiv (0, 0)$ on $[1, \infty)$. The extended coefficients are bounded and continuous, so the equation has a weak solution from every initial condition. Any solution started in $[z^*, 1]$ remains there, by a pathwise argument. Suppose $z_s > 1$ for some s; put $u_1 \equiv \sup\{u \leq s : z_u \leq 1\}$, so that $z_{u_1} = 1$ and $z_u > 1$ on $(u_1, s]$; there $a = b = 0$, so $z_s = z_{u_1} = 1$, a contradiction. Suppose instead $z_s < z^*$; put $u_1 \equiv \sup\{u \leq s : z_u \geq z^*\}$, so that $z_{u_1} = z^*$ and $z_u < z^*$ on $(u_1, s]$; there $a = 0$ and $b = b(z^*) > 0$, so $z_s = z^* + b(z^*)(s - u_1) > z^*$, again a contradiction. Hence $z_t \in [z^*, 1]$ for all t, and if z reaches 1 it is absorbed there, consistent with Lemma OA.3. The induced control pair $(\beta(z_t), \phi(z_t))$ is adapted and bounded with $\beta \in [0, 1]$ and $\phi \leq 0$, hence restricted admissible.

Verification. The hypotheses of Lemma OA.17 are now satisfied. The extension is restricted admissible and state constrained, with $(\beta, \phi)(1) = (1, 0)$ absorbing at the upper endpoint. Lower-end smooth fit, the Hamiltonian identity $(r + \lambda)\bar{v}(z^*) = f(\beta^*, z^*)$, the policy limit $(\beta, \phi)(z^*) = (\beta^*, 0)$, and the ordering $z^* > \beta^*$ follow from Lemmas OA.18 and OA.19, Corollary OA.21, and Lemma OA.22. The upper endpoint, if attained, is absorbing and satisfies $(r + \lambda)\bar{v}(1) = f(1, 1)$ by Lemma OA.3. Lemma OA.17 therefore yields the attainment identity and, together with the sandwich (OA.1-6), the equalities (OA.2-59). Initialization at the largest maximizer z^* attains $\max_z \bar{v}(z)$, so the optimal contract is the stationary Markov contract started at z^*. ∎

OA.3 Policies

This section derives the policy functions reported in Propositions 3 and 5. Theorem OA.26 implies that the relaxed, restricted, and incentive-compatible values coincide on $[z^*, 1]$; throughout this section, write

$$v \equiv v^- = v^{\mathrm{IC}} = \bar{v} \qquad \text{on } [z^*, 1]. \tag{OA.3-1}$$

The interior feedback formulas follow from the attained HJB. Lemma OA.27 establishes finite curvature at z^*, and Lemma OA.28 gives the right derivative of current PPS at the initial state. Propositions OA.31 and OA.33 provide sufficient conditions, expressed entirely in primitive parameters, under which the PPS cap binds above a unique threshold z^{**} and the contract switches from back-loading to front-loading at a unique threshold $\widehat{z}$.

OA.3.1 Interior and Endpoint Policies

On $(z^*, 1)$, Lemma OA.16 gives $v \in C^2$, $v' < 0$, and $v'' < 0$. Hence the attained HJB holds classically:

$$(r+\lambda)v(z) = \max_{\beta\in[0,1],\,\phi\in\mathbb{R}} \left\{ f(\beta, z) + k\left(z - \beta + \sigma^2\gamma r\beta\phi\right) v'(z) + \frac{1}{2}k^2\sigma^2\phi^2 v''(z) \right\},$$

where $k = r + \rho + \lambda$. For a fixed β, the objective is a strictly concave quadratic in ϕ, so the first-order condition gives

$$\phi(z) = -\frac{\gamma r\beta(z)v'(z)}{kv''(z)}. \tag{OA.3-2}$$

Since $v'(z) < 0$ and $v''(z) < 0$, any policy with $\beta(z) > 0$ has $\phi(z) < 0$.

Substituting (OA.3-2) into the Hamiltonian reduces the maximization over β to

$$\max_{\beta\in[0,1]} \left\{ A(z)\beta - \frac{1}{2}\Gamma(z)\beta^2 \right\}, \tag{OA.3-3}$$

up to terms independent of β, where

$$A(z) \equiv \frac{1}{\theta} - kv'(z) > 0, \tag{OA.3-4}$$

$$\Gamma(z) \equiv \frac{1}{\theta} + \sigma^2\gamma r + \sigma^2\gamma^2 r^2 \frac{v'(z)^2}{v''(z)}. \tag{OA.3-5}$$

Therefore the interior policy is

$$\beta(z) = \begin{cases} \min\left\{1, \dfrac{A(z)}{\Gamma(z)}\right\}, & \Gamma(z) > 0, \\[2ex] 1, & \Gamma(z) \le 0, \end{cases} \tag{OA.3-6}$$

and $\phi(z)$ is then given by (OA.3-2). The maximizer is unique: if $\Gamma(z) > 0$ the reduced objective is strictly concave, while if $\Gamma(z) \le 0$ its derivative $A(z) - \Gamma(z)\beta$ is strictly positive on $[0,1]$, so the cap $\beta = 1$ binds. In particular,

$$\beta(z) > 0, \qquad \phi(z) < 0, \qquad z \in (z^*, 1). \tag{OA.3-7}$$

Corollary OA.21 and Lemma OA.22 give the lower-endpoint policy:

$$\lim_{z \downarrow z^*} \beta(z) = \beta^*, \qquad \lim_{z \downarrow z^*} \phi(z) = 0, \qquad z^* > \beta^*. \tag{OA.3-8}$$

Accordingly, set $\beta(z^*) = \beta^*$ and $\phi(z^*) = 0$. At the upper endpoint, Lemmas OA.24 and OA.25 give $\beta(z) = 1$ for all z sufficiently close to 1 and $\phi(z) \to 0$ as $z \uparrow 1$. Thus $\beta(1) = 1$ and $\phi(1) = 0$.

