

Buying In and Selling Out: The Dynamic Returns to Investing in Expertise*

Felix Zhiyu Feng[†] Mark M. Westerfield[‡]

October 2026

Abstract

We study medium-term optimal contracts for a knowledge worker with hidden, durable, and portable expertise who is matched with a sequence of principals through the labor market. Employment relationships are sufficiently persistent for incentives to evolve but remain temporary, allowing future contracts to shape current incentives. The model therefore characterizes career trajectories within and across jobs in a unified framework. Within a job, the worker may be promoted from salary-like compensation to equity-like pay after sufficient tenure, or begin as a residual claimant and sell part of his stake following good performance. Across jobs, the division of expected match surplus generates distinct career regimes: a serial founder repeatedly begins with full ownership, whereas a serial salaried worker begins with full insurance and waits for promotion. Because the labor market links incentives across successive contracts, allocating all match surplus to either party need not maximize that party's welfare.

JEL Classification: D82, D86, J24, J41

Keywords: dynamic agency, persistent private information, expertise, human capital, medium-term contracts, serial employment, labor market

*We thank Philip Bond, Roozbeh Hosseini, Martin Gervais, Fahad Khalil, Rachel Kranton, R. Vijay Krishna, Jacques Lawarrée, Benjamin Lester, Adrien Matray, Andrew McClellan, Marcus Opp, Stefano Pegoraro, Marzena Rostek, Peter Sorensen, Curtis Taylor, Xu Tian, Noah Williams, Huseyin Yildirim, and Erina Ytsma, seminar participants at Duke University, University of Georgia, University of Washington, and participants at the Atlanta Finance Workshop, the Econometric Society North American Summer Meeting, the St. Gallen Workshop on Dynamic Contracts and Games, the Stony Brook International Conference on Game Theory, and the Finance Theory Group Member Meeting at NYU for helpful comments and suggestions. Robin Yifan Luo provided excellent research assistance.

[†]Terry College of Business, University of Georgia. Email: ffeng@uga.edu

[‡]Foster School of Business, University of Washington. Email: mwesterf@uw.edu

1 Introduction

Dynamic incentives are typically studied at two extremes: long-term contracts with a single principal, where skills accumulate over time, or a sequence of short-term engagements, where skills are fixed or evolve exogenously. Yet, most career paths lie between these extremes, as employment relationships in practice are neither permanent nor purely spot-market interactions. In the United States, the median job length is about 3.9 years, and individuals work an average of 12.9 jobs between ages 18 and 58.¹ This pattern indicates that employment relationships are *medium-term*: long enough for each job to shape workers’ incentives and skills, but short enough that the prospect of future employment is also an important driver of incentives.

As individuals work in a sequence of finite-term roles, they gain knowledge and experience, hone their skills, and build networks, all of which improve productivity and contribute to the accumulation of *expertise*. Critically, such expertise is usually hidden, durable, and inalienable; it is only partially observable to employers, and workers take it with them when moving on to their next job. A biotech scientist whose clinical program is terminated, a petroleum engineer whose field project is canceled, or a software manager whose startup shuts down due to industry shocks moves to another employer in a similar role, carrying tacit knowledge and accumulated expertise into the next match.² Founders and small-business owners often re-enter entrepreneurship after an exit or closure, bringing with them prior operating experience, investor relationships, and reputational capital to their new businesses.³ Because agents carry their expertise from one job to another, realistic career paths typically combine two trajectories. Within an employer, individuals can progress from entry-level positions to managerial roles and sometimes become part owners, with each stage carrying its own incentive and compensation structure. Across employers, individuals can move to new firms or found their own ventures. In some settings, employees work their way up within the same firm as their responsibilities change. In others, switching employers allows them to realize the returns to human capital, and launching new enterprises is common.

Motivated by this combination of career paths and the varied means of realizing the value of expertise, this paper studies medium-term contracts with persistent private information in the form of the agent’s expertise. Each contract offers dynamic incentives that respond to cash-flow shocks until an exogenous project-level shock makes the current match unproductive and the agent is redeployed to a new principal. This contracting environment is

¹Bureau of Labor Statistics, U.S. Department of Labor, 2025.

²See, e.g., Maital and Barzani (2021); Sun, Gao, Hu, and Zhou (2022) for more detailed facts and statistics.

³Lafontaine and Shaw (2016) find that more than a quarter of retail businesses were started by owners who had started at least one business before. Shaw and Sørensen (2019) report that serial entrepreneurs account for 19% of entrepreneurs and operate 29% of new firms.

materially different from the long-term contracts studied in Williams (2011) and Marinovic and Varas (2019), which focus on a single principal, and from the short-term contracts in Gibbons and Murphy (1992) and Holmström (1999), in which the agent works with multiple principals but only for a single period. In contrast, our framework allows us to study *mixed career paths* in which an agent can advance under a single employer, changing his role and responsibilities as his expertise accumulates, or can change responsibilities by moving to another employer. Within a single framework, the model generates distinct career patterns commonly observed in practice. When future surplus splits favor principals, the agent begins each relationship with low-powered incentives and waits for his expertise to be rewarded through future promotion-like incentive pay. When future surplus splits favor the agent, he begins each relationship as a residual claimant but gradually sells part of the claim after good performance. Thus, the same contracting environment can resemble either repeated salaried jobs with promotion to joint ownership or repeated founding of new ventures and partial exit, depending on how future surplus from expertise is shared.

We analyze the contractual relationships of a knowledge worker, or agent, who is employed by a sequence of principals. Each principal operates a project whose output depends on the agent’s expertise and an idiosyncratic cash-flow shock. The agent can make costly investments in expertise that raise future productivity. Principals are risk neutral, while the agent has CARA utility. Compensation can therefore insure the agent, reimburse investment costs, and expose him to output risk through pay-performance sensitivity (PPS).⁴

Crucially, the agent’s relationship with each principal is not permanent, as each project is only productive for a stochastic amount of time. When the project becomes unproductive, the agent moves on to new employment and begins a new contract, carrying his expertise with him. This means that expertise is valuable beyond the current project, but the current principal cannot fully capture its future value. This creates a dynamic externality. A principal who incentivizes investment today raises the agent’s expertise not only for the current project but also for future projects with other principals. At the same time, the agent’s willingness to invest today depends partly on what he expects future principals to pay for that expertise. Thus, incentives within the current contract and incentives from future labor-market opportunities are jointly determined.

To make this interaction tractable, we introduce a reduced-form labor market in which the agent is paired with a new principal as soon as he separates from the previous one.⁵ When the agent and the new principal meet, they write a contract that splits the expected surplus of their employment relationship. We show that in equilibrium, the agent’s share of this future

⁴We restrict attention to contracts in which the agent’s share of cash-flow shocks lies between zero and one, so pay-performance sensitivity can be interpreted as an inside-equity or wealth-performance share. A value of zero corresponds to wage-like compensation, while a value of one corresponds to full residual claimancy.

⁵In other words, we abstract away from any search friction in the labor market.

surplus determines the labor market value of expertise, which in turn must be consistent with the incentives embedded in the contracts that future principals offer. Therefore, each contract contains two linked sets of promises. The first is the agent’s continuation utility, which determines his consumption and insurance. Its sensitivity to output shocks is summarized by current PPS. The second is the stock of incentives: the expected present value of all future PPS, normalized by the marginal value of savings. This stock of incentives is the shadow value of expertise inside the contract. Because investment in expertise affects output persistently, the agent’s investment decision depends not just on current incentives but also on the entire future path of incentives.

The distinction between current incentives and the stock of future incentives is central. A contract can provide low current performance pay while still giving the agent strong incentives through anticipated future performance pay. Conversely, a contract can front-load incentives by making the agent a residual claimant today while promising weaker incentives later. The current principal chooses the path of incentives, but the starting point of that path is pinned down by the labor market: when the agent enters a new contract, the value he receives for his expertise must equal the stock of incentives required to create and truthfully report that expertise.

Our model reveals three sets of results, each matching key features of common career paths while generating novel predictions for career dynamics within and across jobs. First, we show that the optimal sequence of contracts has endogenous incentive *regions* that generate a form of gradual “buying in.” Early in a job, the agent may receive salary-like compensation, with no exposure to output risk. As incentives accumulate, compensation can become equity-like: positive and negative cash-flow shocks are partly passed through to the agent, making his payoff resemble a partial ownership claim. When the stock of incentives is sufficiently high, the agent’s current PPS equals one, so the agent becomes a full residual claimant on cash flows, resembling an owner.

We find that the stock of incentives responds negatively to output shocks, so bad performance raises continuation incentives. The mechanism comes from persistence in the conditional mean, not autocorrelation in the underlying Brownian shocks. If the agent overstates his durable expertise, the innovation inferred from output acquires a persistent negative drift because the reporting error survives over time. Under the deviation, negative innovations therefore predict further negative realizations, in contrast to learning models in which the agent can manipulate the principal’s beliefs, as in [He, Wei, Yu, and Gao \(2017\)](#) or [Cisternas \(2018b\)](#). In those models, positive performance raises beliefs and may consequently predict later negative innovations. Thus, revelation and learning generate different relationships between current performance and the conditional distribution of future incentives and output.

Second, recontracting with new principals links contractual incentives to the agent’s surplus share, generating distinct *regimes* with different policy combinations. When the agent receives a small fraction of the surplus created by his expertise in the labor market, the optimal contracts begin with low levels of incentives and low performance sensitivity. The agent is paid a wage that grows with his expertise, resembling the compensation of a skilled professional, but his compensation is not sensitive to performance. We thus refer to this as the “serial salaried-worker regime”. If the agent remains matched to the current principal long enough, he enters a region with positive PPS. In other words, *the agent can be promoted* into a position with a share of cash flows (e.g., equity) and incentives to invest in expertise. In contrast, when the agent receives a large fraction of the surplus in the labor market, he starts each relationship with a new principal as a full residual claimant on cash flows. We thus refer to this as the “serial founder regime”. However, because the agent’s stock of incentives is negatively correlated with cash-flow shocks, a series of positive cash-flow shocks reduces the agent’s pay-performance sensitivity, implying that the agent *sells part of the firm* after sufficiently good performance.

Critically, the agent’s welfare is not always maximized in economies in which he receives all the surplus, as in economies that closely resemble competitive screening. This is because optimal contracts link surplus in the labor market to incentives during employment, and full surplus for the agent implies maximum incentives. However, exposing the risk-averse agent to the entire volatility of cash flows prevents risk sharing. The agent would like to face a labor market that grants him the full surplus while contracts provide only moderate incentives, but that is not possible because future labor market rents constitute most of the agent’s incentives. Similarly, the principal’s welfare is not always maximized when she receives all the surplus, because that implies no incentives for the agent. The principal would prefer to face a labor market that grants her the full surplus while still incentivizing investment, but this is also not possible.

Finally, these regimes correspond to different stationary distributions of expertise in our numerical analysis. The serial salaried-worker regime corresponds to a low level of initial incentives following each separation and produces a stationary distribution with a low mean concentrated at lower levels of expertise. In the serial founder regime, each separation resets the agent’s stock of incentives to a high level, and the incentives from future contracts are sharper than those from the current contract following positive cash-flow shocks. Importantly, the agent’s return to investing in expertise arises not only from the current contract but also from contracts with all future principals. Analytically, the invariant interval containing expertise shifts upward as the agent’s share of the surplus increases; numerically, the stationary distribution shifts with it. Consequently, policies that aim to increase the fluidity of the labor market may not always improve workers’ welfare or increase the aggre-

gate accumulation of human capital. If workers do not expect to capture a sufficiently large fraction of the surplus generated by their expertise each time they re-enter the labor market, a high-turnover labor market may discourage investment in expertise by repeatedly resetting workers to low-powered contracts and truncating the build-up of within-job incentives.

Relation to Existing Literature

Our paper bridges several strands of literature, including dynamic contracting models with persistent private information and career-concerns models. These dynamic contracting models can be further divided into two categories based on how persistence is introduced. In one category, the agent possesses private information about an exogenous state, such as a persistent underlying state or serially correlated shocks.⁶ The second category includes Williams (2011, 2015), Sannikov (2014), Marinovic and Varas (2019), Di Tella and Sannikov (2021), and Bloedel, Krishna, and Strulovici (2023) and assumes that the agent can take a private action that has long-term effects. Feng, Luo, and Westerfield (2024) analyze how the incentives for such an action interact with the incentives for short-term investments that only boost current output in a long-lived firm. Our key innovation is to put the agent in a series of contracts that generates endogenous terminal payoffs and incentives. Incentive consistency jointly restricts both the recontracting stage and the sequence of incentive contracts. These restrictions generate different types of contracts (e.g., serial salaried worker versus serial founder) with variation across types in policies and incentives.

Our paper is also related to the literature on investment incentives stemming from future employment, particularly the seminal work on career concerns by Holmström (1999) and Gibbons and Murphy (1992). These models usually feature an agent with an innate type (e.g., a level of expertise) who can exert costly effort to influence the output of a production technology but not his type. The other parties offering compensation contracts to the agent cannot observe the agent’s type or effort, but they can infer the type from the agent’s observable performance. Thus, these models illustrate how future reputational returns provide incentives for agents to exert effort even without an explicit long-term contract, and how contract offers take into account the implicit incentives generated by career concerns.⁷

⁶Examples include Fernandes and Phelan (2000), Kennan (2001), Zhang (2009), Kapička (2013), He, Wei, Yu, and Gao (2017), Cisternas (2018b), Fu and Krishna (2019), Krasikov and Lamba (2021), and Bloedel, Krishna, and Leukhina (2025). Special cases of this literature are the weakly time-separable problems studied in Pavan, Segal, and Toikka (2014) and Bergemann and Strack (2015). In these problems, although the agent’s private information is persistent, the optimal mechanism can be reduced to a static screening problem at time zero.

⁷Notable recent work on similar topics includes Bonatti and Hörner (2017), Hörner and Lambert (2021), and Ke, Li, and Safronov (2023) studying career concerns with Poisson signals, ratings, and endogenous task choice, respectively. Häfner and Taylor (2023) consider the substitution between explicit incentives from the current employer and implicit incentives from future employers in a serial employment market without expertise. Finally, Kaniel and Orlov (2025) consider the case in which the principal (e.g., a financial

Our model shares with the career-concerns literature the feature that the agent’s incentives for investing in his expertise are jointly determined by the optimal contract with the current principal and implicit incentives from future principals. However, the agent in our model knows his own type. Thus, we study the interaction of investment and revelation, rather than investment and learning. In our model, future principals are not merely trying to infer a fixed type from past performance. Instead, the agent privately accumulates expertise through costly investment, and this expertise affects current and future output. Reports about expertise are therefore payoff-relevant over time. Because expertise is durable, a misreport today is also durable: it changes the principal’s interpretation of output innovations not only in the current relationship, but also in later relationships if the report is carried forward.⁸ Compared with the model in [Cisternas \(2018a\)](#), our model additionally requires contracts to induce truthful revelation. Unlike [Board and Meyer-ter-Vehn \(2015\)](#), continuation incentives arise through bilateral contracts and the agent’s verifiable disclosure of past reports, rather than through inference by a large competitive reputation market.

Moreover, the endogenous interaction between repeated contracting and the labor market creates regimes that differ in their starting points and in how incentives evolve over time, distinguishing our model from existing studies in this area that examine either worker promotion or changes in ownership.⁹ Our paper establishes a unified framework in which *both* the serial salaried-worker and serial founder regimes arise endogenously depending on the distribution of surplus in future contracts. Moreover, most existing studies consider either static or finite-period models in which promotion or the distribution of ownership occurs only once. Our model is fully dynamic, implying that the agent may be promoted from salary to equity-based compensation or may take full ownership repeatedly, and it is precisely the anticipation of those events that provides the agent with different incentives to invest in his expertise. Depending on his performance history, the agent may exit either the salary region or the full-ownership region, or both, during the same employment relationship, generating practical implications for the dynamics of expertise accumulation and the distribution of ownership within the firm. In related work, [Fuchs and Garicano \(2010\)](#) study whether hidden expertise is better allocated within firms or through markets, while [Fuchs, Garicano, and Rayo \(2015\)](#) derive firm-like contracts with fixed pay, performance bonuses, and expert residual claimancy. Unlike these static matching models, ours studies endogenous expertise

intermediary) is more informed about the agent’s type than the outside market and uses such an informational advantage to incentivize the agent to work for her at a low wage.

⁸To focus on this persistence of private information, we assume that the agent can verifiably disclose past reports and contract terms to future principals, even though principals cannot directly communicate with one another or write contracts that bind future principals. This structure partially mitigates adverse selection across jobs, but it also means that lies about expertise persist across jobs.

⁹See, e.g., [Zabojnik and Bernhardt \(2001\)](#), [Golan \(2005\)](#), and [Fudenberg and Rayo \(2019\)](#), or [Aghion and Tirole \(1997\)](#) and [Pagano and Röell \(1998\)](#).

and the evolution of incentives within and across employment relationships.

Finally, in our model, each principal and the agent commit to the contract during their employment relationship. In that regard, our model is categorically different from those in the relational contracting literature with persistent private information. In those studies, one or both contracting parties can renege on existing contractual terms via unilateral termination or renegotiation, and the option to do so continuously typically results in spot or stationary contracts.¹⁰ We follow the convention of dynamic contracting by assuming commitment within each employment relationship. Anticipated separation and the contract within each employment relationship jointly determine the agent’s incentives to invest in expertise.

2 The Model

We study the problem of an agent (he) who faces a series of principals (she). The agent is risk averse and has durable, hidden, and inalienable expertise and financial savings, while principals are risk neutral, have deep pockets, and own real assets. Expertise and real assets can be combined by principal-agent pairs in projects that produce risky output. Each project terminates following a public technological shock, after which the agent is rematched to start a new project. A competitive fringe of low-quality principals determines the parties’ outside options and therefore the net surplus divided when a new productive pair forms. Time is continuous, and principals and the agent share a discount rate $r > 0$.

2.1 Technology and Preferences

Throughout his lifetime, an agent possesses expertise $H_t \geq 0$, which grows with costly investment $h_t \geq 0$ and depreciates at rate $\rho > 0$:

$$dH_t = (h_t - \rho H_t)dt. \tag{1}$$

Expertise can be combined with a principal’s real assets in productive projects that generate cumulative output Y_t according to

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

¹⁰Notable examples include [Malcomson \(2016\)](#), [Kwon \(2016\)](#), and [Martimort, Semenov, and Stole \(2017\)](#). Note that without commitment, the specific properties of the equilibrium contract depend crucially on the exact assumptions about which parts of the contract are enforceable or renegotiable, and on how equilibrium is defined. See, e.g., [Watson, Miller, and Olsen \(2020\)](#), [Pegoraro \(2020\)](#), and [Strulovici \(2024\)](#) for more comprehensive discussions.

where $\sigma > 0$ is the volatility, and Z_t is a standard Brownian motion. We normalize the cash flow from expertise by $(r + \rho)$ so that one unit of expertise used in perpetuity, net of discounting and depreciation, generates a present value of \$1.

To model the need for serial employment, we assume each productive project is temporary and becomes unproductive following exogenous, public technological shocks arriving via a Poisson process N_t with intensity $\lambda \geq 0$. These shocks represent shifts in the best use of any particular expertise as a result of technological change. When such a shock arrives, the principal and the agent separate. The principal exits the game and the agent moves on to a new employer.

The agent possesses a savings account S_t that funds consumption c_t and receives compensation dM_t from his current principal. We assume that there are public asset markets that are complete with respect to public shocks but not private shocks (i.e., N_t but not Z_t), and the common discount rate r is also the expected return on savings and traded assets. Consequently, the agent's savings S_t evolve according to

$$dS_t = (rS_t - c_t)dt + dM_t + \nu_t(dN_t - \lambda dt), \quad (3)$$

where ν_t is the agent's public asset portfolio policy. In addition to being realistic, allowing trade in claims contingent on public shocks simplifies the equilibrium because it fixes the price of insurance against public shocks. The alternative is to force the agent to seek insurance for these shocks from the principal through the contract, and for the principal to profit by re-trading this insurance in public markets.

The agent has CARA utility over consumption c and investment in expertise h ,

$$u(c, h) = -\frac{1}{\gamma} \exp \left[-\gamma \left(c - \frac{\theta}{2} h^2 \right) \right], \quad (4)$$

where $\gamma > 0$ is the coefficient of absolute risk aversion and $\theta > 0$ is the cost parameter for investment in expertise. The quadratic cost can be interpreted as the financial or opportunity cost of education and training or a disutility from the effort to develop skills.¹¹

2.2 The Agency Problem

The agent's expertise, investment in expertise, savings, and consumption are hidden. A principal can directly observe only project output and compensation, and only for the project in which she participates.