OA.3.2 Lower-Endpoint Asymptotics

Lemma OA.27 (Finite lower-end curvature) *The lower-end curvature has the finite one-sided limit*

$$\lim_{z \downarrow z^*} v''(z) = X_* \in (-\infty, 0), \tag{OA.3-9}$$

with

$$X_* = -\frac{1 - z^*}{\theta k (z^* - \beta^*)}. \tag{OA.3-10}$$

Proof. Write $h = z - z^*$, $P(z) \equiv -v'(z) > 0$, $R(z) \equiv P'(z) = -v''(z) > 0$, and

$$\Delta(z) \equiv (r + \lambda) v(z) - H_0(z, v'(z)), \qquad W(z) \equiv \frac{\Delta(z)}{P(z)^2}. \tag{OA.3-11}$$

Let

$$c_0 \equiv \frac{1 - z^*}{\theta}, \qquad y^* \equiv z^* - \beta^*, \qquad L \equiv \frac{c_0}{k y^*}, \qquad g_* \equiv \frac{\sigma^2 \gamma^2 r^2}{2} \beta^{*2}. \tag{OA.3-12}$$

Corollary OA.21 gives $y^* > 0$, $P(z) \to 0$, $q(z) \to 0$, and $\beta_0(v'(z); z) \to \beta^*$ as $z \downarrow z^*$. The boundedness argument in the proof of Lemma OA.22 gives an upper bound $W(z) \le M$ near z^*.

First, $P(z) = O(h)$. Since $\Delta = W P^2 \le M P^2$, and since the gap derivative identity (OA.2-32) gives

$$\Delta(z) = \int_{z^*}^{z} \left[\rho P(s) - \frac{1 - s}{\theta} + k\, y(s) R(s) \right] ds, \qquad y(s) \equiv s - \beta_0(v'(s); s), \tag{OA.3-13}$$

we have, for z close to z^*,

$$\Delta(z) \ge -Ch + \frac{k y^*}{2} P(z) \tag{OA.3-14}$$

for some finite constant C. If $P(z)/h$ were unbounded along a sequence approaching z^*, then the last display and $\Delta \le MP^2$ would imply $(ky^*/4)P(z) \le MP(z)^2$ along that sequence, impossible because $P(z) \to 0$. Thus $P(z) \le C_P h$ near z^*.

Next, define

$$g(z) \equiv \frac{G(z, v'(z), q(z))}{q(z)}, \qquad q(z) = \frac{P(z)^2}{R(z)}. \tag{OA.3-15}$$

By the envelope formula for G, $g(z) \to g_*$. Since

$$\Delta(z) = G(z, v'(z), q(z)) = g(z)\frac{P(z)^2}{R(z)}, \tag{OA.3-16}$$

we have the exact relation

$$R(z)W(z) = g(z). \tag{OA.3-17}$$

Let $W_* \equiv g_*/L$. We show $W(z) \to W_*$. Fix $\epsilon > 0$. For z sufficiently close to z^*, (OA.3-17) and the convergence of g, y, P, and $(1-z)/\theta$ imply:

$$W(z) \ge W_* + \epsilon \quad \Longrightarrow \quad \Delta'(z) \le -\delta, \tag{OA.3-18}$$

and

$$W(z) \le W_* - \epsilon \quad \Longrightarrow \quad \Delta'(z) \ge \delta, \tag{OA.3-19}$$

for some $\delta > 0$. Indeed, these implications are just the gap derivative identity

$$\Delta'(z) = \rho P(z) - \frac{1-z}{\theta} + k\,y(z)R(z) \tag{OA.3-20}$$

and the fact that $R = g/W$ is below $L = c_0/(ky^*)$ by a fixed amount in the first case and above it by a fixed amount in the second.

Since $\Delta = WP^2$,

$$W'(z) = \frac{\Delta'(z)}{P(z)^2} - \frac{2W(z)R(z)}{P(z)}. \tag{OA.3-21}$$

If $W \ge W_* + \epsilon$, then (OA.3-21) gives $W' \le -\delta/P^2 < 0$. If this inequality held at a point z_0 arbitrarily close to z^*, then moving left from z_0 the function W cannot cross below $W_* + \epsilon$; while it remains in this region,

$$W'(z) \le -\frac{\delta}{P(z)^2} \le -\frac{\delta}{C_P^2(z-z^*)^2}. \tag{OA.3-22}$$

Integrating backward to z^* would force $W(z) \to +\infty$, contradicting the upper bound $W \le M$. Hence $\limsup_{z \downarrow z^*} W(z) \le W_*$.

Similarly, if $W \le W_* - \epsilon$ near z^*, then $WR = g$ is bounded and (OA.3-21) gives

$$W'(z) \ge \frac{\delta}{P(z)^2} - \frac{C}{P(z)} \ge \frac{\delta}{2P(z)^2} > 0 \tag{OA.3-23}$$

for z close enough to z^*. If this inequality held at a point z_0 arbitrarily close to z^*, then moving left from z_0 the function W cannot cross above $W_* - \epsilon$; while it remains in this region,

$$W'(z) \ge \frac{\delta}{2C_P^2(z-z^*)^2}. \tag{OA.3-24}$$

Integrating backward to z^* would force W to become negative, impossible because $\Delta > 0$. Thus $\liminf_{z\downarrow z^*} W(z) \ge W_*$. Therefore $W(z) \to W_*$.

Finally, (OA.3-17) gives

$$R(z) = \frac{g(z)}{W(z)} \to \frac{g_*}{W_*} = L. \tag{OA.3-25}$$

Since $R = -v''$, (OA.3-10) follows. ■

Lemma OA.28 (Lower-end slope of current PPS) *Current PPS has the lower-end expansion*

$$\beta(z) = \beta^* + \beta^* \frac{1-z^*}{z^*-\beta^*}(z-z^*) + o(z-z^*), \qquad z \downarrow z^*. \tag{OA.3-26}$$

Equivalently,

$$\beta'_+(z^*) = -\frac{kX_*}{1/\theta + \sigma^2\gamma r} = \beta^* \frac{1-z^*}{z^*-\beta^*} > 0. \tag{OA.3-27}$$

Proof. Let $h = z - z^* > 0$. Lemma OA.27 and smooth fit give

$$p(z) \equiv v'(z) = X_* h + o(h), \qquad X_* = -\frac{1-z^*}{\theta k(z^*-\beta^*)} < 0. \tag{OA.3-28}$$

Corollary OA.21 gives $\beta(z) \to \beta^* < 1$, so the cap is slack near z^* and $\beta = A/\Gamma$. Also

$$q(z) \equiv -\frac{v'(z)^2}{v''(z)} = O(h^2). \tag{OA.3-29}$$

Thus

$$A(z) = \frac{1}{\theta} - kX_*h + o(h), \qquad \Gamma(z) = \frac{1}{\theta} + \sigma^2\gamma r + O(h^2). \tag{OA.3-30}$$

Therefore

$$\beta'_+(z^*) = \frac{A'_+(z^*)\Gamma(z^*) - A(z^*)\Gamma'_+(z^*)}{\Gamma(z^*)^2} = \frac{-kX_*}{1/\theta + \sigma^2\gamma r}. \tag{OA.3-31}$$

Substituting the expression for X_* from Lemma OA.27 and using

$$\beta^* = \frac{1/\theta}{1/\theta + \sigma^2\gamma r} \tag{OA.3-32}$$

gives (OA.3-27). Since $z^* > \beta^*$ and $z^* < 1$, this derivative is strictly positive, and (OA.3-26) follows. ■

OA.3.3 Uniqueness of z^{**}

This subsection gives a sufficient condition, stated in terms of the model parameters alone, under which the region where the incentive cap binds is a single interval: there is a unique threshold $z^{**} \in (z^*, 1)$ with $\beta(z) < 1$ on (z^*, z^{**}) and $\beta(z) = 1$ on $[z^{**}, 1)$. Economically, the condition guarantees that the state space is split into exactly two ordered regions, so the maximal-incentive region is entered only after promised incentives have accumulated past a single threshold. In view of Theorem OA.26, we write v for the common value.

Throughout, $z \in (z^*, 1)$, and we use the interior objects established above: $p \equiv v' < 0$, $P \equiv -v' > 0$ with $P' = -v'' \equiv R > 0$, $q \equiv P^2/R > 0$, the gap $\Delta(z) \equiv (r+\lambda)v(z) - H_0(z, v'(z)) > 0$ (Lemma OA.15), the exact gap identity $\Delta = G(z, p, q)$ (Lemma OA.16 and (OA.1-86)), and the deterministic maximizer $\beta_0(p) = \min\{1, (1/\theta + kP)/C\}$, where

$$k \equiv r + \rho + \lambda, \qquad C \equiv \frac{1}{\theta} + \sigma^2\gamma r, \qquad S \equiv \sigma^2\gamma^2 r^2 = 2s'. \tag{OA.3-33}$$

Since $v'^2/v'' = -q$, the coefficients of the policy formula (OA.3-6) are $A(z) = 1/\theta + kP(z)$ and $\Gamma(z) = C - Sq(z)$, and $\beta = 1$ exactly when $A \geq \Gamma$ (either $\Gamma \leq 0$, or $\Gamma > 0$ and $A/\Gamma \geq 1$). Hence

$$\beta(z) = 1 \quad \iff \quad kP(z) + Sq(z) \geq \sigma^2\gamma r. \tag{OA.3-34}$$

We first record a priori bounds that control the endogenous objects by primitives. They are consequences of the endpoint identity $(r+\lambda)v(z^*) = H_0(z^*, 0) = f(\beta^*, z^*)$ established in Lemma OA.19: this is where the attainment analysis feeds the comparative structure of the policies.