¹¹In equilibrium, the principal reimburses the agent for recommended investment in expertise and training. Thus, our model is equivalent to one in which the principal pays for training and the agent chooses whether to shirk.

We use plain letters to denote true processes (e.g., H_t) which constitute the agent's information set. We use hats to denote requests from the principal or reports by the agent (e.g., $\hat{H}_t$), which are part of the principal's information set. The agent takes hidden actions that are not observable by any principal (c_t , h_t , and ν_t) and has persistent private information about his own expertise and savings (H_t and S_t). Because output (2) is observable, the agent's reports of expertise imply a path of shocks, which we denote $d\hat{Z}_t$. We write $E_t[(\cdot)_s]$ for an expectation conditioned on the true state of the economy and $\hat{E}_t[(\cdot)_s]$ for one conditioned on the current principal's information set.

In a slight abuse of notation, for any particular principal-agent pairing, we will use $t = 0$ to mean the initiation time and $t = \tau$ to mean the termination time. The economy starts at some unspecified time $t = -T$ in the past with known and public initiation values H_{-T} and S_{-T} . For example, it may be common knowledge that the agent begins life with zero expertise.

We further simplify the setting by restricting the contract so that the agent's share of cash flows is between zero and one. This restriction is the dynamic analogue of the standard static requirement that both the principal's and the agent's payoffs be non-decreasing in cash flows to prevent manipulation of those cash flows, as in [Innes \(1990\)](#). Formally, we require that there is some $\beta_t \in [0, 1]$ such that the output-contingent component of the agent's compensation process is $dM_t = (\dots)dt + \beta_t dY_t$.¹²

We assume the agent can be punished sufficiently severely that he never makes verifiable deviations from a particular contract and that lies cannot be unboundedly large:

Assumption 1 *The agent can commit not to quit early, and his reports must be outcomes that could have happened without deviations under some set of shocks. In addition, there are fixed finite constants $\bar{S}$ and $\bar{H}$, which may be chosen arbitrarily large, such that $|S_t - \hat{S}_t| \leq \bar{S}$ and $|H_t - \hat{H}_t| \leq \bar{H}$.*

Remark 1 *To prevent the agent from making verifiable deviations, the principal could require that the agent post a bond that is returned with interest when the project ends. In fact, as we show in [Proposition 3](#), the agent's wealth process (as opposed to consumption or continuation utility) is only pinned down when projects begin and end. Thus, a bond is simply a change in wealth that reverts when the project ends. Such a bond would act as a commitment device against observable deviations, and it can be made arbitrarily large. Examples used in practice include restricted stock grants or equity lock-ups (i.e., golden handcuffs). To simplify the discussion, we assume that the agent can commit to avoid verifiable deviations.*

¹²These bounds can be micro-founded. For example, if the agent has the ability to destroy cash flows, then $\beta \geq 0$ under the optimal contract. If the agent can also augment cash flows from private borrowing, then $\beta \leq 1$ under the optimal contract.

2.3 Labor Market and Objective Functions

Because the agent's contract with each productive principal ends when the technological shock arrives (i.e., at the stopping time τ satisfying $dN_\tau = 1$), we need to specify how the agent begins his next employment. To do so while maintaining tractability, we adopt a reduced-form labor market in which the agent, once separated from a productive principal, can be immediately paired with a new principal that operates the next productive project. That is, we abstract away from any search frictions in the labor market. Furthermore, we assume that a current contract cannot bind a future principal, nor can a principal make demands of a future or past principal or observe cash flows from other projects.

Nevertheless, the agent can relay information from his earlier work to his later principal. We assume that the agent's reports to an earlier principal can be verifiably revealed by the agent to any future principal, and similarly for other contractual details. Such disclosure is complete rather than selective.¹³ Thus, expertise is hidden, but the agent's reports of his expertise to one principal can be verifiably revealed to another. For example, if the agent documents expected output, that document also serves to implicitly report his expertise through (2), and that document can be verifiably reported to future principals. In this example, the difference between expected and realized output constitutes the agent's implicit report of the output shocks.

The agent's objective function is his continuation utility:

$$W_t = \max E_t \left[\int_t^\tau e^{-r(s-t)} u(c_s, h_s) ds + e^{-r(\tau-t)} R_\tau \right] \quad (5)$$

$$R_\tau = \max E_\tau \left[\int_\tau^\infty e^{-r(s-\tau)} u(c_s, h_s) ds \right] \quad (6)$$

where R_τ denotes the welfare the agent obtains when he separates from the current principal (i.e., the agent's residual continuation utility). It will be convenient to express the agent's continuation utility in financial terms as certainty-equivalent wealth, V_t^A :

$$\frac{1}{r} u(rV_t^A, 0) = W_t. \quad (7)$$

The principal's objective function is the expected discounted value of net output from

¹³This assumption can be interpreted as certified employment records, audited performance reports, vesting and ownership records, or other formal documentation that can be transferred across employers. It aims to circumvent known difficulties associated with multiple principals, such as those in common agency (e.g., Peters, 2001) or sequential contracting (e.g., Pavan and Calzolari, 2009). In contrast to those studies, we simplify the strategic revelation between principals in order to focus on the evolution of expertise and contracts over time. In our model, the agent can verifiably report past contracts to future principals. Consequently, there is no hidden information about previous contracts, only previous actions and states.

the project, taken under the principal's conditional expectation $\hat{E}_t$:

$$V_t^P = \max \hat{E}_t \left[\int_t^\tau e^{-r(s-t)} (dY_s - dM_s) \right]. \quad (8)$$

Each time the agent meets a new productive principal, they split the expected net surplus created by their relationship. To generate a notion of such surplus, we assume that there is a competitive fringe of low-quality principals that operate projects producing output without any effective use of the agent's expertise. Specifically, if the agent and the principal with the productive project do not agree on a contract, the agent is instead paired with a fringe principal, generating output $dY_t = \sigma dZ_t$ until the next separation shock arrives. Naturally, in equilibrium, the agent only works for principals with productive projects and never for a fringe principal.

We denote the agent's certainty-equivalent value at contract initiation from one spell with the competitive fringe, followed by equilibrium productive pairings, as $V_0^{\chi A}$. Because fringe principals are assumed to be competitive, they earn zero profits from matching with the agent, i.e., $V_0^{\chi P} = 0$. Thus, the *net* surplus created by the agent's employment relationship with a productive principal, evaluated at contract initiation, is $V_0^A + V_0^P - V_0^{\chi A}$. We assume the agent receives a fraction $\kappa \in [0, 1]$ of this net surplus and the productive principal receives the remaining fraction $1 - \kappa$. That is,

$$\hat{E}_0 [V_0^A - V_0^{\chi A}] = \kappa \hat{E}_0 [V_0^P + V_0^A - V_0^{\chi A}] \quad (9)$$

$$\hat{E}_0 [V_0^P] = (1 - \kappa) \hat{E}_0 [V_0^P + V_0^A - V_0^{\chi A}]. \quad (10)$$

Notice that the split takes place under the principal's measure, so the pair divides the surplus evaluated under the agent's equilibrium reports, which equals the actual surplus in equilibrium. In the ensuing analyses, we take κ as a parameter of the model, and explore in Section 5 the implications for the equilibrium contract and incentive dynamics as κ varies from 0 to 1.

2.4 Equilibrium in Contracts

The agent's decisions are motivated by his lifetime incentives, so incentive compatibility for the current contract must take account of anticipated future contracts during the agent's life. The agent's current incentives are the sum of *explicit* incentives from the current contract and *implicit* incentives from the labor market. In addition, each contract begins with the agent reporting his activities from his previous employment relationship. But those reports are only truthful if the *combination* of incentives from the previous contract and the new

contract induces accurate reporting. Thus, we will refer to incentive compatibility and optimality as properties of the sequence of contracts as opposed to individual contracts.

In addition, we appeal to standard efficiency arguments and the revelation principle to restrict our attention in two ways. First, the principal will always require the agent to reveal his entire history of reports, including the agent's most recent reports of $\hat{H}_{0-}$ and $\hat{S}_{0-}$ to the previous principal (as described in Section 2.3 above). Revealing this information is free and removes a layer of adverse selection. Second, contracts induce a zero private portfolio position ($\nu = 0$). Any exposure to publicly traded assets that the agent could obtain privately can be replicated through contractual transfers by the risk-neutral principal, who has access to the same assets. A private portfolio position therefore adds no value.

Definition 1 (Contracts) *A contract $\mathcal{C}$ specifies a recommended path of actions $\{\hat{c}_t, \hat{h}_t, \hat{\nu}_t\}$ and a cumulative compensation process M_t . These recommendations and compensation are conditioned on the reported histories of $\{\hat{H}_t, \hat{S}_t\}$ over $(-T, t]$, past contracting details, and the observed histories of Y_t and N_t , which together comprise the principal's information set.*

A sequence of contracts is incentive compatible if the agent is induced to accurately report the path of $\{H_t, S_t\}$ and to choose $\{c_t, h_t, \nu_t\} = \{\hat{c}_t, \hat{h}_t, \hat{\nu}_t\}$ with $\hat{\nu}_t = 0$ over his lifetime.

A sequence of contracts is optimal if (i) the sequence is incentive compatible, (ii) each contract maximizes the respective principal's welfare taking future contracts in the sequence as given, and (iii) the agent receives his promised utility at each contract initiation, with the promised level determined by the distribution of surplus in the labor market, (9) and (10).

Incentives and welfare require the principal and the agent to forecast what the agent will receive when the current contract ends — to look forward to the next principal-agent pairing as part of evaluating the current contract. We require those forecasts to be consistent with the equilibrium distribution, and we will search for equilibria in which the value of expertise is *linear*:

$$R_t = R(H_t, S_t) = -\frac{1}{\gamma r} \exp \left[-\gamma r (\bar{V} + S_t + \zeta H_t) \right], \quad (11)$$

where R is the agent's residual utility from (6), and ζ and $\bar{V}$ are constants. We will refer to ζ as the shadow value of expertise from future contracts, while $\bar{V}$ is the remaining certainty equivalent wealth after accounting for H_t and S_t .

Now we can define an equilibrium for the economy:

Definition 2 (Equilibrium Contracts) *A linear equilibrium is an optimal sequence of contracts such that (11) holds and surplus is distributed as in (9) and (10).*

To assess the existence and properties of an equilibrium, we use a guess-and-verify approach: we assume (11) is a correct forecast of the future and show that it gives rise to a valid

equilibrium. Notice that the assumed form for the residual utility must also hold at the initiation of any given contract, because the initial value of a new contract is the residual value of the old one. In that sense, we are looking for a stationary equilibrium.

Finally, we impose conditions that ensure positive social value from employing the agent and provide the curvature restriction used in the policy analysis:

Assumption 2 $\sigma^2\gamma r\theta < \left(\frac{r+\rho+\lambda\zeta}{r+\rho+\lambda}\right)^2 - \left(\frac{\lambda\zeta}{r+\rho+\lambda}\right)^2$ and $\sigma^2\gamma^2r^2 < r + \rho$.

The conditions are easily satisfied, for example, if γr is sufficiently small.¹⁴

2.5 Benchmark Economies

In this section, we analyze two benchmark economies without agency problems. The first benchmark is the first-best, which maximizes social surplus without hidden information. In the second, the agent owns the productive technology and does not require a principal. In this ownership benchmark, the risk-averse agent absorbs output volatility, but there is no agency problem and the project does not end.

Proposition 1 *In the first-best economy, the agent's continuation utility is constant (i.e., $dW_t = 0$), and the combined value to the sequence of principals is given by¹⁵*

$$V_t^P = H_t + \frac{1}{\gamma r} \ln(-\gamma r W_t) + \frac{1}{2\theta r}. \quad (12)$$

If the agent owns the firm, the agent's certainty equivalent wealth from (7) is

$$V_t^A = S_t + H_t + \frac{1}{2\theta r} - \frac{1}{2}\gamma\sigma^2, \quad (13)$$

and $dW_t = -\sigma\gamma r W_t dZ_t$ and $dS_t = (rS_t - c_t + (r + \rho)H_t)dt + \sigma dZ_t$.

In both economies,

$$h_t = \frac{1}{\theta} \quad (14)$$

$$c_t = -\frac{1}{\gamma} \ln(-\gamma r W_t) + \frac{1}{2\theta}. \quad (15)$$

The derivations of these benchmarks are standard and presented in the Online Appendix [OA.1](#).

¹⁴Although the first condition involves the equilibrium quantity ζ , Theorem 1 shows that $\zeta = \kappa$, so the condition can be read as a restriction on primitives. Moreover, because the right side is increasing in ζ , setting $\zeta = 0$ gives a sufficient condition in primitives that covers all $\zeta \in [0, 1]$: $\sigma^2\gamma r\theta < \left(\frac{r+\rho}{r+\rho+\lambda}\right)^2$.

¹⁵This refers to the social value of production net of the agent's promised utility for the full sequence of principals. Thus, project-ending shocks only reassign production across principals, and λ does not appear.

The principal’s insurance of the agent in the first-best enhances welfare, but it does not affect the choice of investment relative to a model in which the agent owns the full output of his expertise. The agent’s continuation utility is constant in the first-best, but when the agent owns the firm, W_t has positive volatility as noisy output is added to savings. Similarly, consumption is fully smoothed in the first-best and fully exposed to shocks if the agent owns the firm. Despite this, the participants’ valuations (12) and (13) show that the two economies place the same (full) value on expertise, H_t . Investment h_t is chosen to set the marginal benefit — an additional unit of H_t , valued at 1 — equal to the marginal cost, θh_t . Thus, both economies have the same high and constant optimal investment in expertise despite the different risk-sharing arrangements.

The analysis below shows that moving from principal to principal in a labor market leaves the linear structure of V_t^A and V_t^P intact but induces variation in the relative price of expertise.

3 The Agent’s Problem

In this section, we analyze the agent’s optimal response to any given set of incentives when working for a sequence of principals, with each principal-agent pair taking outside options as given. The results lead to the next sections in which we describe how the principal chooses incentives and then how equilibrium contracts and surplus are determined.

We proceed in two steps. First, we analyze a relaxed version of the agent’s problem: given that the agent has reported truthfully so far and is engaged in a sequence of productive projects, what are the necessary and sufficient conditions for a lifetime path of recommendations to be incentive compatible? We characterize what any such lifetime looks like. Second, we analyze the compensation package and the agent’s future opportunities as part of a lifetime of incentive-compatible contracts. Proofs of the results for this section are in Appendix A.

3.1 Lifetime Incentives

The agent maximizes his lifetime utility (5) and (6), subject to the production technologies (1), (2), and (3), and the sequence of contracts he expects. There are two state variables that together characterize the incentives that the contract provides for the agent. The first state variable is continuation utility, $W_t \leq 0$, which summarizes promises from the sequence of principals to the agent, both explicit through the current contract and implicit through outside options. The second state variable is z_t , which is the shadow value of expertise in financial terms. As formalized below, one can think of z_t as $\frac{\partial W_t}{\partial H_t} / \frac{\partial W_t}{\partial S_t}$, the marginal rate

of substitution between expertise and savings. For comparison, ζ is the shadow value of expertise in financial terms offered by the labor market, $\frac{\partial R_t}{\partial H_t} / \frac{\partial R_t}{\partial S_t}$ using $R(H, S)$ in (11). In the benchmark economies of Section 2.5, these marginal rates of substitution were constant at 1.

We can use Hamiltonian methods to find necessary and sufficient conditions for the agent to report truthfully and to follow the recommended levels of investment and consumption:

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

$$dW_t = -\sigma\gamma r W_t \beta_t d\hat{Z}_t \quad (16)$$

$$\hat{h}_t = \frac{1}{\theta} z_t \quad (17)$$

$$rW_t = u(\hat{c}_t, \hat{h}_t) \quad (18)$$

$$dz_t = (r + \rho) (z_t + \sigma^2 \gamma r \beta_t \phi_t - \beta_t) dt + (r + \rho) \sigma \phi_t d\hat{Z}_t + (\zeta - z_t) (dN_t - \lambda dt) \quad (19)$$

where $0 \leq \beta_t \leq 1$, $\frac{\lambda\zeta}{r+\rho+\lambda} \leq z_t \leq \frac{r+\rho+\lambda\zeta}{r+\rho+\lambda}$, and W_t is a martingale.

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

We go through each result in sequence. First, W_t is the agent's continuation value, with the law of motion given by (16):

$$dW_t = -\sigma\gamma r W_t \beta_t d\hat{Z}_t = -\gamma r W_t \beta_t (dY_t - (r + \rho) \hat{H}_t dt) \quad (20)$$

where $-\gamma r W_t \beta_t \geq 0$ measures pay-performance sensitivity (PPS). When realized output exceeds its conditional expectation, the agent is rewarded with more continuation utility in proportion to $-\gamma r W_t \beta_t$ and the unexpected shock.

The agent's ability to hide savings and consumption forces W_t to be a martingale (16). To prevent the agent from undoing incentives using hidden savings and consumption, any changes to the agent's marginal utility must be entirely unexpected: the principal must recommend a level of consumption that makes the agent's marginal utility a martingale. Similarly, any Poisson jump in the agent's marginal utility could be offset by trading in public markets.¹⁶ Because the agent has CARA utility, his marginal utility is proportional to his utility; both are martingales, and we have $u(\hat{c}, \hat{h}) = rW_t$. The absence of wealth effects

¹⁶If the agent's marginal utility changed in expectation, he could benefit from manipulating his savings and consumption to increase utility. The agent would consume less when marginal utility was expected to increase and consume more when marginal utility was expected to decrease. Similarly, because the agent is risk averse, he will undo any Poisson jump in marginal utility by trading privately using public claims on technological change. Thus, no sequence of contracts with a jump in continuation utility can be incentive compatible.

with CARA utility means that we do not have to keep track of a second set of incentives regarding savings. Instead, W_t captures continuation utility and $\partial W_t / \partial S_t = -\gamma r W_t > 0$ captures the incremental impact of savings on continuation utility.

If the effect of the agent's hidden investment in expertise were instantaneous, then characterizing the law of motion for continuation utility (16) would be sufficient for the agent's problem. However, expertise is a hidden and persistent source of output. To measure the agent's incentives to invest in expertise, we require the agent's gain from additional expertise.

If H_t is increased by one unit, then the drift of output increases by $(r + \rho)$, with the difference discounted at rate r and decaying at rate ρ . The agent receives a $-\gamma r W_t \beta_t \geq 0$ portion of that output as additional continuation utility. In addition, the agent gains the marginal value of H from his next interaction with the labor market, $R_H(H_\tau, S_\tau)$. Adding everything up, we obtain the change in the agent's continuation utility in integral form:

$$\frac{\partial W_t}{\partial H_t} = \hat{E}_t \left[\int_t^\tau e^{-(r+\rho)(s-t)} (-\gamma r (r + \rho) W_s \beta_s) ds + e^{-(r+\rho)(\tau-t)} R_H(H_\tau, S_\tau) \right] \quad (21)$$

Thus, the marginal value of expertise is the expected discounted value of all future pay-performance sensitivity (PPS) during the current contract plus its value under future contracts.

Then, since $R_H(H_\tau, S_\tau) = -\gamma r \zeta R(H_\tau, S_\tau)$ from (11), and equilibrium implies $W_\tau = R_\tau$ from Definition 2, we can write $R_H = -\gamma r \zeta W_\tau$. Taking the ratio of the marginal value of expertise to the marginal value of savings ($\partial W_t / \partial S_t = -\gamma r W_t$, above), we obtain

$$\frac{\partial W_t / \partial H_t}{\partial W_t / \partial S_t} = z_t = \hat{E}_t \left[\int_t^\tau e^{-(r+\rho)(s-t)} (r + \rho) \frac{W_s}{W_t} \beta_s ds + e^{-(r+\rho)(\tau-t)} \zeta \right] \quad (22)$$

which is the integral form of (19). Critically, the agent's incentive to invest in expertise today includes *explicit* incentives from the current contract ($s \in [t, \tau)$) plus *implicit* incentives from anticipated future contracts ($s \in [\tau, \infty)$).

One can think of W_t and z_t as a joint account of promises over the lifetime of the agent. W_t is total promised continuation utility, while z_t aggregates promised PPS over the agent's lifetime. Because PPS is generated through the volatility of W_t , the agent's incentive to invest in expertise z_t also aggregates promised exposure of continuation utility to output shocks. A principal controls the contract and the agent's incentives by controlling the evolution of W_t and z_t . However, the duration of this pairing and thus any particular principal's control is limited, because z_t jumps to ζ every time the agent changes principals.