Lemma OA.29 (A priori bounds) *For every $z \in (z^*, 1)$:*

(i) $\beta(z) > \beta^*$;

(ii) *with* $I(z) \equiv \int_{z^*}^{z} \frac{1-\zeta}{\theta}\, d\zeta$,

$$\Delta(z) \;\leq\; k(z-\beta^*)P(z) - I(z) \;\leq\; k(z-\beta^*)P(z), \tag{OA.3-35}$$

hence

$$P(z) > \frac{I(z)}{k(z-\beta^*)}; \tag{OA.3-36}$$

(iii) $\Delta(z) \geq s'\beta_0(p(z))^2\, q(z)$;

(iv) the optimal initial state satisfies $z^* \geq \underline{\zeta}$, where

$$\underline{\zeta} \equiv 1 - \frac{1-\beta^*}{\sqrt{1+\beta^*}} \;\in\; \left(\beta^*,\; z^*_{stop}\right). \tag{OA.3-37}$$

Proof. (i) Where the cap binds, $\beta = 1 > \beta^*$. Elsewhere $\beta = A/\Gamma$ with $A > 1/\theta$ and $0 < \Gamma < C$, so $\beta > (1/\theta)/C = \beta^*$.

(ii) Choosing $\beta = \beta^*$ in the maximum defining H_0 gives $H_0(z,p) \geq f(\beta^*, z) + k(z-\beta^*)p = f(\beta^*, z) - k(z-\beta^*)P$. Since z^* maximizes v and $(r+\lambda)v(z^*) = f(\beta^*, z^*)$ (Lemma OA.19),

$$\Delta(z) \leq (r+\lambda)v(z^*) - f(\beta^*, z) + k(z-\beta^*)P(z) = k(z-\beta^*)P(z) - \int_{z^*}^{z} \frac{1-\zeta}{\theta}\, d\zeta,$$

using $f_z(\beta^*, \zeta) = (1-\zeta)/\theta$. The slope bound follows from $\Delta > 0$.

(iii) In the proof of Lemma OA.14, $\partial G/\partial q = s'\beta_q(p)^2$ with $\beta_q(p)$ nondecreasing in q and equal to $\beta_0(p)$ at $q = 0$; integrating from 0 to q gives $G(z,p,q) \geq s'\beta_0(p)^2 q$, and $G(z,p,q) = \Delta(z)$.

(iv) Using $\sigma^2\gamma r\, \beta^* = (1-\beta^*)/\theta$, direct substitution gives the closed forms

$$f(\beta^*, \zeta) = \frac{\beta^* + 2\zeta - \zeta^2}{2\theta}, \tag{OA.3-38}$$

$$\max_{\zeta\in[0,1]} f(\zeta, \zeta) = f\left(z^*_{stop}, z^*_{stop}\right) = \frac{z^*_{stop}}{\theta}, \tag{OA.3-39}$$

$$z^*_{stop} = \frac{2\beta^*}{1+\beta^*}. \tag{OA.3-40}$$

By Lemma OA.19 and Lemma OA.1, $f(\beta^*, z^*) = (r+\lambda)v(z^*) \geq (r+\lambda)v_{stop}(z^*_{stop}) = f(z^*_{stop}, z^*_{stop})$. Since $f(\beta^*, \cdot)$ is strictly increasing on $[0,1)$, it follows that $z^* \geq \underline{\zeta}$, where $\underline{\zeta}$ solves $\beta^* + 2\zeta - \zeta^2 = 2z^*_{stop}$, that is, $(1-\zeta)^2 = 1 + \beta^* - 2z^*_{stop} = (1-\beta^*)^2/(1+\beta^*)$, which is (OA.3-37). Finally, $\underline{\zeta} > \beta^*$ because $\sqrt{1+\beta^*} > 1$, and $\underline{\zeta} < z^*_{stop}$ because $f(\beta^*, z^*_{stop}) > f(z^*_{stop}, z^*_{stop})$, the maximizer β^* of $f(\cdot, z^*_{stop})$ being unique and distinct from z^*_{stop}. ■

The first hypothesis of the theorem below is $\underline{\zeta} \geq 1/2$. By (OA.3-37) this is a restriction on β^*, hence on the single parameter combination $\sigma^2\gamma r\theta$:

$$\underline{\zeta} \geq \tfrac{1}{2} \quad\Longleftrightarrow\quad \beta^* \geq \frac{9-\sqrt{33}}{8} \quad\Longleftrightarrow\quad \sigma^2\gamma r\theta \leq \frac{3+\sqrt{33}}{6} \approx 1.457. \tag{OA.3-41}$$

Together with the maintained primitive restriction $\sigma^2\gamma r\theta \geq 1$ in Section 2.2, the admissible band is $1 \leq \sigma^2\gamma r\theta \leq (3+\sqrt{33})/6$; the baseline calibration has $\sigma^2\gamma r\theta = 1.2$.

The characterization (OA.3-34) compares $kP + Sq$ with $\sigma^2\gamma r$, but q involves v'', and existence of v''' has not been established. The next lemma recasts the comparison through the gap Δ, which *is* continuously differentiable, and computes the derivative of the resulting margin at a crossing.

Lemma OA.30 (Cap margin) *On $\mathcal{O} \equiv \{z \in (z^*, 1) : kP(z) < \sigma^2\gamma r\}$, define*

$$\tilde{q}(z) \equiv \frac{\sigma^2\gamma r - kP(z)}{S}, \qquad m(z) \equiv \Delta(z) - G\big(z, p(z), \tilde{q}(z)\big). \tag{OA.3-42}$$

Then: (i) $m \in C^1(\mathcal{O})$; (ii) for $z \in \mathcal{O}$ the cap binds at z if and only if $m(z) \geq 0$, while for $z \notin \mathcal{O}$ it binds automatically; (iii) at any $z_0 \in \mathcal{O}$ with $m(z_0) = 0$,

$$m'(z_0) = \rho P(z_0) - \frac{1 - z_0}{\theta} + \frac{k(2z_0 - 1)}{2} R(z_0), \quad R(z_0) = \frac{S\, P(z_0)^2}{\sigma^2 \gamma r - kP(z_0)}. \tag{OA.3-43}$$

Proof. (i) $v \in C^2(z^*, 1)$ makes p, P, Δ, and $\tilde{q}$ continuously differentiable, and G is C^1 by Danskin's theorem, its maximizer β_q being unique and continuous.

(ii) If $kP \geq \sigma^2 \gamma r$, then (OA.3-34) holds because $Sq > 0$. For $z \in \mathcal{O}$, (OA.3-34) reads $q \geq \tilde{q}$, and since $G(z, p, \cdot)$ is strictly increasing (Lemma OA.14), this is equivalent to $G(z, p, q) \geq G(z, p, \tilde{q})$, that is, to $\Delta(z) \geq G(z, p, \tilde{q})$.

(iii) At z_0, $m(z_0) = 0$ forces $q(z_0) = \tilde{q}(z_0)$ by strict monotonicity of $G(z, p, \cdot)$, hence $R(z_0) = P^2/\tilde{q}$ as displayed, and the cap binds exactly, with maximizer $\beta_{\tilde{q}} = 1$. Differentiate the two pieces of m. First, by the identity (OA.2-32), valid on all of $(z^*, 1)$,

$$\Delta'(z_0) = \rho P - \frac{1 - z_0}{\theta} + k\big(z_0 - \beta_0(p)\big) R. \tag{OA.3-44}$$

Second, by Danskin's theorem applied to (OA.1-84): $\partial_z G = 0$ (the z-derivative $f_z + kp$ of both maxima is independent of β), $\partial_p G = k\big(\beta_0(p) - \beta_{\tilde{q}}\big) = k\big(\beta_0(p) - 1\big)$, and $\partial_q G = s' \beta_{\tilde{q}}^2 = s'$; with $p' = -R$ and $\tilde{q}\,' = -kR/S$,

$$\begin{aligned} \frac{d}{dz} G\big(z, p(z), \tilde{q}(z)\big)\Big|_{z_0} &= k\big(\beta_0(p) - 1\big)(-R) + s'\Big(-\frac{kR}{S}\Big) && \text{(OA.3-45)} \\ &= k\big(1 - \beta_0(p)\big) R - \frac{kR}{2}. && \text{(OA.3-46)} \end{aligned}$$

Subtracting, the β_0-terms cancel and (OA.3-43) follows. ■

Proposition OA.31 (Uniqueness of z^{**}) *Define, for $z \in [\underline{\zeta}, 1)$ and $0 < x < \sigma^2 \gamma r$,*

$$\begin{aligned} U(z, x) &\equiv \frac{\rho x}{k} - \frac{1 - z}{\theta} + \frac{(2z - 1)\, S\, x^2}{2k\,(\sigma^2 \gamma r - x)}, && \text{(OA.3-47)} \\ \omega(z) &\equiv \min \Big\{ x \in (0, \sigma^2 \gamma r) : \Big(\tfrac{1}{\theta} + x\Big)^2 (\sigma^2 \gamma r - x) \leq 2C^2 (z - \beta^*)\, x \Big\}. && \text{(OA.3-48)} \end{aligned}$$

Suppose that

$$\text{(H1)} \quad \sigma^2 \gamma r \theta \leq \frac{3 + \sqrt{33}}{6}, \qquad \text{(H2)} \quad U\big(z, \omega(z)\big) > 0 \ \ \textit{for all } z \in [\underline{\zeta}, 1). \tag{OA.3-49}$$

Then the cap region is a single interval: there exists a unique $z^{**} \in (z^*, 1)$ such that $\beta(z) < 1$ for $z \in (z^*, z^{**})$ and $\beta(z) = 1$ for $z \in [z^{**}, 1)$.