In addition, there is a consistency condition imposed by lifetime incentive compatibility. We can split continuation utility into two components, utility over the current contract (explicit promises) and residual utility (implicit promises from future principals), as in the

objective function (5). Consistency implies that the implicit promises from future principals are the expected value of future explicit promises. A similar logic applies to the marginal value of expertise. In (21), we split this value into explicit (current contract) and implicit (future contracts) terms. Consistency therefore allows us to express z_t entirely in terms of lifetime PPS:

$$z_t = \hat{E}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} (r + \rho) \frac{W_s}{W_t} \beta_s ds \right] \quad (23)$$

At separation, $z_\tau = \zeta$, so the labor-market value of expertise also determines the initial incentives in the next contract.

Recall that we restricted the contract so that the agent's share of cash flows is between zero and one. That is, $\beta \in [0, 1]$ (see footnote 12 for the micro-foundation for this restriction). Combined with the fact that W_t is a martingale, this implies a lower bound and an upper bound for z_t , denoted by z_L and z_R :

$$z_L = \hat{E}_t \left[\int_t^\tau e^{-(r+\rho)(s-t)} \frac{W_s}{W_t} \cdot 0 \cdot (r + \rho) ds + e^{-(r+\rho)(\tau-t)} \zeta \right] = \frac{\lambda \zeta}{r + \rho + \lambda} \quad (24)$$

$$z_R = \hat{E}_t \left[\int_t^\tau e^{-(r+\rho)(s-t)} \frac{W_s}{W_t} \cdot 1 \cdot (r + \rho) ds + e^{-(r+\rho)(\tau-t)} \zeta \right] = \frac{r + \rho + \lambda \zeta}{r + \rho + \lambda} \quad (25)$$

The bounds follow directly from $\beta_t \in [0, 1]$ and the martingale property of W_t , which implies $\hat{E}_t \left[\frac{W_s}{W_t} \right] = 1$. Their economic motivation comes from the manipulation argument underlying the restriction on β : negative PPS would induce the agent to destroy output, while PPS above one would induce him to augment output privately. Restricting PPS to $[0, 1]$ therefore also bounds the agent's incentives to create expertise. Because the agent expects to receive ζ from future opportunities, his current lifetime incentives cannot be too far from that value.

The condition $\phi_t \leq 0$ means that the forward-looking stock of incentives is negatively related to output, and it is the dynamic analog to a more standard static adverse selection condition that payoffs should be convex in the underlying type. Suppose the agent overreports his expertise, meaning $\hat{H} - H > 0$. Then output surprises will be negative on average and durable. Put differently, lies generate persistent distortions in reported output innovations. With $\phi \leq 0$, the evolution of z_t (19) implies that z_t will increase after negative surprises. Thus, incentives become sharper after the principal sees an indication that the agent has overrepresented his expertise, so upward lies result in more punishing incentives over time. Similarly, if the agent lies downward, output surprises will be positive, and z_t will decrease. Either way, incentives move against the predictable effect on cash flows from the agent's lies.¹⁷ While this is currently stated as a sufficient condition for the agent's in-

¹⁷If cash flows are lower than expected, it could be because of chance or because the agent has lied about

centive compatibility, as we will show in Section 4, $\phi_t \leq 0$ is indeed a feature of the optimal contracts.

3.2 Compensation and Future Opportunities

Once we have the laws of motion for the two sets of promises, W_t and z_t , we can derive the savings process S_t needed to support them, and therefore the compensation process for any incentive-compatible contract:

Proposition 3 (Compensation and Consumption) *Under equilibrium contracts, $W_\tau = R(H_\tau, S_\tau)$. There are many welfare-equivalent compensation processes, and we select one that maintains $W_t = R(H_t, S_t)$ for all t . Then, compensation, consumption, and certainty-equivalent wealth are given by*

$$V_t^A = S_t + \bar{V} + \zeta H_t, \quad (26)$$

$$\hat{c}_t = rV_t^A + \frac{\theta}{2}\hat{h}_t^2, \quad (27)$$

$$dM_t = \left(r\bar{V} + \frac{\theta}{2}\hat{h}_t^2 + \frac{\sigma^2\gamma r}{2}\beta_t^2 \right) dt + \zeta \left((r + \rho)\hat{H}_t - \hat{h}_t \right) dt + \sigma\beta_t d\hat{Z}_t. \quad (28)$$

For productive projects after the first one, there are no discrete transfers at contract initiation and $z_0 = z_\tau = \zeta$.

An important subtlety is that there are many welfare-equivalent compensation processes because incentive compatibility only determines S_t when a contract ends. We have chosen the path for which the agent is always indifferent between his contractually specified continuation utility and his outside option. One can always add a transfer process to a given compensation process, as long as the addition sums to zero by the time the contract ends. This is a slightly revised framing of the standard result that borrowing and lending between the principal and agent at the common discount rate is irrelevant if it does not change consumption.¹⁸

Under our continuous savings process, β_t governs the exposure of both compensation and continuation utility to output innovations. Thus, β can be interpreted as either pay-for-performance sensitivity or wealth-performance sensitivity. If $\beta_t = 0$, compensation is smooth and resembles an expertise-dependent salary. If $\beta_t > 0$, one can implement such a contract

H , with the critical difference that lies generate persistent surprises. Because the agent must be compensated for volatility, it is cheaper to offer incentives that weigh more heavily on paths induced by lying, meaning volatility (and hence PPS) increases after failures and decreases after successes.

¹⁸This is also the source of the potential bonding mechanism in the motivation of Assumption 1 that allows the agent to commit not to quit early. In addition, keeping the agent indifferent between his inside and outside values implies that the compensation process does not jump in the transition between principals. There may be a single jump at initiation in out-of-equilibrium pairs, e.g. when calculating the welfare the agent would get when pairing with the competitive fringe.

with equity and option grants, so that β is the sensitivity of those grants to realized cash flow shocks. This compensation process keeps the agent's certainty equivalent wealth equal to the sum of savings (S_t), the value of a consumption perpetuity ($\bar{V}$), and the value the agent obtains from expertise (ζH_t). The agent consumes this sum as if it were a perpetuity, with an addition from the principal to offset the cost of investment in expertise, totaling $rV_t^A + \frac{1}{2}\theta\hat{h}_t^2$.

The agent's compensation (28) simply offsets the required changes in the agent's certainty equivalent wealth. The first set of parentheses captures the consumption perpetuity ($r\bar{V}$), the cost of investment ($\frac{1}{2}\theta\hat{h}_t^2$), and compensation for risk ($\frac{1}{2}\sigma^2\gamma r\beta_t^2$). The second set of parentheses captures direct compensation for expertise. The term $(r + \rho)\hat{H}_t - \hat{h}_t$ is the net cash flow from expertise, net of investment, evaluated at the price offered by future principals, ζ . Thus, ζ is the outside value of H_t (from residual utility), which is also the present value of the agent's *fixed salary*. In contrast, z_t is the inside value of H_t (from the contract), which is also the present value of the agent's *incentive payments*.

4 The Principal's Problem

The principal maximizes the expected discounted value of net output from the current project (8). Her problem is conceptually simple: she controls the shadow value of expertise (z_t) by controlling how strong incentives are today (β_t) and the relationship between the stock of future incentives and output shocks (ϕ_t). We can use dynamic programming techniques to characterize the principal's value function and policy regions:

Proposition 4 (The Principal's Problem) *If the principal-agent pair forms a productive project, then*

$$V^P = v(z) + \frac{(r + \rho)(1 - \zeta)}{r + \lambda + \rho}H - \frac{r}{r + \lambda}\bar{V}, \quad (29)$$

where $v(z)$ is the principal's value function and is a viscosity solution of the principal's Hamilton-Jacobi-Bellman equation (B-11) and satisfies the boundary conditions (B-9) and (B-10) given in Appendix B. $v(z)$ is continuous, strictly positive, and strictly concave on (z_L, z_R) , with a unique interior maximizer $z^* \in (z_L, z_R)$. The principal's HJB implies the following policy characterization:

- $\beta(z) = \phi(z) = 0$ if $z \in (z_L, z^*)$.
- $\beta(z) > 0$ and $\phi(z) \leq 0$ if $z \in (z^*, z_R)$.¹⁹

¹⁹We can further strengthen this result to $\phi(z) < 0$ at every point at which $v(z)$ is twice-differentiable.




Figure 1: The Principal's Value and Policy Functions

The left panel of this figure plots the principal's value function $v(z)$ under the optimal contract. The top-right panel plots the optimal pay-performance sensitivity (PPS) $\beta(z)$ and $\beta_{stop}(z)$, defined in (30). The bottom-right panel plots the optimal relationship between the stock of future incentives and output shocks $\phi(z)$. z^* represents the point at which $v(z)$ is maximized, and z^{**} represents the left boundary of the $\beta_t = 1$ region. Parameter values are $r = 0.02$, $\rho = 0.5$, $\gamma = 3$, $\sigma = 12$, $\theta = 0.06$, $\lambda = 0.2$, and $\zeta = 0.5$.

- There exists $z^{**} \in (z^*, z_R)$ such that $\beta(z) = 1$ for all $z \in (z^{**}, z_R)$.

Appendix B derives the value-function representation and boundary conditions. Online Appendix OA.3 completes the proof. The principal's value function and policy characterization are illustrated in Figure 1.

The principal's value is maximized when negative output shocks induce harsher incentives: $\phi \leq 0$. The intuition arises from simple cost minimization: the agent's incentives are the *marginal-utility-weighted* pay-performance sensitivity (PPS), $\frac{W_s}{W_t} \beta_s$ in (22), but the principal only needs to compensate the agent for *unweighted* PPS, $-\frac{\sigma^2 \gamma r}{2} \beta_s^2$ in (B-6). Thus, incentives are most cost-effective when PPS is matched to high values of marginal utility $\frac{W_s}{W_t}$, meaning after negative cash flow shocks.²⁰ If incentives intensify after bad shocks, then z_t covaries negatively with output innovations ($\phi_t < 0$).

We will say that incentives are front-loaded or back-loaded based on the relative size of

We show in Online Appendix OA.3.2 that v is twice differentiable for Lebesgue-almost every $z \in (z^*, z_R)$. In Online Appendix OA.4 we further extend this result to all $z \in (z^*, z_R)$ under an additional technical assumption of a local nondegeneracy bound on $|\phi|$ that does not change the policy characterization stated here. Numerically, we have not found a case in which $v(z)$ is *not* smooth.

²⁰Under CARA, utility and marginal utility are proportional with opposite signs, and incentive compatibility implies $rW_t = u(c, h)$ (18). Thus, a high $|W_s|$ is equivalent to high marginal utility or low utility.

PPS today (β_t) compared to the total promised *future* incentives (z_t).²¹ Define

$$\beta_{stop}(z) = z + \frac{\lambda}{r + \rho}(z - \zeta) \quad (30)$$

which is the level of PPS that would be incentive-compatible and integrate to z_t if it were applied constantly for the remainder of the current contract. The subscript “stop” reflects that this policy keeps z_t constant until separation. We will say that PPS is front-loaded if $\beta_t > \beta_{stop}(z)$ because β_t is higher than its long-term requirement. Similarly, PPS is back-loaded if $\beta_t < \beta_{stop}(z)$. We plot $\beta(z)$ and $\beta_{stop}(z)$ in Figure 1.

The state space for incentives has two boundaries and several policy regions. At the boundaries, $z_t = z_L$ and $z_t = z_R$ generate $dz_t = 0$, so z_t does not move until the next technological jump. These are the states at which the principal has promised zero or full PPS for the remainder of the contract. At both points, incentives stop evolving. Between these boundaries, incentives evolve through distinct deterministic and stochastic policy regions, separated by z^* and z^{**} as illustrated in Figure 1.

We will refer to $z < z^*$ as the “*salary* region”, because the agent’s compensation package is smooth ($\beta_t = 0$) and resembles a salary. Total incentives are weak (since $z_t < z^*$) and are back-loaded, deterministic, and increasing (since $\beta_t = 0 < \beta_{stop}(z)$, $\phi_t = 0$, and z_t rises deterministically). The agent builds up his expertise in anticipation of future PPS and the principal’s value increases as the distance to that future time declines ($v'(z) > 0$ and z_t rises deterministically). In other words, the agent is *promoted* after a delay.

We will refer to $z > z^*$ as the “*equity* region” because the agent has an equity stake in the firm ($\beta > 0$). Total incentives are strong (since $z_t > z^*$) and decline following positive output innovations ($\phi \leq 0$). In particular, when z is sufficiently high ($z_t > z^{**}$), $\beta(z) = 1$, and the agent becomes a full residual claimant on output. We refer to this as the “*founder* region” because the agent’s incentives resemble those of an owner, and the region resembles the benchmark case in which the agent owns the firm. However, this full residual ownership is temporary because total incentives are still volatile ($\phi < 0$). After sufficiently large positive output surprises, z_t can still fall below z^{**} and $\beta(z_t)$ falls below one. In other words, *the agent sells part of the firm after positive output shocks*.

The nature of serial employment creates two dynamic externalities, which we illustrate with comparative statics with respect to ζ and λ in Figure 2. The first externality is that the principal gains from incentives offered by the labor market. In fact, she strictly benefits from stronger incentives: the principal’s flow value in (B-6) is strictly increasing for $z \in [z_L, z_R]$.

²¹Recall that both z_t and β_t are within $[0, 1]$ and z_t is the marginal-utility-weighted average of all future β_t , as shown in Equation (23). The benchmark $\beta_{stop}(z)$ translates these promised future incentives into the current PPS that would keep z_t constant, accounting for separation and resetting.




Figure 2: $v(z)$ with different ζ and λ

The left panel of this figure plots the principal's value function $v(z; \zeta)$ for three different levels of ζ . The right panel plots $v(z; \lambda)$ for different λ . Parameter values are $r = 0.02$, $\rho = 0.5$, $\gamma = 3$, $\sigma = 12$, and $\theta = 0.06$. For the left panel, $\lambda = 0.2$, and ζ takes three values 0.2, 0.5, and 0.8. For the right panel, $\zeta = 0.5$, and λ takes three values 0.1, 0.2, and 0.3.

Conversely, she strictly suffers from having to compensate the agent for PPS offered today ($-\frac{\sigma^2 \gamma r}{2} \beta_t^2$). The principal's highest flow value is from situations in which the agent has high incentives *today* because *future* principals will offer PPS. Today's principal benefits from a labor market that offers the agent a higher share of the value from his expertise.

This externality can be further illustrated by decomposing the value function v into the contribution of explicit incentives, or z within the contract, and implicit incentives, supplied by the prospect of re-joining the market at separation. The latter is captured by z_L in (24), which equals the discounted value of ζ if the remainder of the current contract offers zero PPS (i.e., $\beta = 0$ until the current contract ends). Thus, the former can be written as $z - z_L(\zeta)$. Write $v(z; \zeta)$, $z_L(\zeta)$, and $z_R(\zeta)$ to make their dependence on ζ explicit. We have the following de-composition:

Lemma 1 (Separation of Implicit and Explicit Incentives) *Let $x = z - z_L(\zeta)$ denote the explicit incentives of the current contract and define*

$$x_R \equiv z_R(\zeta) - z_L(\zeta) = \frac{r + \rho}{r + \rho + \lambda}, \quad (31)$$

as its upper bound, which does not depend on ζ . There is a function $\tilde{v} : [0, x_R] \rightarrow \mathbb{R}$ that does not depend on ζ either, such that

$$v(z; \zeta) = \frac{z_R(\zeta)^2}{2\theta(r + \lambda)} + \tilde{v}(x), \quad \forall z \in [z_L(\zeta), z_R(\zeta)]. \quad (32)$$

The proof is in Appendix B.3. Implicit incentives are a level that the current principal takes as given. Holding turnover λ fixed, a higher market price of expertise ζ raises both the

floor z_L and the ceiling z_R by the same amount, $\lambda/(r + \rho + \lambda)$. Consequently, the range of explicit incentives that each contract controls, $x \in [0, x_R]$, and the principal's value of explicit incentives $\tilde{v}(x)$, are the same in every economy. A higher ζ raises the principal's value by the constant $z_R(\zeta)^2/(2\theta(r + \lambda))$ and shifts her policy functions horizontally by $z_L(\zeta)$, as the left panel of Figure 2 illustrates.

The second dynamic externality is that the current principal does not realize all of the cash flows from the expertise that she incentivizes because her interests end at τ , while the agent's expertise generates cash flows past τ . Adding the agent's value of expertise in (26) and the principal's value in (29) yields a total value of

$$\zeta + \frac{(r + \rho)(1 - \zeta)}{r + \lambda + \rho} = \frac{r + \rho + \lambda\zeta}{r + \rho + \lambda} \leq 1. \quad (33)$$

Thus, the sum of the values of expertise to the two participants is weakly below its cash flow value of 1 (strictly below when $\lambda > 0$ and $\zeta < 1$). Serial employment is the source of this externality: if $\lambda = 0$, there would be no such mismatch, and the mismatch increases with λ .

5 Equilibrium

In Sections 3 and 4, we characterized the incentives and payoffs for the current contract once underway. In this section, we describe how contracts are initiated, how surplus is distributed, and what the implications are. The formal equilibrium results are proved in Appendix C.

We are now ready to state our main analytical result: the equilibrium contracts characterized in Definition 2. When the agent is paired with a productive principal, the labor market splits the pair's expected surplus according to the rules in (9) and (10). A fraction κ of that surplus goes to the agent and the rest goes to the principal. Based on the analysis in the previous sections, Theorem 1 summarizes the equilibrium of the labor market and its implications for the resulting optimal contract:

Theorem 1 (Equilibrium Contracts) *There is a unique linear and time-invariant set of equilibrium contracts, up to welfare-equivalent changes to the compensation process (Proposition 3). That equilibrium has*

$$\zeta = \kappa \quad (34)$$

$$\bar{V} = \kappa \frac{r + \lambda}{r} v(\kappa) + (1 - \kappa) \frac{1}{2\theta r} \left(\frac{\lambda\kappa}{r + \lambda + \rho} \right)^2 \quad (35)$$

and all contracts after the first begin with $z_0 = \kappa$. All participants attain weakly positive gains from participation. The surplus from a productive pairing over a pairing with the competitive

fringe is

$$v(\kappa) - \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\kappa}{r + \lambda + \rho} \right)^2 + \frac{r + \rho}{r + \lambda + \rho} H_0, \quad (36)$$

which is nonnegative and strictly positive whenever $\kappa > 0$ or $H_0 > 0$ (Lemma C.3). The first contract (at time $-T$) begins with $z_{-T} = z^* = \arg \max v(z)$ and a one-time payment to the agent of $\kappa(v(z^*) - v(\kappa))$.

Theorem 1 shows that the distribution of surplus in the labor market cannot be separated from the incentives to create expertise. In post-initial contracts, incentive compatibility requires $\zeta = z_0$. The former represents the labor-market value of expertise at initiation (from $R(H, S)$ in (11)), and the latter is the agent's incentive to create expertise at that moment. Any difference between these two could be exploited by lying. For example, if $\zeta > z_0$, then the agent would gain ζ from exaggerating H at signing while only losing z_0 from lower cash flows later. This lie would show up in *both* the current contract and the previous one. First, the agent would lie in the previous contract, expecting a payment based on that lie when moving to a new principal. Second, because both H and lies about H are durable, the lie would continue into the current contract.

The surplus of the principal-agent relationship arises from two sources. The first is the productive use of expertise. As the agent retains more of the value of future pairings (i.e., higher κ), his returns to expertise are higher in the labor market *and* his total dynamic incentives to add to that expertise increase. Thus, $\kappa = \zeta = z_0$ in equilibrium, so the distribution of surplus determines both the labor market's contribution to incentives and each contract's initial incentive level.

The second source of surplus is risk sharing and future growth in expertise, captured in the principal's value $v(\kappa)$. This surplus is split and awarded to the agent in the form of a consumption perpetuity with total value $\bar{V}$, as shown in (26) and (27). Equation (35) decomposes this perpetuity into a component proportional to the discounted value of a productive pairing and a component tied to the lowest incentive state, $z_t = z_L$, associated with the competitive fringe.

6 Implications

In this section we elaborate on the various implications of Theorem 1 including career paths, investment incentives, expertise accumulation, promotion and turnover, and welfare surplus.