Proof. *The floor at a crossing.* Let $z_0 \in \mathcal{O}$ with $m(z_0) = 0$, and write $x_0 \equiv kP(z_0) \in (0, \sigma^2 \gamma r)$. On $\mathcal{O}$ the deterministic maximizer is interior, $\beta_0(p(z_0)) = (1/\theta + x_0)/C < 1$, so combining Lemma OA.29(iii) at $q = \tilde{q}(z_0)$ with Lemma OA.29(ii),

$$\frac{(1/\theta + x_0)^2}{2C^2} \big(\sigma^2 \gamma r - x_0\big) = s' \beta_0\big(p(z_0)\big)^2\, \tilde{q}(z_0) \ \leq \ \Delta(z_0) \ \leq \ (z_0 - \beta^*)\, x_0. \tag{OA.3-50}$$

Thus x_0 belongs to the constraint set in (OA.3-48), so $x_0 \geq \omega(z_0)$. To see that the minimum in (OA.3-48) is well defined, subtract the right side of its inequality from the left and extend the resulting continuous function to $[0, \sigma^2\gamma r]$. Its value at zero is $\sigma^2\gamma r/\theta^2 > 0$, whereas its value at $\sigma^2\gamma r$ is strictly negative because $z_0 > z^* > \beta^*$. The feasible set is therefore nonempty, closed, bounded away from zero, and contains a left neighborhood of $\sigma^2\gamma r$. Its minimum consequently lies in $(0, \sigma^2\gamma r)$ and is strictly positive.

Sign of the margin derivative. By (H1) and (OA.3-41), $\underline{\zeta} \geq 1/2$, so Lemma OA.29(iv) gives $z_0 > z^* \geq \underline{\zeta} \geq 1/2$. Both $x \mapsto \rho x/k$ and $x \mapsto x^2/(\sigma^2\gamma r - x)$ are strictly increasing. Hence $U(z_0, \cdot)$ is strictly increasing on $(0, \sigma^2\gamma r)$, and (OA.3-43) reads $m'(z_0) = U(z_0, x_0)$. Therefore

$$m'(z_0) = U(z_0, x_0) \geq U\big(z_0, \omega(z_0)\big) > 0 \tag{OA.3-51}$$

by (H2), since $z_0 \in (z^*, 1) \subseteq [\underline{\zeta}, 1)$.

Single crossing. Suppose the cap binds at z_1 and is slack at some $z_2 \in (z_1, 1)$. Slackness gives $kP(z_2) < \sigma^2\gamma r$ and, by Lemma OA.30(ii), $m(z_2) < 0$. Since P is increasing, $[z_1, z_2] \subset \mathcal{O}$, and the cap binding at z_1 gives $m(z_1) \geq 0$. Let $z_3 \equiv \sup\{z \in [z_1, z_2] : m(z) \geq 0\}$. By continuity, $m(z_3) = 0$ and $z_3 < z_2$, with $m < 0$ on $(z_3, z_2]$; hence $m'(z_3) \leq 0$, contradicting the previous step. The cap region is therefore an upper set in $(z^*, 1)$. It is nonempty (Lemma OA.24) and its complement is nonempty ($\beta(z) \to \beta^* < 1$ as $z \downarrow z^*$, Corollary OA.21), so $z^{**} \equiv \inf\{z \in (z^*, 1) : \beta(z) = 1\}$ lies in $(z^*, 1)$, with $\beta = 1$ on $[z^{**}, 1)$ by the upper-set property and continuity of β, and $\beta < 1$ on (z^*, z^{**}). ■

Three remarks. First, both hypotheses are conditions on the primitives $(r, \rho, \lambda, \gamma, \sigma, \theta)$ alone: $\underline{\zeta}$, U, and ω are explicit, and checking (H2) reduces to evaluating the one-dimensional function $z \mapsto U(z, \omega(z))$ on $[\underline{\zeta}, 1)$. The construction eliminates every endogenous object: at a crossing, q and R are pinned by the crossing equation, the slope P is bounded below by the primitive floor ω, and the location of the continuation region is bounded below by the primitive $\underline{\zeta}$. Second, since $\beta_0 \geq \beta^*$, the constraint set in (OA.3-48) yields the closed-form but more conservative floor $\omega(z) \geq \sigma^2\gamma r\, \beta^{*2}/\big(\beta^{*2} + 2(z - \beta^*)\big)$; replacing ω by this expression gives a fully closed-form version of (H2), at the cost of a smaller admissible parameter region. Third, at the baseline calibration $(r, \rho, \lambda, \gamma, \sigma, \theta) = (0.05, 0.1, 0, 1, 4, 1.5)$, $\sigma^2\gamma r\theta = 1.2$ satisfies (H1) with $\underline{\zeta} \approx 0.548$. Numerical minimization of $z \mapsto U(z, \omega(z))$ on $[\underline{\zeta}, 1)$ gives a minimum of approximately 0.09 near $z \approx 0.8$, so the baseline satisfies (H2) to numerical precision and exhibits a unique z^{**}.

OA.3.4 Single-Crossing Threshold

We next give a primitive sufficient condition under which $\beta(z)$ crosses z exactly once. We construct a continuously differentiable margin whose sign agrees with that of $\beta(z) - z$ and evaluate its derivative only at a crossing.

For $z \in [\underline{\zeta}, 1]$, define

$$D(z) \equiv C(z - \beta^*), \tag{OA.3-52}$$

$$\eta(z) \equiv \frac{C}{2}\left[\sqrt{z^2 + 4\beta^*(z - \beta^*)} - z\right], \tag{OA.3-53}$$

$$\Psi(z, x) \equiv \frac{Sz^2x^2}{k\big(D(z) - x\big)} + \left(\frac{2\rho}{k} - 1\right)x - \frac{3 - 2z}{\theta}, \tag{OA.3-54}$$

and let

$$\mathcal{K} \equiv \left\{(z, x) : z \in [\underline{\zeta}, 1],\ \eta(z) \le x < D(z)\right\}. \tag{OA.3-55}$$

All of these objects depend only on the primitives.

Lemma OA.32 (The crossing margin) *Let $x(z) \equiv kP(z)$ and*

$$\mathcal{J} \equiv \{z \in (z^*, 1) : x(z) < D(z)\}. \tag{OA.3-56}$$

On $\mathcal{J}$, define

$$\widehat{q}(z) \equiv \frac{D(z) - x(z)}{Sz}, \qquad n(z) \equiv \Delta(z) - G\big(z, p(z), \widehat{q}(z)\big). \tag{OA.3-57}$$

Then $n \in C^1(\mathcal{J})$ and

$$n(z) \gtreqless 0 \iff \beta(z) \gtreqless z. \tag{OA.3-58}$$

For $z \notin \mathcal{J}$, $\beta(z) > z$. At any $z_0 \in \mathcal{J}$ such that $n(z_0) = 0$, writing $x_0 = x(z_0)$,

$$n'(z_0) = \frac{1}{2}\Psi(z_0, x_0), \qquad R(z_0) = \frac{Sz_0x_0^2}{k^2\big(D(z_0) - x_0\big)}. \tag{OA.3-59}$$

Moreover, every such crossing satisfies

$$x_0 \ge \eta(z_0). \tag{OA.3-60}$$

Proof. On $\mathcal{J}$, $\widehat{q} > 0$ is the unique value of q for which the unconstrained maximizer in (OA.1-90) equals z:

$$\frac{1/\theta + x}{C - S\widehat{q}} = z. \tag{OA.3-61}$$

Because the maximizer is increasing in q and $G(z, p, \cdot)$ is strictly increasing, (OA.3-58) follows from $\Delta = G(z, p, q)$. If $x \ge D(z)$, the deterministic maximizer is at least z, and the strict inequality $q > 0$ makes the full maximizer strictly larger than z; hence $\beta(z) > z$ outside $\mathcal{J}$. The regularity of n follows as in Lemma OA.30.

At a zero of n, strict monotonicity of G gives $q = \widehat{q}$ and hence $\beta = z$. Since $q = P^2/R$

and $x = kP$, this gives the expression for R in (OA.3-59). Also

$$\widehat{q}\,'(z_0) = \frac{1/\theta + x_0 - kR(z_0)z_0}{Sz_0^2}. \tag{OA.3-62}$$

At the crossing, Danskin's theorem gives

$$G_p = k(\beta_0 - z_0), \qquad G_q = \frac{S}{2}z_0^2, \qquad G_z = 0. \tag{OA.3-63}$$

Using $p' = -R$ and the gap derivative identity (OA.2-32), differentiation of n therefore yields

$$n'(z_0) = \rho P - \frac{1-z_0}{\theta} - \frac{1}{2}\left(\frac{1}{\theta} + x_0 - kRz_0\right) \tag{OA.3-64}$$

$$= \frac{1}{2}\left[kRz_0 + 2\rho P - x_0 - \frac{3-2z_0}{\theta}\right] = \frac{1}{2}\Psi(z_0, x_0), \tag{OA.3-65}$$

where the last equality uses the displayed expression for R.

It remains to establish the primitive floor. At a crossing, the deterministic maximizer is

$$\beta_0 = \beta^* + \frac{x_0}{C} < z_0, \tag{OA.3-66}$$

and direct evaluation of the two quadratic maxima defining G gives

$$\Delta(z_0) = G\big(z_0, p(z_0), \widehat{q}(z_0)\big) = \frac{1}{2}\left(\beta^* + \frac{x_0}{C}\right)\big(D(z_0) - x_0\big). \tag{OA.3-67}$$

The upper bound in Lemma OA.29(ii) gives $\Delta(z_0) \le (z_0 - \beta^*)x_0$. Consequently,

$$\left(\beta^* + \frac{x_0}{C}\right)\big(D(z_0) - x_0\big) \le 2(z_0 - \beta^*)x_0, \tag{OA.3-68}$$

or, equivalently,

$$x_0^2 + Cz_0x_0 - \beta^*C^2(z_0 - \beta^*) \ge 0. \tag{OA.3-69}$$

The positive root is precisely $\eta(z_0)$ in (OA.3-53), proving (OA.3-60). ■

Proposition OA.33 (Unique back-loading threshold) *Suppose (H1)–(H2) in (OA.3-49) and the additional primitive condition*

$$\text{(H3)} \qquad \inf_{(z,x)\in\mathcal{K}} \Psi(z, x) > 0. \tag{OA.3-70}$$

Then there exists a unique $\widehat{z} \in (z^*, z^{**})$ *such that*

$$\beta(z) < z, \qquad z \in (z^*, \widehat{z}), \tag{OA.3-71}$$

$$\beta(\widehat{z}) = \widehat{z}, \tag{OA.3-72}$$

$$\beta(z) > z, \qquad z \in (\widehat{z}, 1). \tag{OA.3-73}$$

Thus the optimal contract is back-loaded below $\widehat{z}$ and front-loaded above it.