6.1 Incentive Regimes and Career Paths

Our model of medium-term contracts enables us to jointly analyze within-employment dynamics and across-employment incentives. Our equilibrium can be summarized by *regimes* based on κ , the fraction of net employment surplus that accrues to the agent. The regions in the principal's problem correspond to regimes in equilibrium pairings. The key addition is that we can now demonstrate the equilibrium relationship between the sharing rule and the incentive structure. Unless stated otherwise, the entry descriptions below concern contracts after the first; the first pair can choose its starting incentives because initial expertise is public.

Because $\kappa = \zeta = z_0$ in equilibrium, the region in which each contract begins depends on where κ lies relative to the thresholds z^* and z^{**} . Lemma 1 directly implies the following:

Proposition 5 (Regime Thresholds) *Define*

$$\underline{\kappa} \equiv \frac{r + \rho + \lambda}{r + \rho}(z^* - z_L), \quad \bar{\kappa} \equiv \frac{r + \rho + \lambda}{r + \rho}(z^{**} - z^L). \quad (37)$$

Then $0 < \underline{\kappa} \leq \bar{\kappa} < 1$, and for $\kappa \in (0, 1)$, every contract after the first begins

1. in the salary region with $\beta_0 = 0$, if $\kappa < \underline{\kappa}$;
2. in the equity region with $z^* < z_0 < z^{**}$ and $\beta_0 > 0$, if $\underline{\kappa} < \kappa < \bar{\kappa}$;
3. in the full-ownership region with $\beta_0 = 1$, if $\kappa > \bar{\kappa}$.

Equivalently, $\kappa < \underline{\kappa}$ if and only if $\kappa < z^*$, and $\kappa > \bar{\kappa}$ if and only if $\kappa > z^{**}$

When the agent receives a relatively small fraction of the surplus, i.e., $\kappa < z^*$, every contract begins in the salary region. We refer to this regime as the “serial salaried-worker regime” because each job starts with a low level of incentives to invest in expertise. In contrast, when the agent receives a relatively large fraction of the surplus, i.e., $\kappa > z^{**}$, every contract begins in the founder region. We call this the “serial founder regime” because the agent begins each new relationship with a temporary full residual claim on output, $\beta_0 = 1$. In the middle, $z^* < \kappa_0 < z^{**}$, the agent begins each relationship with positive equity share. We refer to this as the “serial partner regime” because the agent and the principal shares the residual claims.

The first row of Figure 3 illustrates simulated paths of z_t (blue solid lines) and β_t (red dashed lines) for all three regimes. The second row illustrates their corresponding wage M_t (black solid lines). In the serial salaried worker regime (column 1), $\beta_0 = 0$ at the start of each employment. The agent is compensated through smooth wages alone. The wage is higher when the agent has a larger stock of expertise, resembling the compensation of a

high-paying professional, but the wage is not sensitive to performance. Instead, the agent is promised the possibility of performance incentives (e.g., an equity stake) in the future. The stock of such promises (z_t) *grows smoothly* over time. If the agent survives in the current project long enough, the contract will reach z^* and positive pay-performance sensitivity will begin. In other words, *the agent can be promoted* into a joint-owner role that is compensated with a share of cash flows. Within each contract, incentives are back-loaded at initiation. Yet, the joint-ownership role is temporary: when the project ends, the agent moves to a new principal, and he will be reset to a low-powered contract, receiving only a wage and waiting to be promoted again. This transition delays explicit incentives within each job, while the labor-market component continues to contribute through the reset value; the relative importance of the two depends on λ .

In the serial founder regime. The agent begins with full residual claim on output, $\beta_0 = 1$. If the project lasts long enough, optimal incentives ($\phi_t < 0$) imply that a series of positive output innovations (or simply drift over time) will cause z_t to *decline*. When z_t declines below z^{**} , the agent's claim on output falls below one, $\beta_t < 1$, in the numerical solutions. That is, the agent *sells part of the firm* after sufficiently good performance. The founder becomes a partner. However, every time the project ends, the agent begins again with a temporary full share of output, and each contract is initiated with front-loaded incentives: $\beta_0 > z_0$.

The second row of Figure 3 also decompose the agent's wage into the expertise salary and the risk premium. The former equals $\kappa(r + \rho)H_t$ because under the implementation in Proposition 3, it pays the market price $\kappa = \zeta$ for the expected output flow $(r + \rho)H_t$ of the agent's expertise, independently of the agent's performance. This compensation is on average higher as κ increases because of two effects: the direct effect because the market price of expertise is high, and an indirect price because a higher κ increases the agent's incentives to build expertise, so that his average expertise is higher, as illustrated by the third row of Figure 3. The risk premium directly depends on β_t^2 and is thus also higher when z_t is higher. Therefore, in the serial salaried worker regime, the agent begins each employment with compensation that is low both in levels and in its volatility, and his compensation is generally expected to grow over time. In contrast, in the serial salaried worker regime, the agent begins each employment with compensation that is high both in levels and in its volatility, and his compensation is generally expected to decline over time until the next market reset.

In short, our medium-term optimal contracting equilibrium generates dynamics across incentive regions within each employment relationship, while serial employment generates repeated entry into a particular region across relationships. Expertise and the salary it commands carry over from one job to the next, while the incentive structure resets. The




Figure 3: Simulated Career Paths under the Three Incentive Regimes

Each column shows a simulated stretch of a career after the first job: the serial salaried-worker regime ($\kappa = 0.1 < \underline{\kappa}$, left), the serial partner regime ($\underline{\kappa} < \kappa = 0.5 < \bar{\kappa}$, middle), and the serial founder regime ($\kappa = 0.9 > \bar{\kappa}$, right), where $\underline{\kappa} \approx 0.34$ and $\bar{\kappa} \approx 0.81$ are the thresholds in Corollary 5. Top row: the stock of incentives z_t and PPS β_t . The horizontal lines mark κ , which equals ζ and z_0 in equilibrium, and the thresholds z^* and z^{**} . Investment is proportional to the stock of incentives, $h_t = z_t/\theta$. Middle row: expertise H_t , which is continuous across jobs. Bottom row: the fixed component of pay (the dt term of dM_t in (28)), the expertise salary $\kappa(r + \rho)H_t$, and the risk premium $\frac{\sigma^2 \gamma r}{2} \beta_t^2$. This decomposition of fixed pay is specific to the implementation in Proposition 3, whereas β_t is not. The vertical dashed lines mark separations, after which the agent is re-paired with a new principal. Job lengths are fixed for illustration rather than drawn from the Poisson process, and initial expertise is obtained by simulating 60 years of the model with Poisson separations. Other parameter values are the same as those used in Figure 1.

same contracting environment can therefore generate repeated employment and promotion, repeated partnership, or repeated founding and partial exit, depending on how future surplus from expertise is shared.

6.2 Turnover, Promotion, and Expertise Accumulation

How does the expected duration of a relationship affect the start of a contract and the prospect of promotion? Table 1 reports the cutoffs from Proposition 5 based on λ . Higher λ shrinks the serial salaried-worker range and slightly expands the serial founder range, and increases the entry PPS at the intermediate share $\kappa = .5$.

Table 1: Turnover and Contract Entry

λ	$\underline{\kappa}$	$\bar{\kappa}$	β_0 at $\kappa = .5$	Promotion delay	Promotion probability
.05	.4192	.8192	.175	2.514	.882
.10	.3908	.8150	.224	2.199	.803
.15	.3642	.8117	.265	1.929	.749
.20	.3400	.8092	.299	1.700	.712

Promotion delay and promotion-before-separation probability use $\kappa = .1$. Cutoffs and entry PPS use 1,200 state steps and 2,000 policy steps. $\kappa = .1, .5, .9$ correspond respectively to salary, partner, and founder entry. $r = .02$, $\rho = .5$, $\gamma = 3$, $\sigma = 12$, and $\theta = .06$.

Moreover, if $\kappa < \underline{\kappa}$, the agent begins each employment with $\beta_0 = 0$ and waits to be promoted with equity compensation, $\beta_t > 0$. Since z_t evolves deterministically whenever $\beta_t = 0$, we can explicitly calculate T_p , the time the contract takes from κ to enter the equity region, and

$$T_p(\kappa, \lambda) = \frac{\log(\underline{\kappa}(\lambda)/\kappa)}{r + \rho + \lambda}, \quad 0 < \kappa < \underline{\kappa}(\lambda). \quad (38)$$

An increase in λ shortens the promotion delay T_p because it reduces the promotion cutoff $\underline{\kappa}$. Meanwhile, the probability that the current relationship survives to promotion is $\Pr(\tau > T_p) = e^{-\lambda T_p}$ and it is decreasing in λ . In other words, for agents in the serial salaried-worker regimes, promotion happens sooner but is overall less likely within a given employment when separation becomes more frequent.

The likelihood of turnover also changes the agent's incentives to invest in his expertise. Immediately after a reset, the mean rate of change in expertise is

$$\left. \frac{d}{dt} \mathbb{E}[H_t] \right|_{0+} = \frac{\kappa}{\theta} - \rho \mathbb{E}[H_0].$$

$\kappa/\theta - \rho \mathbb{E}[H]$. Figure 4 starts a new job with expertise drawn from the simulated stationary distribution and conditions on no subsequent separation.

At the low and middle shares, investment is initially too low to replace depreciation: mean expertise declines, then recovers as within-job incentives develop. At $\kappa = .1$, the mean loss reaches 8.160 after 2.2 years when $\lambda = .05$, versus 3.439 after 1.6 years when $\lambda = .20$. The latter economy already has a lower pre-reset stock, so its smaller loss per reset does not imply a higher stationary stock. At $\kappa = .9$, the initial mean drift is positive, but the early increase unwinds with tenure; by year 10, expertise is below its entry mean. Repeated high-state incentive resets replenish investment, rather than making expertise itself jump or producing perpetual within-job growth.




Figure 4: Expertise Adjustment After a Job Change

Mean changes relative to expertise at entry, conditional on no further separation; shaded bands are pointwise 95% intervals from independent trajectory-level changes. Entry expertise comes from year-200 stationary simulations; only z resets to κ , so the stock change at entry is exactly zero. Blue: $\lambda = .05$; red: $\lambda = .20$. Other parameters are as in Figure 5. These profiles are conditional tenure responses, not unconditional career paths.

6.3 The Stationary Distribution of Expertise

For $\lambda > 0$, the stationary Markov implementation used in the numerical analysis makes z_t a regenerative process: it evolves on the compact interval $[z_L, z_R]$ within a contract and resets to $\zeta = \kappa$ at independent Poisson times. Its invariant law is obtained by mixing the no-reset transition law starting from κ over the exponentially distributed age of the current contract. Specifically, the corresponding stationary stock of expertise has the representation

$$H_t = \frac{1}{\theta} \int_{-\infty}^t e^{-\rho(t-s)} z_s ds. \quad (39)$$

Because $z_s \in [z_L, z_R]$, this representation gives the invariant bounds

$$H_t \in \left[\frac{1}{\rho\theta} \frac{\lambda\kappa}{r + \rho + \lambda}, \frac{1}{\rho\theta} \frac{r + \rho + \lambda\kappa}{r + \rho + \lambda} \right]. \quad (40)$$

Both invariant bounds increase with κ , and the interval narrows as λ rises: its lower bound increases and its upper bound decreases. Algebraically, as $\lambda \rightarrow \infty$, it collapses to $\kappa/(\rho\theta)$, subject to the formal-limit qualification in Section 6.5. These are analytical statements about a containing interval, not proofs that the stationary distribution or its mean shifts monotonically.

Figure 5 combines the stationary mean benchmarks in Section ?? with independently simulated medians and 10th–90th percentiles. When λ rises from .05 to .20, the corresponding median changes are 20.289 to 12.236 at $\kappa = .1$, 21.910 to 19.169 at $\kappa = .5$, and 23.652 to 24.721 at $\kappa = .9$. Figure 6 reports the full empirical CDF comparisons. Neither the ordering of selected quantiles nor the larger mean is used to assert stochastic dominance or a welfare improvement.




Figure 5: Mean Expertise and Stationary Quantiles

Blue lines show the stationary mean benchmarks in (??). Dashed lines show medians and shaded regions the 10th–90th percentile ranges, using one year-140 stock observation from each of 1,500 independent trajectories per case, after a 60-year burn-in and an 80-year observation window. Expertise is accumulated on the live .005-year time grid and remains continuous at exact Poisson reset times. Parameters are as in Figure 7, except for the displayed separation intensities. Quantiles are numerical estimates, not analytical support bounds.

A focused robustness check lowers risk aversion from $\gamma = 3$ to $\gamma = 2$, retaining the other primitives and only the two endpoint turnover rates. Table 2 shows that the directions of all three mean responses survive. The new salary/partner and partner/founder cutoffs are approximately (.482, .843) at $\lambda = .05$ and (.404, .833) at $\lambda = .20$; thus the three representative shares retain their entry regimes, although the thresholds change. The low- and middle-share




Figure 6: Turnover and the Stationary Distribution of Expertise

Blue: $\lambda = .05$; red: $\lambda = .20$. Each CDF uses one stock observation per independent trajectory, with the simulation design and primitives in Figure 5. The distribution-free 95% uniform CDF half-width is .0351 for each curve separately; bands are omitted for readability. These finite-sample comparisons do not establish stochastic dominance.

decline/recovery profiles and the high-share initial growth followed by reversal also survive. This is one parameter check, not a claim of universal comparative-statics signs.

κ	Entry regime	$\gamma = 3$		$\gamma = 2$	
		$\lambda = .05$	$\lambda = .20$	$\lambda = .05$	$\lambda = .20$
.1	Salary	19.260	12.246	20.850	13.263
.5	Partner	22.210	19.416	24.019	20.845
.9	Founder	24.022	24.966	25.785	26.285

Table 2: Stationary Mean Expertise: A Risk-Aversion Check

Means use (??) and the exponential-age resolvent of the no-reset policy generator. The $\gamma = 2$ exercise requires [one policy solve per turnover rate \(Lemma 1\)](#) plus grid validation, and six equilibrium stock-simulation cases. All other primitives are unchanged; entry regimes are classified using each parameterization's own cutoffs.

Numerical checks. The stock simulations use exact depreciation with linearly interpolated investment within each numerical segment, including separate segments on either side of every Poisson reset. Their maximum discrepancy from the stationary mean benchmark is .0336, within the trajectory-based uncertainty check. Paired .005/.01-year steps change mean expertise by at most .00183 and the reported quantiles by .0763; representative doubled state/policy grids change means by .00636 and quantiles by .0118. Extending the burn-in from 60 to 120 years, with an 80-year observation window in each case, passes the trajectory-based stability checks. Different initial stocks have an exactly exponentially decaying effect, bounded by 2.85×10^{-12} after the main burn-in when initialized anywhere in the invariant

interval. No stock-bound violations or incentive-state clamping occur. These checks support the reported numerical findings, but are not continuous-problem error bounds.

6.4 Surplus

Surprisingly, the agent's lifetime value is *not* always maximized in economies in which he receives all of the surplus from projects. Similarly, the principal's value is *not* always maximized when she receives all of the surplus:

Proposition 6 (Surplus and Welfare) *For any contract after the first one, fix $H_0 > 0$ and S_0 , and evaluate V_0^A , V_0^P , and total surplus (36) at $z_0 = \zeta = \kappa$. If $\sigma^2\gamma r\theta$ is sufficiently high relative to H_0 ,²² then V_0^A , V_0^P , and total surplus (36) are lower at $\kappa = 1$ than for κ sufficiently close to 1 from below. For any parameter values, V_0^A and total surplus (36) are higher for sufficiently small $\kappa > 0$ than at $\kappa = 0$. If, in addition, $H_0 < 1/(\theta(r + \lambda))$, then V_0^P is also higher for sufficiently small $\kappa > 0$ than at $\kappa = 0$.*

We illustrate the results in Figure 7 by plotting $v(\cdot)$ and $\bar{V}$, which are equilibrium components of the agent's welfare (26), the principal's welfare (29), and total net surplus (36).




Figure 7: $v(\kappa; \kappa)$, $v(z; \kappa)$, and $\bar{V}$

Left: $v(\kappa; \kappa)$ for $\kappa \in [0, 1]$; dotted curves reproduce $v(z; \kappa)$ from Figure 2, while the solid curve connects the points $v(z = \kappa; \kappa)$, showing the incentive component of pair value in each economy. Right: $\bar{V}$, the constant component of optimal compensation from (35). Parameters are $r = 0.02$, $\rho = 0.5$, $\gamma = 3$, $\sigma = 12$, $\theta = 0.06$, and $\lambda = 0.2$.

Why might the agent prefer an economy in which he receives less than all of the surplus? Giving the agent a full share of the value of expertise implies very high incentives, which are expensive because they require exposing the risk-averse agent to volatile output. When $\kappa = \zeta = z_0 = 1$, Equation (23) implies that the agent receives a full claim on output

²²Such values of $\sigma^2\gamma r\theta$ are consistent with Assumption 2 precisely when $\theta r H_0 < ((r + \rho)^2 + (r + \rho)\lambda - \lambda^2) / (r + \lambda + \rho)^2$. The condition holds, for example, under the parameter values of Figure 7 for any H_0 below approximately 538.

for his entire life ($\beta_t = 1$). However, this is the right boundary of the principal's concave value function $v(z)$, an over-incentivized state with value below the interior maximum and especially costly when $\sigma^2\gamma r\theta$ is high enough. This can generate a similar low point in the total pair surplus (36). What the agent would like to do is take the entire value of expertise in the labor market ($\kappa = 1$) but enter into contracts that only provide moderate incentives ($z_0 = z^*$). But this is impossible precisely because the incentives to create expertise cannot be separated from the distribution of the surplus from expertise.

Similarly, giving the principal a full share of expertise implies that the agent has zero incentive to develop expertise in contracts after the first ($\kappa = 0$ implies $z_0 = z_t = 0$). In these contracts, neither the current principal nor future principals leave the agent a marginal return to expertise; this does not rule out investment in the first contract, whose starting incentives can be chosen freely. If the agent's stock of expertise is low, the lack of incentives implies little surplus for the principal to extract. What the principal would like to do is take the entire value of expertise today, but enter into contracts with moderate incentives and the knowledge that *future* principals would not extract the full value. However, this is impossible precisely because the incentives to create expertise cannot be separated from the distribution of the surplus from expertise, including across principals.

6.5 Unique Properties of Medium-Term Contracting

To further demonstrate the value of explore medium-term contracting, this section compares its properties with two benchmarks, long-term contract without termination ($\lambda = 0$), and serial spot contracting ($\lambda \rightarrow \infty$).

Everything in Sections 3 and 4 remains valid with $\lambda = 0$. Hidden, persistent expertise still generates time-varying incentive pay and an incentive stock that rises after bad performance. The critical distinction is that when $\lambda = 0$, there are no implicit incentives, no condition that links incentives and rents, and no reset-induced salary phase at entry. Instead, the agent starts at the pair-value-maximizing level of $z_0 = z^*$, the boundary between the salary and equity regions, and his career evolves with time-varying performance pay. Medium-term contracting changes where each job starts and repeatedly returns the agent to that starting point; it does not create the entire within-contract incentive mechanism.

The entire value of the long-term contract could be given to the agent or kept by the principal through a consumption perpetuity. Under CARA, such transfers redistribute rents without changing the optimal incentive policy. With serial contracting, however, the surplus share also fixes starting incentives: $z_0 = \kappa$ after the first job. Thus, the distribution of surplus selects the entry regime and affects subsequent investment and expertise accumulation. The same restriction underlies the entry loss in Section 6.6 and the welfare comparisons in

Section 6.4.

The resulting career mechanisms are therefore tied to repeated entry, rather than to a different technology for building expertise. Low-share workers repeatedly restart in the salary region and wait for promotion; intermediate-share workers restart with a partial performance claim in the numerical solutions; high-share workers repeatedly begin with full PPS and can subsequently sell part of that claim. Section 6.2 quantifies entry and promotion, while Section ?? shows that mean PPS differences across starting regimes fade with tenure in the simulations. These within-job profiles do not remove the consequences of repeated resets. Section ?? separates contract adjustment from reset frequency: across the displayed rates, higher turnover lowers career-average investment at low and middle shares but raises it at the high share. Sections ?? and 6.3 show how these incentives translate into gradual stock adjustment and stationary expertise. Expertise itself remains continuous at a job change, and greater turnover is not uniformly beneficial or harmful for its accumulation in these examples.