Proof. Corollary OA.21 gives $\beta(z^*) = \beta^* < z^*$, whereas Lemma OA.24 gives $\beta(z) = 1 > z$ sufficiently close to one. Continuity therefore guarantees at least one crossing.

We first show that the set $\{z : \beta(z) \geq z\}$ is an upper set. Suppose instead that $\beta(z_1) \geq z_1$ and $\beta(z_2) < z_2$ for some $z_1 < z_2$. The latter inequality implies $z_2 \in \mathcal{J}$ and $n(z_2) < 0$. Consider the connected component of $\mathcal{J}$ containing z_2. If it also contains z_1, then $n(z_1) \geq 0$. Otherwise its left endpoint lies weakly above z_1 and strictly above z^*; as that endpoint is approached from the right, $x \to D(z)$, $\widehat{q} \to 0$, and $n \to \Delta > 0$. In either case, continuity produces a point z_3 in this component at which $n(z_3) = 0$, with $n < 0$ immediately to its right. Hence $n'(z_3) \leq 0$.

At this crossing, Lemma OA.29(iv) and (OA.3-60) imply $(z_3, x(z_3)) \in \mathcal{K}$. Lemma OA.32 and (H3) therefore give

$$n'(z_3) = \frac{1}{2}\Psi\big(z_3, x(z_3)\big) > 0, \tag{OA.3-74}$$

a contradiction. The set $\{z : \beta(z) \geq z\}$ is consequently an upper set. Its complement is nonempty near z^* and the set itself is nonempty near one, so its lower boundary $\widehat{z}$ is an interior crossing. The inequalities are strict on either side: equality at any later point would, by the same positive-derivative argument, require $n < 0$ immediately to its left, contradicting the upper-set property. Finally, Proposition OA.31 gives $\beta(z^{**}) = 1 > z^{**}$, so $\widehat{z} < z^{**}$. ■

Condition (H3) is a two-dimensional region check involving only primitive parameters. If $\rho \geq r + \lambda$, then $2\rho/k - 1 \geq 0$ and $x \mapsto \Psi(z, x)$ is strictly increasing on $[\eta(z), D(z))$. In that case, the simpler one-dimensional condition

$$\Psi\big(z, \eta(z)\big) > 0 \qquad \text{for every } z \in [\underline{\zeta}, 1] \tag{OA.3-75}$$

implies (H3). These conditions are sufficient rather than necessary because the primitive region $\mathcal{K}$ contains slope–state pairs that need not arise along the optimal policy.

OA.4 Utility Weighting and Tail Outcomes

This section focuses on the baseline model without project termination ($\lambda = 0$) and examines whether the utility weighting in the stock of promised future incentives is driven primarily by rare tail histories. Under the optimal contract, (10) implies

$$W_t = W_0 \exp\left(-\frac{1}{2}\sigma^2\gamma^2 r^2 \int_0^t \beta_s^2\, ds - \sigma\gamma r \int_0^t \beta_s\, d\hat{Z}_s\right). \tag{OA.4-1}$$

The optimal policy satisfies $\beta_t \in [\beta^*, 1]$. Hence the quadratic variation of the stochastic integral in (OA.4-1) grows at least linearly, and the stochastic-exponential law of large numbers gives

$$W_t \to 0 \qquad \text{almost surely.} \tag{OA.4-2}$$

Because zero is the upper bound on the agent's utility, the agent's utility therefore approaches its upper bound along almost every sample path as $t \to \infty$. At the same time, W_t is a martingale and $\mathrm{E}[W_t] = W_0 < 0$ at every finite date. The family $\{W_t\}_{t \geq 0}$ is consequently not uniformly integrable: its typical realization approaches zero, while its constant mean is sustained by a vanishing probability of increasingly extreme negative outcomes. This produces a wedge between the mean and typical outcomes.

This asymptotic behavior raises a potential concern about the economic mechanism. Because z_t is the utility-weighted average

$$z_t = \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} (r+\rho) \frac{W_s}{W_t} \beta_s \, ds \right], \tag{OA.4-3}$$

rare histories with very negative W_s receive large relative weight. One might therefore conjecture that promised future incentives operate mainly through remote states in which the agent suffers extreme losses. For comparison, define the unweighted average of future PPS by

$$\beta_t^{\mathrm{u}} \equiv \hat{\mathrm{E}}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} (r+\rho) \beta_s \, ds \right]. \tag{OA.4-4}$$

If rare tail histories dominated the incentive calculation, z_t would substantially exceed β_t^{u}.

The right panel of Figure 2 shows that z_t does exceed β_t^{u}, consistent with future PPS being concentrated in low-utility, high-marginal-utility states. The gap is modest throughout the baseline state space, however. The reason is that the mean–typical-outcome wedge becomes pronounced primarily at long horizons, whereas the definition of z_t discounts future PPS at rate $r+\rho$. The principal's objective also discounts future compensation at rate r. Thus, the state-contingent placement of future PPS is economically relevant, but its incentive effect comes primarily from moderate states in the near future rather than from vanishingly rare disasters in the distant future.