6.5.1 The Obstruction to Long-Term Replication

To see the obstruction to carrying a long-term incentive process across successive principals, consider what it would take to keep z_t continuous as the agent moves from one principal to another. Under the entry condition, each agent's surplus share and initial stock of incentives would have to equal the inherited z_t and therefore respond to *past realizations of the idiosyncratic shocks* Z . In other words, there could not be a common return to expertise. This explains why common market pricing prevents such continuity; it is not a construction establishing welfare equivalence. Nor does separation mechanically destroy expertise or first-best surplus: the shocks reassign production across principals. We therefore do not infer a lifetime-welfare ranking between the long-term and medium-term economies from this comparison.

6.5.2 The Formal Spot Limit

The serial spot contracting limit, $\lambda \rightarrow \infty$, is instead degenerate. The lower and upper bounds on aggregate incentives, z_L and z_R in (24) and (25), both converge to $\zeta = \kappa$. In other words, if the agent is an instant away from pairing with another principal, the incentives in the current contract cannot differ from those offered by the future principal. The bounds imply a formal limit of $h = \kappa/\theta$, with vanishing within-job variation in z . At fixed positive risk primitives, this limit eventually violates Assumption 2: its first right-hand side tends to zero while $\sigma^2\gamma r\theta$ remains positive. It is therefore a comparison of algebraic bounds, not an

equilibrium-convergence result under unchanged maintained assumptions.²³

6.5.3 Outside Options and Contract Distortions

These results sharpen the distinction between two ways in which future employment affects a current contract. In [Guerrieri \(2008\)](#), contracts offered by other firms change incumbent workers' outside options and hence the distortions associated with private information. In [Chemla, Rivera, and Shi \(2025\)](#), attractive outside options weaken the threat of termination, leading to excessive compensation, insufficient deferral, and excessive tenure relative to shareholder-value-maximizing contracts. Here separation is exogenous, so termination is not an incentive instrument. Instead, future contracts price a privately accumulated productive stock. Lifetime incentive compatibility ties that price to the incentive stock at entry, $z_0 = \kappa$. The resulting restriction selects the entry regime and can put investment below or above the current pair's free-entry optimum (Section 6.6); it is not a result about excessive compensation or inefficient retention.

Future employment can also substitute for current performance pay without an accumulated productive stock. In [Häfner and Taylor \(2023\)](#), employers reward successful production with references that improve subsequent hiring prospects. In their firm-optimal equilibrium, references support higher effort with lower bonuses. Our implicit incentives instead reward investment through the earnings generated by portable expertise in later jobs; they do not operate through preferential hiring. Moreover, a fixed market price of expertise implies a constant implicit component within a job, $z_L = \lambda\kappa/(r + \rho + \lambda)$. Thus, the time variation in investment incentives comes from the current contract's explicit component, even though future contracts determine its starting point (Section ??).

The investment effect of mobility likewise depends on who invests and how future returns are appropriated. In [Shi \(2023\)](#), firms invest in general human capital, and noncompete buyouts let them recover returns from future employers at the cost of restricting worker reallocation. Here the worker privately chooses investment, expertise is fully portable, and separation neither destroys that stock nor is chosen through a mobility restriction. Turnover changes both the weight on future earnings and the frequency of incentive resets. The numerical results in Section ?? show that more frequent resets can reduce career-average investment at low surplus shares but raise it at high shares. These are incentive-timing effects, not a policy ranking of mobility restrictions.

Finally, the partnership interpretation concerns compensation rather than control rights. In [Picariello \(2026\)](#), portability can induce owners who cannot commit to task assignments to sacrifice productivity to reduce retention costs; sharing control and profits through an

²³Inducing welfare to converge in expectation, rather than pointwise, requires additional assumptions about the evolution of z_t as λ grows.

equity partnership can restore efficient task allocation. Our serial partner regime does not rely on task choice or a change in organizational form. It describes a partial performance claim that adjusts with output realizations and resets at the next job. Likewise, entry in the founder regime gives a temporary full residual claim, not a distinct production technology. Medium-term contracting therefore explains transitions among compensation arrangements within a fixed technology, rather than the optimal legal organization of the firm.

6.6 The Cost of the Entry-Incentive Restriction

Are starting incentives too weak or too strong for the current pair? The relevant benchmark holds future labor-market opportunities fixed at $\zeta = \kappa$ and lets the pair choose its initial incentive state. This is the problem in the first contract, when initial expertise is public: the pair chooses $z^*(\kappa) = \arg \max_z v(z; \kappa)$. Later contracts instead must start at $z_0 = \kappa$. They are already optimal conditional on that restriction. The comparison below measures the cost of the entry restriction relative to the free-entry benchmark, not a gain that a later pair could obtain by renegotiating while retaining all the lifetime incentive-compatibility constraints.

Corollary 1 (Cost of the Entry-Incentive Restriction) *Fix $\lambda > 0$ and the other primitives, and suppose Assumption 2 holds for all $\kappa \in [0, 1]$. Write $x_R = (r + \rho)/(r + \rho + \lambda)$, let x^* maximize the principal's value of explicit incentives $\tilde{v}$ in Lemma 1, and recall $\underline{\kappa} = x^*/x_R$. Then*

$$\kappa - z^*(\kappa) = x_R(\kappa - \underline{\kappa}), \quad (41)$$

$$D(\kappa) \equiv v(z^*(\kappa); \kappa) - v(\kappa; \kappa) = \tilde{v}(x^*) - \tilde{v}(x_R \kappa) \geq 0. \quad (42)$$

The entry-restriction loss D is continuous and strictly convex in κ , with its unique zero and minimum at $\underline{\kappa}$. It strictly decreases below that share and strictly increases above it. Relative to the free-entry benchmark, initial investment is too low when $\kappa < \underline{\kappa}$ and too high when $\kappa > \underline{\kappa}$, with investment gap $x_R(\kappa - \underline{\kappa})/\theta$. The loss and investment gap do not depend on initial expertise or savings.

The proof in Appendix C.5 uses only strict concavity and Lemma 1; it does not require smoothness of the value function. One surplus share performs two tasks: it allocates surplus to the agent and fixes the incentives with which the next job begins. A low share leaves investment incentives below the level the current pair would prefer; a high share requires stronger incentives than the pair finds worthwhile, given the costs of investment and risk exposure.

At the baseline parameters, $\underline{\kappa} \approx .34$, so entry at $\kappa = .1$ is under-incentivized and entry at .5 or .9 is over-incentivized relative to each pair's own free-entry benchmark. In particular,

entry can be over-incentivized in the partner regime, not only under full ownership. These are comparisons of total investment incentives and initial investment, not claims that current PPS is too high or too low, or that incentives remain excessive or insufficient throughout the job.

Theorem 1 already identifies the free-entry gain in the first contract: the agent’s signing payment is $\kappa D(\kappa)$, while the principal retains $(1 - \kappa)D(\kappa)$ of the additional surplus. The gap D itself is therefore not the signing payment. Nor does its minimum identify the welfare-maximizing surplus share: changing κ also changes the market value of expertise and other components of equilibrium surplus. Section 6.4 concerns these equilibrium welfare comparisons, rather than the conditional entry comparison.

7 Conclusion

A defining feature of career paths is that they combine two trajectories. Workers often spend substantial time with a single employer, developing skills and accumulating expertise, and may eventually move to another employer or launch a new venture, carrying that expertise with them. Employment relationships are therefore neither purely long-term attachments to a single principal nor purely spot-market transactions. They are medium-term arrangements in which incentives to accumulate expertise arise both from the current employment relationship and from the anticipated value of expertise in future employment. This structure creates a dynamic externality: current principals provide incentives for investment in expertise that they may not fully capture, while future principals benefit from expertise created by earlier contracts and, through expected future rewards, help shape incentives today.

This paper studies medium-term optimal contracting with persistent private information and stochastic transitions that periodically reset contract terms. The equilibrium features dynamic incentives for investment and revelation within each contract (the explicit incentives) as well as incentives from anticipated labor market outcomes across contracts (the implicit incentives). Because the model concerns the revelation of hidden, endogenously accumulated expertise rather than learning about or manipulating beliefs about a fixed type, misreports are durable and persistently distort reported cash-flow innovations. Because we model medium-term contracts, the agent’s total incentives (total value of H) differ from both the agent’s implicit incentives (labor market value of H) and explicit incentives (contracting value of H). The combination of these features generates endogenous regions with distinct incentive dynamics within each contract as well as distinct career regimes resembling serial salaried employment and serial entrepreneurship. The equilibrium link between the distribution of surplus and incentives also has welfare consequences: the agent’s welfare is not always maximized in economies in which he captures all of the surplus from his expertise,

because receiving the full surplus implies maximal and potentially very costly incentives. Our numerical results further show that greater labor market fluidity promotes the aggregate accumulation of expertise only when workers expect to capture a sufficiently large share of the surplus that their expertise creates.

Our framework can be extended in several directions. One possibility is to explore stationary but nonlinear equilibria, meaning that $\zeta(H)$ is not constant.²⁴ Linearity provides tractability but is not essential to the central mechanism: with a nonlinear labor-market value of expertise, initial incentives would vary with the agent’s state but remain linked to that value through incentive consistency. Such an extension may produce richer implications, but we expect our main findings regarding the interaction between explicit and implicit incentives to remain robust. Another natural extension is to consider a more elaborate labor market with search frictions, endogenous matching, or multiple simultaneous principals. We leave these topics for future research.

²⁴For example, one might want to allow the agent’s share of surplus to vary with his level of expertise. Our methods appear able to accommodate this approach, but at the cost of substantially greater computational complexity. We believe the agent’s problem would be essentially similar, but the principal’s valuation would become a PDE rather than an ODE, as z_t and H_t cannot be separated. Finding an equilibrium, or proving its existence, would require solving a fixed-point problem within a system of related PDEs instead of simply matching linear coefficients.

References

- Aghion, P. and J. Tirole (1997). Formal and real authority in organizations. *Journal of Political Economy* 105(1), 1–29.
- Bergemann, D. and P. Strack (2015). Dynamic revenue maximization: A continuous time approach. *Journal of Economic Theory* 159, 819–853.
- Bloedel, A. W., R. V. Krishna, and O. Leukhina (2025). Insurance and inequality with persistent private information. *Econometrica* 93(3), 821–857.
- Bloedel, A. W., R. V. Krishna, and B. H. Strulovici (2023). Persistent private information revisited. Working paper.
- Board, S. and M. Meyer-ter-Vehn (2015). Relational contracts in competitive labor markets. *Review of Economic Studies* 82, 490–534.
- Bonatti, A. and J. Hörner (2017). Career concerns with exponential learning. *Theoretical Economics* 12(1), 425–475.
- Chemla, G., A. Rivera, and L. Shi (2025). Too much, too soon, for too long: The dynamics of competitive executive compensation. *Journal of Finance* 80(5), 2921–2970.
- Cisternas, G. (2018a). Career concerns and the nature of skills. *American Economic Journal: Microeconomics* 10(2), 152–189.
- Cisternas, G. (2018b). Two-sided learning and the ratchet principle. *Review of Economic Studies* 85(1 (302)), 307–351.
- Crandall, M. G., H. Ishii, and P.-L. Lions (1992). Users guide to viscosity solutions of second order partial differential equations. *Bulletin of the American mathematical society* 27(1), 1–67.
- Di Tella, S. and Y. Sannikov (2021). Optimal asset management contracts with hidden savings. *Econometrica* 89(3), 1099–1139.
- Feng, F. Z., R. Y. Luo, and M. M. Westerfield (2024). The optimal duration of incentives and cash flows. Working Paper.
- Fernandes, A. and C. J. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory* 91(2), 223–247.
- Fleming, W. H. and H. M. Soner (2006). *Controlled Markov processes and viscosity solutions*, Volume 25. Springer Science & Business Media.
- Fu, S. and R. V. Krishna (2019). Dynamic financial contracting with persistent private information. *RAND Journal of Economics* 50(2), 418–452.
- Fuchs, W. and L. Garicano (2010). Matching problems with expertise in firms and markets. *Journal of the European Economic Association* 8(2–3), 354–364.

- Fuchs, W., L. Garicano, and L. Rayo (2015). Optimal contracting and the organization of knowledge. *Review of Economic Studies* 82(2), 632–658.
- Fudenberg, D. and L. Rayo (2019). Training and effort dynamics in apprenticeship. *American Economic Review* 109(11), 3780–3812.
- Gibbons, R. and K. J. Murphy (1992). Optimal incentive contracts in the presence of career concerns: Theory and evidence. *Journal of Political Economy* 100(3), 468–505.
- Golan, L. (2005). Counteroffers and efficiency in labor markets with asymmetric information. *Journal of Labor Economics* 23(2), 373–393.
- Guerrieri, V. (2008). Inefficient unemployment dynamics under asymmetric information. *Journal of Political Economy* 116(4), 667–708.
- Haadem, S., B. Øksendal, and F. Proske (2013). Maximum principles for jump diffusion processes with infinite horizon. *Automatica* 49(7), 2267–2275.
- Häfner, S. and C. R. Taylor (2023). Working for references. *American Economic Journal: Microeconomics* 15(3), 33–77.
- He, Z., B. Wei, J. Yu, and F. Gao (2017). Optimal long-term contracting with learning. *Review of Financial Studies* 30(6), 2006–2065.
- Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *Review of Economic studies* 66(1), 169–182.
- Hörner, J. and N. S. Lambert (2021). Motivational ratings. *Review of Economic Studies* 88(4), 1892–1935.
- Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory* 52(1), 45–67.
- Kaniel, R. and D. Orlov (2025). Intermediated asymmetric information, compensation, and career prospects. *American Economic Review* 115(10), 3638–3674.
- 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.
- Ke, T. T., C. Li, and M. Safronov (2023). Learning by choosing: Career concerns with observable actions. *American Economic Journal: Microeconomics* 15(2), 536–567.
- 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.
- Kwon, S. (2016). Relational contracts in a persistent environment. *Economic Theory* 61, 183–205.

- Lafontaine, F. and K. Shaw (2016). Serial entrepreneurship: Learning by doing? *Journal of Labor Economics* 34(S2), S217–S254.
- Maital, S. and E. Barzani (2021). Why startups fail: a survey of empirical studies. *Samuel Neaman Institute for National Policy Research Technion-Israel Institute of Technology*.
- Malcomson, J. M. (2016). Relational incentive contracts with persistent private information. *Econometrica* 84(1), 317–346.
- Marinovic, I. and F. Varas (2019). CEO Horizon, Optimal Pay Duration, and the Escalation of ShortTermism. *Journal of Finance* 74, 2011 – 2053.
- Martimort, D., A. Semenov, and L. Stole (2017). A theory of contracts with limited enforcement. *Review of Economic Studies* 84(2), 816–852.
- Pagano, M. and A. Röell (1998). The choice of stock ownership structure: Agency costs, monitoring, and the decision to go public. *Quarterly Journal of Economics* 113(1), 187–225.
- Pavan, A. and G. Calzolari (2009). Sequential contracting with multiple principals. *Journal of Economic Theory* 144, 503–531.
- Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A myersonian approach. *Econometrica* 82(2), 601–653.
- Pegoraro, S. (2020). Flows and performance with optimal money management contracts. In *Essays on Financial Intermediaries and Monetary Policy*. University of Chicago.
- Peters, M. (2001). Common agency and the revelation principle. *Econometrica* 69, 1349–1372.
- Picariello, L. (2026). Organizational design with portable skills. *Journal of Economics & Management Strategy* 35(1), 196–216.
- Sannikov, Y. (2014). Moral hazard and long-run incentives. Working paper.
- Shaw, K. and A. Sørensen (2019). The productivity advantage of serial entrepreneurs. *ILR Review* 72(5), 1225–1261.
- Shi, L. (2023). Optimal regulation of noncompete contracts. *Econometrica* 91(2), 425–463.
- Strulovici, B. (2024). Renegotiation-proof contracts with persistent states. Working Paper.
- Sun, D., W. Gao, H. Hu, and S. Zhou (2022). Why 90% of clinical drug development fails and how to improve it? *Acta Pharmaceutica Sinica B* 12(7), 3049–3062.
- Watson, J., D. A. Miller, and T. E. Olsen (2020). Relational contracting, negotiation, and external enforcement. *American Economic Review* 110(7), 2153–2197.
- 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.
- Zabojnik, J. and D. Bernhardt (2001). Corporate tournaments, human capital acquisition, and the firm size-wage relation. *Review of Economic Studies* 68(3), 693–716.
- Zhang, Y. (2009). Dynamic Contracting with Persistent Shocks. *Journal of Economic Theory* 144(2), 635–675.

Appendix

A The Agent's Problem

We will prove Propositions 2 and 3 as well as some preliminaries for Lemma C.1. Building on the notation in Section 2.2, we use Δ to denote the difference between the agent's choices and the principal's recommendations. For example, $\Delta c = c - \hat{c}$ and $\Delta S = S - \hat{S}$. We denote the Radon-Nikodym derivative between the probability measures induced by the principal's and agent's information sets by ξ_t , so that $E_t[(\cdot)_s] = \hat{E}_t[\frac{\xi_s}{\xi_t}(\cdot)_s]$ for $s \geq t$.

In addition, we use χ_t to indicate the productivity of the principal-agent pair that exists at time t : $\chi_t = 1$ if the project exploits expertise, and $\chi_t = 0$ if the pairing is with the competitive fringe.

Because the recommended actions and the agent's deviations must be consistent with the feasible technologies, we can take differences in (1) and (3) to obtain

$$d\Delta H_t = (\Delta h_t - \rho \Delta H_t) dt \quad (\text{A-1})$$

$$d\Delta S_t = (r\Delta S_t - \Delta c_t) dt + \Delta \nu_t (dN_t - \lambda dt) \quad (\text{A-2})$$

Because output is observable, misreporting expertise or destroying or augmenting output imply an offsetting misreport of the path of the Brownian motion. Thus, dY_t is equal to both $(\chi_t(r + \rho)H_t)dt + \sigma dZ_t$ and $(\chi_t(r + \rho)\hat{H}_t)dt + \sigma d\hat{Z}_t$, and we have

$$(\chi_t(r + \rho)\Delta H_t) dt = -\sigma d\Delta Z_t. \quad (\text{A-3})$$

Thus, the Radon-Nikodym derivative ξ_t between the probability measures induced by Z and $\hat{Z}$ satisfies

$$\xi_t = \xi_0 \exp \left(-\frac{1}{2} \int_0^t \left(\frac{\chi_s(r + \rho)\Delta H_s}{\sigma} \right)^2 ds + \int_0^t \frac{\chi_s(r + \rho)\Delta H_s}{\sigma} d\hat{Z}_s \right) \quad (\text{A-4})$$

$$d\xi_t = \left(\frac{\chi_t(r + \rho)\Delta H_t}{\sigma} \right) \xi_t d\hat{Z}_t. \quad (\text{A-5})$$

All processes satisfy the standard measurability and integrability conditions and are right-continuous with left limits.

A.1 Required Incentives for Any Sequence of Contracts

In this subsection, we derive necessary conditions for incentive compatibility, as in Haadem, Øksendal, and Proske (2013). The agent's goal is to maximize lifetime expected utility, given by (5) and (6), so we write the objective function as

$$E_t \left[\int_t^\infty e^{-r(s-t)} u(c_s, h_s) ds \right] = \hat{E}_t \left[\int_t^\infty e^{-r(s-t)} \frac{\xi_s}{\xi_t} u(\hat{c}_s + \Delta c_s, \hat{h}_s + \Delta h_s) ds \right], \quad (\text{A-6})$$

where we have used the change of measure ξ in (A-5) to go from E_t to $\hat{E}_t$ and, by definition, $c_s = \hat{c}_s + \Delta c_s$. Then, the agent maximizes (A-6) subject to (A-1), (A-2), and (A-5).

We set up the Hamiltonian as

$$\begin{aligned}\mathcal{H}(\Delta c, \Delta h, \xi, \Delta H, \Delta S, p, q) = & \xi u(\hat{c} + \Delta c, \hat{h} + \Delta h) + q^\xi \left(\frac{\chi(r + \rho)\Delta H}{\sigma} \right) \xi \\ & + p^H(\Delta h - \rho\Delta H) + p^S(r\Delta S - \Delta c) + j^S\lambda\Delta\nu\end{aligned}\quad (\text{A-7})$$

Taking first-order conditions and evaluating along the optimal path ($\Delta c^* = \Delta h^* = \Delta\nu = 0$, $\xi^* = 1$, and $\Delta H^* = \Delta S^* = 0$) yields

$$[\Delta c] \quad -\gamma u(\hat{c}, \hat{h}) = p^S \quad (\text{A-8})$$

$$[\Delta h] \quad \theta\gamma\hat{h}u(\hat{c}, \hat{h}) = -p^H. \quad (\text{A-9})$$

Similarly, optimality for $\Delta\nu$ requires $j^S = 0$.

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

$$dp_t^\xi = \left(rp_t^\xi - u(\hat{c}_t, \hat{h}_t) \right) dt + q_t^\xi d\hat{Z}_t + j_t^\xi (dN_t - \lambda dt) \quad (\text{A-10})$$

$$dp_t^H = \left((r + \rho)p_t^H - \chi_t(r + \rho)\frac{q_t^\xi}{\sigma} \right) dt + q_t^H d\hat{Z}_t + j_t^H (dN_t - \lambda dt) \quad (\text{A-11})$$

$$dp_t^S = q_t^S d\hat{Z}_t, \quad (\text{A-12})$$

where dp_t^S has no jump term because the first-order conditions require $j_t^S = 0$.

Integrating (A-10) yields

$$p_t^\xi = \hat{\mathbb{E}}_t \left[\int_t^\infty e^{-r(s-t)} u(\hat{c}_s, \hat{h}_s) ds \right], \quad (\text{A-13})$$

which is the continuation utility W_t for incentive-compatible contracts.

Equation (A-12) indicates that p^S is a local martingale. Equation (A-8) then implies that the flow utility $u(\hat{c}, \hat{h})$ is also a local martingale. Finally, (A-13) implies that p_t^ξ is a local martingale, so (A-10) gives $rp_t^\xi = u(\hat{c}_t, \hat{h}_t)$. Then, from (A-8), $p_t^S = -\gamma r p_t^\xi$, so $j_t^S = 0$ implies $j_t^\xi = 0$. Writing p_t^ξ at the optimum as W_t yields (18) and $p_t^S = -\gamma r W_t$. Combining that result with (A-9) and defining

$$z_t = \frac{p_t^H}{-\gamma r W_t} \quad (\text{A-14})$$

yields (17). Redefining $q_t^\xi = -\sigma\gamma r W_t \beta_t$, we see that $\frac{q_t^\xi}{\sigma} \in [0, p_t^S]$ is equivalent to $0 \leq \beta_t \leq 1$, and we recover the law of motion (16). With β_t bounded, W_t is a true martingale rather than merely a local martingale.

Integrating (A-11), we have

$$p_t^H = \hat{E}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} \chi_s(r+\rho) \frac{q_s^\xi}{\sigma} ds \right] = \hat{E}_t \left[\int_t^\infty e^{-(r+\rho)(s-t)} (-\chi_s(r+\rho) \gamma r W_s \beta_s) ds \right], \quad (\text{A-15})$$

which is the discounted value of future incentives. Since W_t is a martingale and $0 \leq \beta_t \leq 1$, we have $p_t^H \in [0, -\gamma r W_t]$ and $z_t \in [0, 1]$.

A.2 A Sequence of Productive-Pairing Contracts

The results of Appendix A.1 can be specialized to a sequence of productive projects by setting $\chi_s = 1$ and using (11) as the agent moves from one contract to the next. Thus, at any contract separation, we have $W_\tau = R(H_\tau, S_\tau)$ and the Hamiltonian system implies $p_\tau^H = R_H(H_\tau, S_\tau)$. At separation, the post-separation marginal value is

$$p_\tau^H = R_H(H_\tau, S_\tau) = -\gamma r \zeta W_\tau. \quad (\text{A-16})$$

The jump from the pre-separation costate p_t^H to this post-separation value is therefore

$$j_t^H = R_H(H_t, S_t) - p_t^H = -\gamma r \zeta W_t - p_t^H. \quad (\text{A-17})$$

Substituting z_t for p_t^H using (A-14), we have $j_t^H = -\gamma r W_t (\zeta - z_t)$. Applying Itô's lemma to the definition of z_t in (A-14), using (16) for dW_t and (A-11) for dp_t^H , yields (19) after relabeling the volatility as $(r+\rho)\phi_t\sigma$.

Starting with (A-15), applying the value of the jump at $t = \tau$ from (A-17), and dividing through by $-\gamma r W_t$ yields (22). Then, using $0 \leq \beta_t \leq 1$ and the fact that W_t is a martingale, direct integration yields $\frac{\lambda \zeta}{r+\rho+\lambda} \leq z_t \leq \frac{r+\rho+\lambda \zeta}{r+\rho+\lambda}$.

A.3 Proof of Proposition 2

Appendices A.1 and A.2 show the necessary conditions for incentive compatibility. Online Appendix OA.2 shows that if $\phi_t \leq 0$, then W_t is the best the agent can do starting from $\Delta H_t = \Delta S_t = 0$. Since W_t is attained under these necessary conditions, and the sequence of contracts begins with $\Delta H_{-T} = \Delta S_{-T} = 0$, the conditions are also sufficient.

A.4 Proof of Proposition 3

Because all future contracts are expected to be productive pairings regardless of the current pairing, $W_\tau = R(H_\tau, S_\tau)$. Because W_t and H_t do not jump, there is a process for S_t that does not jump and satisfies $W_t = R(H_t, S_t)$. Moreover, because the principal and agent share a discount rate, any set of jump-diffusions with a discounted value of zero over contracts can be added to S_t without affecting welfare. We can therefore choose $S_0 = S_{0-}$ and a process for S_t that is continuous on $[0, \tau]$. This choice implies that there are no discrete transfers at contract initiation for productive pairings.

Solving for the agent's certainty-equivalent wealth in (7) using $W_t = R(H_t, S_t)$ and the form of R in (11) yields (26). Combining the recommended actions $u(\hat{c}_t, \hat{h}_t) = r W_t$ in (18)

and $W_t = R(H_t, S_t)$ with the form of R in (11), and solving for $\hat{c}_t$ yields (27).

Using $W_t = R(H_t, S_t)$ and the form R in (11), we have

$$S_t = -\zeta H_t - \frac{1}{\gamma r} \ln(-\gamma r W_t) - \bar{V}. \quad (\text{A-18})$$

Taking the differential form and using (16) for dW_t , we have

$$dS_t = -\zeta dH_t + \frac{\sigma^2 \gamma r}{2} \beta_t^2 dt + \sigma \beta_t d\hat{Z}_t. \quad (\text{A-19})$$

Substituting (A-18), (A-19), (1), and the values of consumption in (27) and certainty-equivalent wealth in (26) into the law of motion for S_t in (3) and solving for dM_t yields (28).

B The Principal's Problem

We solve the principal's problem using the output process $dY_t = \chi_t(r + \rho)H_t dt + \sigma dZ_t$, where $\chi_t = 1$ indicates a productive project and $\chi_t = 0$ indicates a pairing with the competitive fringe (following the approach of Appendix A). The calculations for both productive and fringe pairings assume that all remaining contracts in the sequence are from productive pairings. These results derive the value-function representation and boundary conditions in Proposition 4 and provide preliminaries for Lemma C.2. The Online Appendix proves the remaining claims in the proposition.

B.1 The Principal's Value Function

Substituting $dY_t = \chi(r + \rho)H_t dt + \sigma dZ_t$ and the consumption and compensation processes in (27) and (28) into the principal's objective function (8), we see that the principal maximizes

$$\mathbb{E}_t \left[\int_t^\tau e^{-r(s-t)} \left((r + \rho)(\chi - \zeta)H_s + \frac{\zeta}{\theta} z_s - \frac{1}{2\theta} z_s^2 - \frac{\sigma^2 \gamma r}{2} \beta_s^2 - r\bar{V} \right) ds \right], \quad (\text{B-1})$$

subject to (1), to (19) when $\chi = 1$, and to constant z_t when $\chi = 0$, as established in Lemma C.1.

Next, using $h_t = \frac{z_t}{\theta}$ from (17) and integrating by parts with respect to H_t , we have

$$\mathbb{E}_t \left[\int_t^\tau e^{-r(s-t)} H_s ds \right] = \mathbb{E}_t \left[\int_t^\infty e^{-(r+\lambda)(s-t)} H_s ds \right] \quad (\text{B-2})$$

$$= \frac{H_t}{r + \lambda + \rho} + \mathbb{E}_t \left[\int_t^\infty e^{-(r+\lambda)(s-t)} \frac{1}{\theta(r + \lambda + \rho)} z_s ds \right] \quad (\text{B-3})$$

$$= \frac{H_t}{r + \lambda + \rho} + \mathbb{E}_t \left[\int_t^\tau e^{-r(s-t)} \frac{1}{\theta(r + \lambda + \rho)} z_s ds \right]. \quad (\text{B-4})$$

Substituting (B-4) into (B-1) and performing the analogous calculation for the constant $\bar{V}$

yields the principal's value function:

$$V(H, z) = v(z) + \frac{(r + \rho)(\chi - \zeta)}{r + \lambda + \rho} H - \frac{r}{r + \lambda} \bar{V} \quad (\text{B-5})$$

$$v(z) = \max E_t \left[\int_t^\tau e^{-r(s-t)} \left(\frac{(r + \rho)\chi + \lambda\zeta}{\theta(r + \lambda + \rho)} z_s - \frac{1}{2\theta} z_s^2 - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right]. \quad (\text{B-6})$$

B.2 Productive Projects

To describe the boundary conditions and find a lower bound for $v(z)$, consider a feasible, potentially suboptimal policy that chooses β and ϕ so that $dz_t = 0$ (i.e., z_t is stopped) before the separation time τ . From (19), we have $\phi_{stop}(z) = 0$ and

$$\beta_{stop}(z) = z + \frac{\lambda}{r + \rho}(z - \zeta). \quad (\text{B-7})$$

Under this policy, the objective in (B-6) attains $v_{stop}(z)$:

$$v_{stop}(z) = \frac{1}{r + \lambda} \left(\frac{r + \rho + \lambda\zeta}{\theta(r + \lambda + \rho)} z - \frac{1}{2\theta} z^2 \right) - \frac{1}{2(r + \lambda)} \sigma^2 \gamma r \beta_{stop}(z)^2. \quad (\text{B-8})$$

Attaining $v_{stop}(z)$ is always feasible, but potentially suboptimal, so $v(z) \geq v_{stop}(z)$.

We also observe that, since $0 \leq \beta \leq 1$, the integral expression for z_t in (22) implies that $z_t \in [z_L, z_R]$, with values given in (24) and (25). Then, the law of motion for z_t in (19) requires $\beta(z_L) = \beta_{stop}(z_L) = 0$ and $\beta(z_R) = \beta_{stop}(z_R) = 1$, so we must have $v(z_L) = v_{stop}(z_L)$ and $v(z_R) = v_{stop}(z_R)$, which yields

$$v(z_L) = \frac{1}{\theta(r + \lambda)} \left(\frac{r + \rho + \lambda\zeta}{r + \lambda + \rho} z_L - \frac{1}{2} z_L^2 \right) \quad (\text{B-9})$$

$$v(z_R) = \frac{1}{\theta(r + \lambda)} \left(\frac{r + \rho + \lambda\zeta}{r + \lambda + \rho} z_R - \frac{1}{2} z_R^2 \right) - \frac{1}{2(r + \lambda)} \sigma^2 \gamma r. \quad (\text{B-10})$$

Assumption 2 is sufficient to show that $v(z_L) = v_{stop}(z_L) \geq 0$ and $v(z_R) = v_{stop}(z_R) > 0$. If $z_L > 0$, then $v_{stop}(z_L) > 0$. Since $v_{stop}(z)$ is concave, $v(z) \geq v_{stop}(z) > 0$ for every $z \in (z_L, z_R]$.

In a productive pairing, the associated HJB equation for $v(z)$ is

$$\begin{aligned} (r + \lambda)v(z) = \max_{\beta, \phi} \left\{ \frac{r + \rho + \lambda\zeta}{\theta(r + \lambda + \rho)} z - \frac{1}{2\theta} z^2 - \frac{\sigma^2 \gamma r}{2} \beta^2 \right. \\ \left. + v'(z) \left((r + \rho)(z + \sigma^2 \gamma r \beta \phi - \beta) + \lambda(z - \zeta) \right) \right. \\ \left. + \frac{1}{2} v''(z) \phi^2 (r + \rho)^2 \sigma^2 \right\} \end{aligned} \quad (\text{B-11})$$

for $z \in [z_L, z_R]$, remembering that $v(z)$ jumps to zero at τ . The first-order conditions are

$$FOC(\beta) = -\sigma^2 \gamma r \beta + v'(z)(r + \rho)(\sigma^2 \gamma r \phi - 1) \quad (\text{B-12})$$

$$FOC(\phi) = v'(z)(r + \rho)\sigma^2 \gamma r \beta + v''(z)\phi(r + \rho)^2 \sigma^2. \quad (\text{B-13})$$

The calculations above derive the value-function representation, boundary conditions, and HJB underlying Proposition 4. The remaining analytical properties of the principal's value function require substantially longer and more technical proofs, which are presented in the Online Appendix. Specifically, OA.3.1 establishes the control representation, viscosity characterization, continuity, positivity, strict concavity, and unique interior maximizer of v . OA.3.2 derives the policy characterization stated in Proposition 4, and OA.3.3 derives the long-lived-principal benchmark. Section OA.4 develops supplementary regularity and feedback-verification results under an additional technical condition.

B.3 Proof of Lemma 1

Fix $\zeta \in [0, 1]$. Online Appendix OA.3.1 represents $v(\cdot; \zeta)$ as the value of a control problem: for $z \in [z_L(\zeta), z_R(\zeta)]$,

$$v(z; \zeta) = \sup E_z \left[\int_0^\infty e^{-(r+\lambda)s} \left(f(z_s; \zeta) - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right], \quad (\text{B-14})$$

where

$$f(z; \zeta) = \frac{r + \rho + \lambda \zeta}{\theta(r + \lambda + \rho)} z - \frac{1}{2\theta} z^2, \quad (\text{B-15})$$

$$dz_s = [(r + \rho)(z_s + \sigma^2 \gamma r \beta_s \phi_s - \beta_s) + \lambda(z_s - \zeta)] ds + (r + \rho) \sigma \phi_s dZ_s, \quad (\text{B-16})$$

and the supremum is taken over admissible controls $\{\beta_s, \phi_s\}$ with $\beta_s \in [0, 1]$ that keep $z_s \in [z_L(\zeta), z_R(\zeta)]$. Define explicit incentives $x_s = z_s - z_L(\zeta)$. We rewrite each ingredient of (B-14) in terms of x_s .

First, the dynamics. Because $(r + \rho + \lambda)z_L(\zeta) = \lambda\zeta$,

$$(r + \rho)z + \lambda(z - \zeta) = (r + \rho + \lambda)z - \lambda\zeta = (r + \rho + \lambda)x, \quad (\text{B-17})$$

so

$$dx_s = [(r + \rho + \lambda)x_s + (r + \rho)(\sigma^2 \gamma r \beta_s \phi_s - \beta_s)] ds + (r + \rho) \sigma \phi_s dZ_s. \quad (\text{B-18})$$

Second, the state constraint. Because $z_R(\zeta) - z_L(\zeta) = x_R$, the constraint $z_s \in [z_L(\zeta), z_R(\zeta)]$ is equivalent to $x_s \in [0, x_R]$. Hence a control process is admissible from z in the economy with ζ if and only if it is admissible from $x = z - z_L(\zeta)$ for (B-18) on $[0, x_R]$. This holds for each bounded-control class $\mathcal{A}^n$ in Online Appendix OA.3.1, and therefore for the admissible class constructed from them.

Third, the flow payoff. Because $\frac{r+\rho+\lambda\zeta}{r+\lambda+\rho} = z_R(\zeta)$, completing the square gives

$$f(z; \zeta) = \frac{1}{\theta} z_R(\zeta) z - \frac{1}{2\theta} z^2 = \frac{z_R(\zeta)^2}{2\theta} - \frac{(z_R(\zeta) - z)^2}{2\theta} = \frac{z_R(\zeta)^2}{2\theta} - \frac{(x_R - x)^2}{2\theta}. \quad (\text{B-19})$$

Combining the three steps, for every admissible control, the payoff in (B-14) equals

$$\frac{z_R(\zeta)^2}{2\theta(r+\lambda)} + \mathbb{E}_x \left[\int_0^\infty e^{-(r+\lambda)s} \left(-\frac{(x_R - x_s)^2}{2\theta} - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right], \quad (\text{B-20})$$

where x_s follows (B-18) from $x = z - z_L(\zeta)$. The expectation term, the dynamics, and the admissible set do not depend on ζ . Defining $\tilde{v}(x)$ as the supremum of the expectation term over admissible controls proves (32).

C Equilibrium Contracts

C.1 Contracts with the Competitive Fringe

Setting $\chi_s = 0$ specializes the results of Appendices A.1 and B.1 to the competitive fringe.

Lemma C.1 (Incentives from the Competitive Fringe) *If the current contract is from the competitive fringe ($\chi_t = 0$) and all future contracts are productive, then, for $s \in [t, \tau)$,*

$$z_s = \frac{\lambda\zeta}{r + \rho + \lambda}, \quad \phi_s = 0, \quad \hat{h}_s = \frac{z_s}{\theta},$$

and all three are constant.

Proof. Setting $\chi_s = 0$ in (A-15) and using the terminal condition (A-17) gives

$$z_t = \hat{\mathbb{E}}_t [e^{-(r+\rho)(\tau-t)} \zeta] = \frac{\lambda\zeta}{r + \rho + \lambda}.$$

Thus z_t is constant. Its diffusion coefficient, $(r + \rho)\sigma\phi_t$, must therefore vanish, while (17) gives $\hat{h}_t = z_t/\theta$. ■

Lemma C.2 (Competitive Fringe Payoffs) *The principal in a competitive fringe pairing offers $z_0 = \zeta^x = \frac{\lambda\zeta}{r+\lambda+\rho}$. The agent's value from signing, including the initial transfer, is*

$$V_0^{\chi A} = S_{0-} + \frac{1}{2\theta(r+\lambda)} \left(\frac{\lambda\zeta}{r+\lambda+\rho} \right)^2 + \frac{\lambda\zeta}{r+\lambda+\rho} H_0 + \frac{\lambda}{r+\lambda} \bar{V}, \quad (\text{C-21})$$

and the principal's value is zero.

Proof. With $\chi = 0$, Lemma C.1 and Equations (B-5) and (B-6) imply $\beta = 0$ and

$$V(H, z) = \frac{1}{2\theta(r+\lambda)} \left(\frac{\lambda\zeta}{r+\lambda+\rho} \right)^2 - \frac{(r+\rho)\zeta}{r+\lambda+\rho} H - \frac{r}{r+\lambda} \bar{V}. \quad (\text{C-22})$$

Zero profit requires this value to be transferred to the agent at signing. Combining that transfer with (26) gives (C-21). ■

The following lemma verifies that the value of a productive pairing exceeds the corresponding value component under a fringe pairing. It is used in the proofs of Theorem 1 and Proposition 6.

Lemma C.3 (Pair Value Net of the Fringe Value) *Let $\zeta = \kappa \in [0, 1]$. Then*

$$v(\kappa) - \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\kappa}{r + \lambda + \rho} \right)^2 \geq 0, \quad (\text{C-23})$$

with equality if and only if $\kappa = 0$.

Proof. Because the stopping policy (B-7) is feasible, $v(\kappa) \geq v_{stop}(\kappa)$, where v_{stop} is evaluated at $z = \zeta = \kappa$ so that $\beta_{stop}(\kappa) = \kappa$. Substituting into (B-8),

$$v_{stop}(\kappa) - \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\kappa}{r + \lambda + \rho} \right)^2 = \frac{\kappa}{\theta(r + \lambda)} \left[\frac{r + \rho + \lambda\kappa}{r + \lambda + \rho} - \frac{\kappa}{2} \left(1 + \sigma^2 \gamma r \theta + \frac{\lambda^2}{(r + \lambda + \rho)^2} \right) \right].$$

Assumption 2 with $\zeta = \kappa$ gives $\sigma^2 \gamma r \theta < [(r + \rho + \lambda\kappa)^2 - \lambda^2 \kappa^2] / (r + \lambda + \rho)^2$. Substituting this upper bound shows that the bracketed term strictly exceeds

$$\begin{aligned} & \frac{2(r + \lambda + \rho)(r + \rho + \lambda\kappa) - \kappa[(r + \lambda + \rho)^2 + (r + \rho + \lambda\kappa)^2 - \lambda^2 \kappa^2 + \lambda^2]}{2(r + \lambda + \rho)^2} \\ &= \frac{(r + \rho)(r + \lambda + \rho - (r + \rho)\kappa - \lambda\kappa^2)}{(r + \lambda + \rho)^2}, \end{aligned} \quad (\text{C-24})$$

where the equality follows from direct expansion using $r + \lambda + \rho = (r + \rho) + \lambda$. Because $\kappa \leq 1$, we have $(r + \rho)\kappa + \lambda\kappa^2 \leq (r + \rho) + \lambda$, so the right side is nonnegative. Hence the left side of (C-23) is strictly positive for $\kappa \in (0, 1]$. At $\kappa = 0$, we have $z = \kappa = z_L = 0$, so $v(0) = v_{stop}(0) = 0$ by (B-9), and both terms in (C-23) vanish. ■

C.2 Proof of Theorem 1

From Propositions 2, 3, and 4, we have

$$V_0^A = S_0 + \zeta H_0 + \bar{V} \quad (\text{C-25})$$

$$V_0^P = v(\zeta) + \frac{(r + \rho)(1 - \zeta)}{r + \lambda + \rho} H_0 - \frac{r}{r + \lambda} \bar{V} \quad (\text{C-26})$$

where we have used $z_0 = \zeta$ (see (19) and the discussion in Section 3.1). Adding the two surpluses and subtracting $V_0^{\chi A}$ using (C-21), we obtain

$$V_0^P + V_0^A - V_0^{\chi A} = v(\zeta) - \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\zeta}{r + \lambda + \rho} \right)^2 + \frac{r + \rho}{r + \lambda + \rho} H_0. \quad (\text{C-27})$$

From the distribution of surplus in (9), using $V_0^{\chi P} = 0$, the agent's welfare is

$$\begin{aligned} V_0^A &= \kappa \left[V_0^P + V_0^A - V_0^{\chi A} \right] + V_0^{\chi A} \\ &= S_0 + \frac{(r + \rho)\kappa + \lambda\zeta}{r + \lambda + \rho} H_0 + \kappa v(\zeta) + (1 - \kappa) \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\zeta}{r + \lambda + \rho} \right)^2 + \frac{\lambda}{r + \lambda} \bar{V}. \end{aligned} \quad (\text{C-28})$$

Using (C-25), we now have two expressions for the agent's welfare. Matching coefficients on H and the constant term, we have

$$\zeta = \frac{(r + \rho)\kappa + \lambda\zeta}{r + \lambda + \rho} \quad (\text{C-29})$$

$$\bar{V} = \kappa v(\zeta) + (1 - \kappa) \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\zeta}{r + \lambda + \rho} \right)^2 + \frac{\lambda}{r + \lambda} \bar{V}. \quad (\text{C-30})$$

Solving, we obtain (34) and (35).

Finally, Lemma C.3 and (35) imply $\bar{V} \geq 0$, with $\bar{V} > 0$ for $\kappa > 0$ and $\bar{V} = 0$ at $\kappa = 0$. Consequently, $V_t^A \geq S_t$ and $V_t^{\chi A} \geq S_t$ for any t and $H_t \geq 0$, with strict inequalities when $\kappa > 0$ (and also $\lambda > 0$ for $V_t^{\chi A}$). Similarly, matching the constant term in V_t^P yields

$$v(\zeta) - \frac{r}{r + \lambda} \bar{V} = (1 - \kappa) \left[v(\zeta) - \frac{1}{2\theta} \frac{1}{r + \lambda} \left(\frac{\lambda\zeta}{r + \lambda + \rho} \right)^2 \right], \quad (\text{C-31})$$

which is nonnegative by Lemma C.3: it is strictly positive for $\kappa \in (0, 1)$ and equals zero at $\kappa = 0$ and $\kappa = 1$.

Next, consider the first contract in the sequence, which differs from the others because H_{-T} and S_{-T} are known (see Section 2.2). The arguments above remain the same, except that the initial value of z_{-T} is a free choice by the principal, and she chooses $z^* = \arg \max v(z; \kappa)$, yielding additional surplus $v(z^*) - v(\kappa)$ in (C-26) for this first contract. This is implemented by a one-time signing payment of $\kappa(v(z^*) - v(\kappa))$. Finally, including the results of Propositions 2, 3, and 4, we have that (11), (34), (35), and the conditions on the first contract are necessary and sufficient to describe a linear and time-invariant equilibrium, up to welfare-equivalent changes to the compensation process (Proposition 3).

C.3 Proof of Corollary 5

We first show that the principal's optimal policies depend on ζ only through explicit incentives $x = z - z_L(\zeta)$. In the proof of Lemma 1, (B-20) holds control by control, with an additive constant that does not depend on the control. Any comparison between controls held fixed over a short horizon and followed by a common continuation, as in the definition of the policy correspondence in Online Appendix OA.3.2, therefore produces the same ranking at z in the economy with ζ and at $z - z_L(\zeta) + z_L(\zeta')$ in the economy with ζ' . Similarly, substituting (32) into the HJB equation (B-11), with $v'(z; \zeta) = \tilde{v}'(x)$ and $v''(z; \zeta) = \tilde{v}''(x)$ wherever these derivatives exist, shows that the Hamiltonian at z equals $z_R(\zeta)^2/(2\theta)$ plus a function of (x, β, ϕ) that does not depend on ζ , so its pointwise maximizers depend on ζ only through x .

Next, the thresholds. By (32), $v(\cdot; \zeta)$ is a constant plus $\tilde{v}(z - z_L(\zeta))$, so $z^*(\zeta)$ maximizes $v(\cdot; \zeta)$ if and only if $z^*(\zeta) - z_L(\zeta)$ maximizes $\tilde{v}$, which does not depend on ζ ; this difference is x^* . Because the policy correspondence depends on ζ only through x , so does the set of x on which $\beta = 1$, and hence so does the left endpoint $x^{**} = z^{**}(\zeta) - z_L(\zeta)$ of its largest interval adjacent to x_R .

Set $\zeta = \kappa$. By the preceding paragraph, $z^*(\kappa) = z_L(\kappa) + x^*$ and $z^{**}(\kappa) = z_L(\kappa) + x^{**}$, with $z_L(\kappa) = \frac{\lambda\kappa}{r+\rho+\lambda}$. Hence

$$\kappa < z^*(\kappa) \iff \kappa - \frac{\lambda\kappa}{r+\rho+\lambda} < x^* \iff \frac{r+\rho}{r+\rho+\lambda}\kappa < x^* \iff \kappa < \underline{\kappa}, \quad (\text{C-32})$$

and, by the same steps, $\kappa < z^{**}(\kappa)$ if and only if $\kappa < \bar{\kappa}$.

Next, for $\kappa \in (0, 1)$, $z_0 = \kappa$ lies in the interior of the state space: $\kappa > z_L(\kappa)$ because $\lambda < r + \rho + \lambda$, and $\kappa < z_R(\kappa)$ because $(r + \rho + \lambda)\kappa < r + \rho + \lambda\kappa$ is equivalent to $\kappa < 1$. If $\kappa < \underline{\kappa}$, then $z_0 \in (z_L, z^*)$, and Proposition 4 gives $\beta_0 = 0$. If $\underline{\kappa} < \kappa < \bar{\kappa}$, then $z_0 \in (z^*, z^{**})$, so $\beta_0 > 0$ by Proposition 4. If $\kappa > \bar{\kappa}$, then $z_0 \in (z^{**}, z_R)$, and $\beta_0 = 1$ by the definition of z^{**} .

Finally, consider the bounds. Proposition 4 gives $z^* \in (z_L, z_R)$, so $x^* > 0$ and $\underline{\kappa} > 0$. Because $\beta = 0$ on (z_L, z^*) , the full-ownership interval adjacent to z_R cannot extend below z^* , so $x^{**} \geq x^*$ and $\bar{\kappa} \geq \underline{\kappa}$. The same proposition guarantees a nonempty interval adjacent to z_R on which $\beta = 1$, so $z^{**} < z_R$; that is, $x^{**} < x_R = \frac{r+\rho}{r+\rho+\lambda}$, and therefore $\bar{\kappa} < 1$.

C.4 Proof of Proposition 6

Let $c(\kappa) = \frac{1}{2\theta(r+\lambda)} \left(\frac{\lambda\kappa}{r+\lambda+\rho} \right)^2$. Feasibility of the stopping policy implies $v(\kappa; \kappa) \geq v_{stop}(\kappa; \kappa)$, with equality at $\kappa \in \{0, 1\}$ because $z = \kappa$ is then the corresponding boundary. Thus, (35) gives

$$\bar{V}(\kappa) \geq \underline{V}(\kappa) \equiv \frac{r+\lambda}{r} [\kappa v_{stop}(\kappa; \kappa) + (1-\kappa)c(\kappa)], \quad (\text{C-33})$$

again with equality at both endpoints.

First, consider $\kappa = 1$. Direct differentiation of the explicit stopping value gives

$$\left. \frac{d}{d\kappa} [S_t + \kappa H_t + \underline{V}(\kappa)] \right|_{\kappa=1} = H_t + \frac{1}{2\theta r} \left[2 - \left(\frac{r+\rho}{r+\lambda+\rho} \right)^2 - 3\sigma^2 \gamma r \theta \right]. \quad (\text{C-34})$$

The right side is negative whenever

$$\sigma^2 \gamma r \theta > \frac{1}{3} \left[2 - \left(\frac{r+\rho}{r+\lambda+\rho} \right)^2 + 2\theta r H_t \right]. \quad (\text{C-35})$$

Under this condition, $S_t + \kappa H_t + \underline{V}(\kappa)$ is larger immediately to the left of $\kappa = 1$ than at $\kappa = 1$. Because it is a lower bound for V_t^A and agrees with V_t^A at $\kappa = 1$, the same endpoint comparison holds for the agent's value without requiring differentiability of $v(\kappa; \kappa)$.

For the remaining comparisons, denote total net surplus and its stopping-policy lower

bound by

$$\mathcal{J}_t(\kappa) = v(\kappa; \kappa) - c(\kappa) + \frac{r + \rho}{r + \lambda + \rho} H_t \quad (\text{C-36})$$

$$\underline{\mathcal{J}}_t(\kappa) = v_{stop}(\kappa; \kappa) - c(\kappa) + \frac{r + \rho}{r + \lambda + \rho} H_t. \quad (\text{C-37})$$

$\mathcal{J}_t(\kappa) \geq \underline{\mathcal{J}}_t(\kappa)$, with equality at both endpoints. In addition, (C-26) and (C-31) imply

$$V_t^P(\kappa) = (1 - \kappa) \mathcal{J}_t(\kappa). \quad (\text{C-38})$$

Lemma C.3 and $H_t > 0$ imply $\mathcal{J}_t(\kappa) > 0$. Hence $V_t^P(1) = 0 < V_t^P(\kappa)$ for $\kappa < 1$. Direct differentiation also gives

$$\left. \frac{d}{d\kappa} \mathcal{J}_t(\kappa) \right|_{\kappa=1} = \frac{1}{\theta(r + \lambda)} \left[\frac{\lambda(r + \rho)}{(r + \lambda + \rho)^2} - \sigma^2 \gamma r \theta \right]. \quad (\text{C-39})$$

If risk is sufficiently high, this derivative and (C-34) are negative. Because both lower bounds bind at $\kappa = 1$, total surplus and agent welfare are larger immediately to its left, completing the comparisons at that endpoint.

At $\kappa = 0$,

$$\left. \frac{d}{d\kappa} \mathcal{J}_t(\kappa) \right|_{\kappa=0} = \frac{r + \rho}{\theta(r + \lambda)(r + \rho + \lambda)} > 0. \quad (\text{C-40})$$

Moreover, the derivative of $S_t + \kappa H_t + \underline{V}(\kappa)$ is $H_t > 0$. Since both bounds bind at $\kappa = 0$, total surplus and agent welfare rise for sufficiently small $\kappa > 0$.

Finally, $(1 - \kappa) \underline{\mathcal{J}}_t(\kappa)$ is a lower bound for $V_t^P(\kappa)$ that agrees with it at $\kappa = 0$, and

$$\left. \frac{d}{d\kappa} [(1 - \kappa) \underline{\mathcal{J}}_t(\kappa)] \right|_{\kappa=0} = \frac{r + \rho}{r + \lambda + \rho} \left[\frac{1}{\theta(r + \lambda)} - H_t \right]. \quad (\text{C-41})$$

This derivative is strictly positive if $H_t < 1/(\theta(r + \lambda))$. The lower bound is then larger for sufficiently small $\kappa > 0$ than at $\kappa = 0$, and so is V_t^P , completing the proof.

C.5 Proof of Corollary 1

Proof. By the proof of Corollary 5 (Appendix C.3), $z^*(\kappa) = z_L(\kappa) + x^*$, with $z_L(\kappa) = \lambda\kappa/(r + \rho + \lambda)$. Thus $\kappa - z_L(\kappa) = x_R\kappa$, which gives (41). The additive term $z_R(\kappa)^2/(2\theta(r + \lambda))$ in (32) cancels from the difference between the two values, yielding (42). Proposition 4 gives continuity, strict concavity, and the unique interior maximizer of v ; hence $\tilde{v}$ has the same properties. Composing $-\tilde{v}$ with the increasing affine map $\kappa \mapsto x_R\kappa$ gives strict convexity of D . The maximizer x^* gives its unique zero at $\kappa = x^*/x_R$. Strict concavity implies that $\tilde{v}$ strictly increases to x^* and strictly decreases thereafter, so D has the stated strict monotonicity on either side of its minimum. No derivatives are needed.

Equation (17) converts the incentive-state gap into the initial investment gap by division by θ . The value representation (29) has expertise and continuation-utility terms independent of z when the future market environment is held fixed; they cancel in the entry comparison.

Under CARA, savings and the level of the agent's promised utility do not affect v or its maximizer. Thus the loss and investment gap are independent of H_0 and S_0 . Finally, the first-contract argument in the proof of Theorem 1 establishes the additional surplus D and its division into κD for the agent and $(1 - \kappa)D$ for the principal. This free-entry choice is feasible when initial expertise is public; it is not a feasible alternative entry state for a post-initial contract subject to $z_0 = \kappa$. ■

Online Appendix for
 Buying In and Selling Out:
 The Dynamic Returns to Investing in Expertise

OA.1 Benchmark Economies

OA.1.1 The First-Best

We maximize the social surplus without hidden information: the expected value of output from all projects in a sequence given the agent's expected utility. The problem is

$$V(H_t, W_t) = \max_{c, h, \beta} E_t \left[\int_t^\infty e^{-r(s-t)} ((r + \rho)H_s - c_s) ds \right] \quad (\text{OA.1-42})$$

$$s.t. \quad dW_t = (rW_t - u(c_t, h_t))dt - \sigma\gamma rW_t\beta_t dZ_t \quad (\text{OA.1-43})$$

$$dH_t = (h_t - \rho H_t) dt \quad (\text{OA.1-44})$$

where W_t is the agent's continuation utility, and W_0 and H_0 are given initial values. We can solve the maximization problem with the standard dynamic programming approach. The social value is

$$V(H_t, W_t) = H_t + \frac{1}{\gamma r} \ln(-\gamma r W_t) + \frac{1}{2\theta r} \quad (\text{OA.1-45})$$

and the solution has constant consumption and investment in expertise:

$$\beta_t = 0, \quad c_t = -\frac{1}{\gamma} \ln(-\gamma r W_t) + \frac{1}{2\theta}, \quad h_t = \frac{1}{\theta} \quad (\text{OA.1-46})$$

OA.1.2 The Agent Owns the Firm

We analyze the case in which the agent owns a long-lived project without requiring a principal, meaning output is paid directly into the agent's savings ($dY_t = dM_t$). The agent's problem is

$$W(H_t, S_t) = \max_{c, h} E_t \left[\int_t^\infty e^{-r(s-t)} u(c, h) ds \right] \quad (\text{OA.1-47})$$

$$s.t. \quad dS_t = (rS_t - c_t + (r + \rho)H_t) dt + \sigma dZ_t \quad (\text{OA.1-48})$$

$$dH_t = (h_t - \rho H_t) dt \quad (\text{OA.1-49})$$

given initial values of S_0 and H_0 . We can solve the agent's problem with the standard dynamic programming approach. The agent's value is

$$W(H, S) = -\frac{1}{\gamma r} \exp \left(-\gamma r S - \gamma r H - \gamma \left[\frac{1}{2\theta} - \frac{1}{2} \gamma r \sigma^2 \right] \right) \quad (\text{OA.1-50})$$

and the solution again has constant investment in expertise:

$$h_t = \frac{1}{\theta} \quad (\text{OA.1-51})$$

$$c_t = -\frac{1}{\gamma} \ln(-\gamma r W_t) + \frac{1}{2\theta} \quad (\text{OA.1-52})$$

If we follow the first-best and write $dW_t = (rW_t - u(c_t, h_t))dt - \sigma\gamma r W_t \beta_t dZ_t$, then the long-lived project's solution yields $\beta_t = 1$. Using $W(H, S)$ to find the certainty equivalent wealth yields V_t^A .

OA.2 A Bound on Gains from Deviation

We construct an upper bound for the payoff after deviation $(\Delta S, \Delta H)$.²⁵ We construct an upper bound of the form:

$$\mathcal{J}(W, \Delta S, \Delta H, z) = W \exp(-\gamma r \Delta S - \gamma r z \Delta H), \quad (\text{OA.2-53})$$

where W is the equilibrium level of continuation utility. Since dY_t is observable, (2) implies $d\hat{Z}_t = \frac{(r+\rho)\Delta H_t}{\sigma} dt + dZ_t$. Then (16) and (19) become

$$dW_t = -\gamma r W_t \beta_t (r + \rho) \Delta H_t dt - \sigma \gamma r W_t \beta_t dZ_t \quad (\text{OA.2-54})$$

$$\begin{aligned} dz_t = & (r + \rho) (z_t + \sigma^2 \gamma r \beta_t \phi_t + (r + \rho) \Delta H_t \phi_t - \beta_t) dt + (r + \rho) \sigma \phi_t dZ_t \\ & + (\zeta - z_t)(dN_t - \lambda dt), \end{aligned} \quad (\text{OA.2-55})$$

We also have laws of motion for $d\Delta H_t$ (A-1) and $d\Delta S_t$ (A-2), remembering that the agent reveals his previous reports to each new principal.²⁶

To show that $\mathcal{J}$ is an upper bound on the agent's payoff, it is sufficient to verify that $u(\hat{c}_t + \Delta c_t, \hat{h}_t + \Delta h_t) - r\mathcal{J} + E[d\mathcal{J}] \frac{1}{dt} \leq 0$ and then to verify the transversality conditions.

²⁵This is analogous to the bound in Sannikov (2014), He et al. (2017), or Marinovic and Varas (2019), extended to hidden expertise and jump-diffusions.

²⁶The bound in (OA.2-53) shows that z_t can be interpreted as the relative value of $H_t - \hat{H}_t$ and $S_t - \hat{S}_t$, i.e., the marginal rate of substitution between *lies* about savings and about expertise. This complements Section 3.1, where z_t is the marginal rate of substitution between true savings and expertise. The bound uses that future contracts discipline current behavior because lies about durable expertise are durable. The agent cannot lie to the current principal and then jump to the truth upon beginning work for a new principal because the agent's reports to the previous principal are verifiable. Otherwise, the agent could periodically wipe out his lies and a new set of incentives would be required to ensure incentive compatibility. Moreover, the agent's welfare from deviations is proportional to zH (inside value) rather than ζH (outside or residual value). This is why we require in Assumption 1 that the agent can commit not to quit early. If the agent shirks a small amount reducing H_t by ϵ , he loses $z_t \epsilon$ in certainty equivalent wealth (from 22). Quitting would instead price this hidden change at ζ . If the agent could quit and re-pair immediately, he could double deviate by shirking and quitting (or work excessively and then quit, depending on the sign of $z_t - \zeta$). Then, the only incentive compatible level of z_t would be ζ .

Calculating from (OA.2-53), we have

$$\begin{aligned}
\mathbb{E}[d\mathcal{J}] \frac{1}{dt} = & (-\gamma r(r + \rho)\beta_t \Delta H_t) \mathcal{J} \\
& + (-\gamma r(r \Delta S_t - \Delta c_t - \lambda \Delta \nu_t)) \mathcal{J} \\
& + (-\gamma r z_t (\Delta h_t - \rho \Delta H_t)) \mathcal{J} \\
& + (-\gamma r \Delta H_t ((r + \rho) (z_t + \sigma^2 \gamma r \beta_t \phi_t \\
& \quad + (r + \rho) \Delta H_t \phi_t - \beta_t) + \lambda(z_t - \zeta))) \mathcal{J} \\
& + \frac{1}{2} \gamma^2 r^2 (\Delta H_t)^2 \sigma^2 (r + \rho)^2 \phi_t^2 \mathcal{J} \\
& + (-\sigma \gamma r \beta_t) \sigma (r + \rho) \phi_t (-\gamma r \Delta H_t) \mathcal{J} \\
& + \lambda (e^{-\gamma r \Delta \nu_t - \gamma r \Delta H_t (\zeta - z_t)} - 1) \mathcal{J}
\end{aligned} \tag{OA.2-56}$$

Similarly, from (18) we have

$$\begin{aligned}
u(\hat{c}_t + \Delta c_t, \hat{h}_t + \Delta h_t) &= r W_t \exp \left(-\gamma \Delta c_t + \frac{\gamma \theta}{2} \Delta h_t^2 + \gamma z_t \Delta h_t \right) \\
&= r \exp \left(-\gamma \Delta c_t + \frac{\gamma \theta}{2} \Delta h_t^2 + \gamma z_t \Delta h_t + \gamma r \Delta S_t + \gamma r z_t \Delta H_t \right) \mathcal{J}
\end{aligned} \tag{OA.2-57}$$

Since $\mathcal{J} < 0$, we divide and reverse the inequality, so $u(\hat{c}_t + \Delta c_t, \hat{h}_t + \Delta h_t) - r \mathcal{J} + \mathbb{E}[d\mathcal{J}] \frac{1}{dt} \leq 0$ if and only if

$$\begin{aligned}
0 \leq & (-\gamma r \beta_t (r + \rho) \Delta H_t) + (-\gamma r (r \Delta S_t - \Delta c_t - \lambda \Delta \nu_t)) + (-\gamma r z_t (\Delta h_t - \rho \Delta H_t)) \\
& + (-\gamma r \Delta H_t ((r + \rho) (z_t + \sigma^2 \gamma r \beta_t \phi_t + \Delta H_t (r + \rho) \phi_t - \beta_t) + \lambda(z_t - \zeta))) \\
& + \frac{1}{2} \gamma^2 r^2 (\Delta H_t)^2 \sigma^2 (r + \rho)^2 \phi_t^2 + (-\sigma \gamma r \beta_t) \sigma (r + \rho) \phi_t (-\gamma r \Delta H_t) \\
& + \lambda (e^{-\gamma r \Delta \nu_t - \gamma r \Delta H_t (\zeta - z_t)} - 1) - r \\
& + r \exp \left(-\gamma \Delta c_t + \frac{\gamma \theta}{2} \Delta h_t^2 + \gamma z_t \Delta h_t + \gamma r \Delta S_t + \gamma r z_t \Delta H_t \right)
\end{aligned} \tag{OA.2-58}$$

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

$$\begin{aligned}
0 \leq & (-\gamma r \beta_t (r + \rho) \Delta H_t) + (-\gamma r (r \Delta S_t - \Delta c_t - \lambda \Delta \nu_t)) + (-\gamma r z_t (\Delta h_t - \rho \Delta H_t)) \\
& + (-\gamma r \Delta H_t ((r + \rho) (z_t + \sigma^2 \gamma r \beta_t \phi_t + \Delta H_t (r + \rho) \phi_t - \beta_t) + \lambda(z_t - \zeta))) \\
& + \frac{1}{2} \gamma^2 r^2 (\Delta H_t)^2 \sigma^2 (r + \rho)^2 \phi_t^2 + (-\sigma \gamma r \beta_t) \sigma (r + \rho) \phi_t (-\gamma r \Delta H_t) \\
& + \lambda (-\gamma r \Delta \nu_t - \gamma r \Delta H_t (\zeta - z_t)) - r \\
& + r \left(1 - \gamma \Delta c_t + \frac{\gamma \theta}{2} \Delta h_t^2 + \gamma z_t \Delta h_t + \gamma r \Delta S_t + \gamma r z_t \Delta H_t \right)
\end{aligned} \tag{OA.2-59}$$

Canceling terms, this is equivalent to

$$0 \leq -\gamma r(\Delta H_t)^2(r + \rho)^2\phi_t + \frac{1}{2}\gamma^2 r^2(\Delta H_t)^2\sigma^2(r + \rho)^2\phi_t^2 + r\frac{\gamma\theta}{2}\Delta h_t^2 \quad (\text{OA.2-60})$$

which is always met if $\phi_t \leq 0$.

Finally, we verify the transversality conditions. The boundedness of ΔH_t in Assumption 1 implies that the market price of risk in (A-5) is bounded, so ξ_t is a true martingale. To justify the terminal term under the change of measure, let η_n be a localizing sequence for the stochastic integrals in W and ξW . Applying the preceding argument to the stopped deviation and then using the change of measure gives

$$\limsup_{s \rightarrow \infty} \hat{\mathbb{E}}_t \left[e^{-r(s-t)} W_{s \wedge \eta_n} (\xi_{s \wedge \eta_n} - 1) \right] \leq 0$$

for each n . Letting $n \rightarrow \infty$ preserves the inequality by Fatou's lemma because $W_t < 0$. Thus,

$$\limsup_{s \rightarrow \infty} \hat{\mathbb{E}}_t \left[e^{-r(s-t)} W_s (\xi_s - 1) \right] \leq 0.$$

It remains to check the terminal terms associated with ΔS_t and ΔH_t . Since $W_t < 0$ and W_t is a true martingale under $\hat{\mathbb{E}}$, $\hat{\mathbb{E}}_t[|W_s|] = -\hat{\mathbb{E}}_t[W_s] = -W_t < \infty$, and hence $\lim_{s \rightarrow \infty} \hat{\mathbb{E}}_t[e^{-r(s-t)}|W_s|] = 0$. The bounds on ΔH_t and ΔS_t (Assumption 1) along with $p_t^S = -\gamma r W_t$ and $p_t^H/W_t \in [-\gamma r, 0]$ imply

$$\lim_{s \rightarrow \infty} \hat{\mathbb{E}}_t \left[e^{-r(s-t)} p_s^S \Delta S_s \right] = \lim_{s \rightarrow \infty} -\gamma r \hat{\mathbb{E}}_t \left[e^{-r(s-t)} W_s \Delta S_s \right] = 0 \quad (\text{OA.2-61})$$

$$\lim_{s \rightarrow \infty} \hat{\mathbb{E}}_t \left[e^{-r(s-t)} p_s^H \Delta H_s \right] = \lim_{s \rightarrow \infty} \hat{\mathbb{E}}_t \left[e^{-r(s-t)} \frac{p_s^H}{W_s} W_s \Delta H_s \right] = 0 \quad (\text{OA.2-62})$$

OA.3 Value Function and Policy Characterization

This section proves the remaining statements of Proposition 4 not established in Appendix B. Subsection OA.3.1 establishes the properties of the value function. Subsection OA.3.2 derives the policy characterization stated in Proposition 4. Subsection OA.3.3 derives the long-lived-principal benchmark for $\lambda = 0$.

OA.3.1 Properties of the Value Function

Substituting (27) and (28) into the principal's objective and integrating by parts gives

$$\mathbb{E}_t \left[\int_t^\tau e^{-r(s-t)} H_s ds \right] = \frac{H_t}{r + \lambda + \rho} + \mathbb{E}_t \left[\int_t^\tau e^{-r(s-t)} \frac{z_s}{\theta(r + \lambda + \rho)} ds \right]. \quad (\text{OA.3-63})$$

Applying the same calculation to the constant $\bar{V}$ yields the stated decomposition of $V^P(H, z)$ and the control representation for v . Integrating over the independent Poisson stopping time

and setting the initial time equal to zero gives the equivalent infinite-horizon representation

$$v(z) = \sup_{\{\beta_s, \phi_s\}} \mathbb{E}_z \left[\int_0^\infty e^{-(r+\lambda)s} \left(f(z_s) - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right], \quad (\text{OA.3-64})$$

where

$$f(z) = \frac{r + \rho + \lambda \zeta}{\theta(r + \lambda + \rho)} z - \frac{1}{2\theta} z^2, \quad (\text{OA.3-65})$$

$$dz_s = \mu(z_s, \beta_s, \phi_s) ds + (r + \rho) \sigma \phi_s dZ_s, \quad (\text{OA.3-66})$$

$$\mu(z, \beta, \phi) = (r + \rho)(z + \sigma^2 \gamma r \beta \phi - \beta) + \lambda(z - \zeta). \quad (\text{OA.3-67})$$

For each $n \geq 1$, let $\mathcal{A}^n$ contain the progressively measurable controls satisfying $\beta_s \in [0, 1]$ and $|\phi_s| \leq n$ that keep $z_s \in [z_L, z_R]$. The admissible class in the control representation above is the finite-horizon payoff closure of $\bigcup_{n \geq 1} \mathcal{A}^n$, with the stopped policy used after localization, and is stable under concatenation at stopping times. These properties give the dynamic programming principle used below.

The function f is a concave quadratic with maximum $f(z_R) = z_R^2/(2\theta)$ on the state space. Since the cost of β is non-positive, every admissible payoff is bounded above by $z_R^2/(2\theta(r + \lambda))$. For the lower bound, fix $z \in [z_L, z_R]$ and choose

$$\phi_s = 0, \quad \beta_s = \beta_{stop}(z) = z + \frac{\lambda}{r + \rho}(z - \zeta). \quad (\text{OA.3-68})$$

This policy is feasible because β_{stop} is affine with endpoint values zero and one, and

$$(r + \rho)(z - \beta_{stop}(z)) + \lambda(z - \zeta) = 0, \quad (\text{OA.3-69})$$

so the state remains constant and the policy delivers $v_{stop}(z)$. Therefore

$$v_{stop}(z) \leq v(z) \leq \frac{z_R^2}{2\theta(r + \lambda)}. \quad (\text{OA.3-70})$$

These bounds prove finiteness. Under Assumption 2, $v_{stop}(z_L) \geq 0$ and $v_{stop}(z_R) > 0$. Strict concavity of v_{stop} and the lower bound give the stated positivity conclusions for v .

Fix $z_0 \in (z_L, z_R)$ and choose an interval $I = [z_0 - \delta, z_0 + \delta] \subset (z_L, z_R)$ and $\eta > 0$ such that $\beta_{stop}(z) \pm \eta \in [0, 1]$ for every $z \in I$. Starting from $z \in I$, set $\phi_s = 0$ and choose $\beta_s = \beta_{stop}(z_s) - \eta$ when $z_s < z_0$ and $\beta_s = \beta_{stop}(z_s) + \eta$ when $z_s > z_0$. This policy moves the state deterministically toward z_0 at speed $(r + \rho)\eta$, so its hitting time satisfies

$$\tau_z \leq \frac{|z - z_0|}{(r + \rho)\eta}. \quad (\text{OA.3-71})$$

After τ_z , switch to an ϵ -optimal policy from z_0 . The steering payoff converges to zero and the discount factor converges to one as $z \rightarrow z_0$, which gives

$$\liminf_{z \rightarrow z_0} v(z) \geq v(z_0) - \epsilon. \quad (\text{OA.3-72})$$

Letting $\epsilon \downarrow 0$ proves that v is lower semicontinuous on the interior.

Fix $z_0 \in (z_L, z_R)$ and let a C^2 function ψ touch v from below at z_0 on a compact neighborhood $[a, b] \subset (z_L, z_R)$. For any fixed $\phi \in \mathbb{R}$, apply the constant controls $(0, \phi)$ until

$$\tau_h = h \wedge \inf\{s \geq 0 : z_s \notin (a, b)\}. \quad (\text{OA.3-73})$$

The dynamic programming principle and Itô's Lemma applied to $e^{-(r+\lambda)s}\psi(z_s)$ give, after division by $E_{z_0}[\tau_h]$ and passage to the limit $h \downarrow 0$,

$$(r + \lambda)v(z_0) \geq f(z_0) + \psi'(z_0)((r + \rho + \lambda)z_0 - \lambda\zeta) + \frac{1}{2}\psi''(z_0)(r + \rho)^2\sigma^2\phi^2 \quad (\text{OA.3-74})$$

for all $\phi \in \mathbb{R}$. Finiteness rules out $\psi''(z_0) > 0$. Thus, whenever a C^2 function ψ touches v from below at a point of (z_L, z_R) , its second derivative at that point is non-positive. Because v is also finite and lower semicontinuous, this property implies that v is concave on (z_L, z_R) . Suppose not. Then there exist $z_1 < z_2$ in (z_L, z_R) and $\bar{z} \in (z_1, z_2)$ with $v(\bar{z}) < \ell(\bar{z})$, where ℓ is the affine function satisfying $\ell(z_i) = v(z_i)$ for $i = 1, 2$. Define

$$q(z) = (z - z_1)(z_2 - z), \quad (\text{OA.3-75})$$

so that $q \geq 0$ on $[z_1, z_2]$, $q(z_1) = q(z_2) = 0$, and $q'' \equiv -2$. Fix

$$0 < \epsilon < \frac{\ell(\bar{z}) - v(\bar{z})}{q(\bar{z})}. \quad (\text{OA.3-76})$$

The function $w(z) = v(z) - \ell(z) + \epsilon q(z)$ is lower semicontinuous on the compact interval $[z_1, z_2]$ and therefore attains its minimum c . At the endpoints $w = 0$, while $w(\bar{z}) < 0$ by the choice of ϵ , so $c < 0$ and every minimizer is interior. Let $z_\epsilon \in (z_1, z_2)$ be a minimizer and set

$$\psi(z) = \ell(z) - \epsilon q(z) + c. \quad (\text{OA.3-77})$$

Then $v - \psi = w - c \geq 0$ on $[z_1, z_2]$ with equality at z_ϵ , so ψ is a C^2 function touching v from below at the interior point z_ϵ . But $\psi''(z_\epsilon) = 2\epsilon > 0$, a contradiction. Hence v is concave on (z_L, z_R) . A finite concave function is locally Lipschitz and continuous on the interior.

Next, for the bounded-control classes defined above, let

$$v^n(z) = \sup_{(\beta, \phi) \in \mathcal{A}^n} E_z \left[\int_0^\infty e^{-(r+\lambda)s} \left(f(z_s) - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right]. \quad (\text{OA.3-78})$$

The associated Hamiltonian is

$$\mathcal{H}^n(z, p, X) = \sup_{\beta \in [0, 1], |\phi| \leq n} \left\{ f(z) - \frac{\sigma^2 \gamma r}{2} \beta^2 + p\mu(z, \beta, \phi) + \frac{1}{2}X(r + \rho)^2\sigma^2\phi^2 \right\}.$$

For each n , the control set is compact, the coefficients are continuous and Lipschitz in the state, and the state interval is viable. Standard dynamic programming results for state-constrained control problems (see, e.g., [Fleming and Soner, 2006](#) and [Crandall, Ishii, and](#)

Lions, 1992) therefore imply that v^n is continuous and satisfies the viscosity equation

$$(r + \lambda)v^n(z) = \mathcal{H}^n(z, (v^n)'(z), (v^n)''(z)) \quad (\text{OA.3-79})$$

for $z \in (z_L, z_R)$, with state-constraint boundary conditions at z_L and z_R . Thus the displayed Hamiltonian characterizes the interior equation; the endpoint restrictions enter through viability and the state-constraint boundary conditions, which are handled below.

The classes $\mathcal{A}^n$ increase with n , so v^n is increasing and is bounded above uniformly by the bound established in the first step. Since the admissible class is the finite-horizon payoff closure of the bounded classes, bounded flow payoffs and discounting imply by definition

$$v^n(z) \uparrow v(z). \quad (\text{OA.3-80})$$

Because the limit v is continuous on the interior, Dini's theorem gives locally uniform convergence on every compact subinterval of (z_L, z_R) . The Hamiltonians $\mathcal{H}^n$ increase to

$$\mathcal{H}(z, p, X) = \sup_{\beta \in [0,1], \phi \in \mathbb{R}} \left\{ f(z) - \frac{\sigma^2 \gamma r}{2} \beta^2 + p\mu(z, \beta, \phi) + \frac{1}{2} X(r + \rho)^2 \sigma^2 \phi^2 \right\}. \quad (\text{OA.3-81})$$

We verify stability separately for upper and lower test functions because the limiting Hamiltonian is extended-valued. Let ψ touch v from above at an interior point z_0 . Local uniform convergence gives points $z_n \rightarrow z_0$ where ψ is an upper test for v^n up to an immaterial vanishing constant. If $\mathcal{H}(z_0, \psi'(z_0), \psi''(z_0)) = +\infty$, the limiting subsolution inequality is automatic. If the Hamiltonian is finite on a neighborhood of the limiting slope and curvature, the maximizing ϕ is locally bounded, $\mathcal{H}^n$ converges locally uniformly to $\mathcal{H}$, and passage to the limit in the subsolution inequalities for v^n gives

$$(r + \lambda)v(z_0) \leq \mathcal{H}(z_0, \psi'(z_0), \psi''(z_0)). \quad (\text{OA.3-82})$$

If the limiting slope and curvature approach values at which the Hamiltonian ceases to be locally finite, the subsolution condition is automatic whenever the upper semicontinuous envelope of the monotone Hamiltonians is infinite. The only remaining finite case is $(\psi'(z_0), \psi''(z_0)) = (0, 0)$. In that case, apply the dynamic programming principle directly, localize the state to a compact neighborhood of z_0 , and use controls whose payoff is within $o(h)$ of the value over a horizon h . The upper-contact property bounds the continuation value by the test function. Applying Itô's Lemma, dividing by h , and letting $h \downarrow 0$ gives

$$(r + \lambda)v(z_0) \leq f(z_0), \quad (\text{OA.3-83})$$

since the flow Hamiltonian is maximized by $\beta = 0$ and equals $f(z_0)$ when the test function has zero slope and curvature. Thus the ordinary subsolution inequality also holds in the zero-slope, zero-curvature case.

Now let ψ touch v from below at z_0 , and choose lower-contact points $z_n \rightarrow z_0$ for v^n . Fix an integer m . For every $n \geq m$, the supersolution inequality and monotonicity of the Hamiltonians give

$$(r + \lambda)v^n(z_n) \geq \mathcal{H}^n(z_n, \psi'(z_n), \psi''(z_n)) \geq \mathcal{H}^m(z_n, \psi'(z_n), \psi''(z_n)). \quad (\text{OA.3-84})$$

Since $\mathcal{H}^m$ is continuous, passage to the limit gives

$$(r + \lambda)v(z_0) \geq \mathcal{H}^m(z_0, \psi'(z_0), \psi''(z_0)). \quad (\text{OA.3-85})$$

Letting $m \rightarrow \infty$ yields

$$(r + \lambda)v(z_0) \geq \mathcal{H}(z_0, \psi'(z_0), \psi''(z_0)). \quad (\text{OA.3-86})$$

The finite left-hand side forces every lower test to have slope and curvature for which the Hamiltonian is finite, consistently with the curvature argument in the first step. Together with (OA.3-83), this proves that v is a viscosity solution of the limiting HJB equation.

Moreover, at either endpoint, the state constraint requires $\phi = 0$, and

$$\mu(z_L, \beta, 0) = -(r + \rho)\beta, \quad \mu(z_R, \beta, 0) = (r + \rho)(1 - \beta). \quad (\text{OA.3-87})$$

The drift restrictions therefore force $\beta = 0$ at z_L and $\beta = 1$ at z_R . The state remains fixed under these stopped policies, so the same viability argument for the original and bounded-control problems gives

$$v(z_L) = v_{\text{stop}}(z_L), \quad v(z_R) = v_{\text{stop}}(z_R), \quad (\text{OA.3-88})$$

$$v^n(z_L) = v_{\text{stop}}(z_L), \quad v^n(z_R) = v_{\text{stop}}(z_R) \quad \text{for every } n. \quad (\text{OA.3-89})$$

The stopped-policy lower bound and continuity of v_{stop} give

$$\liminf_{z \downarrow z_L} v(z) \geq v_{\text{stop}}(z_L), \quad \liminf_{z \uparrow z_R} v(z) \geq v_{\text{stop}}(z_R). \quad (\text{OA.3-90})$$

For the reverse inequalities, use barriers that apply uniformly to the bounded-control problems. Define

$$a = r + \rho + \lambda, \quad k = \sigma^2 \gamma^2 r^2, \quad (\text{OA.3-91})$$

and choose $\alpha \in (0, 1)$ satisfying

$$\alpha \left(a + \frac{k}{2(1 - \alpha)} \right) < r + \lambda. \quad (\text{OA.3-92})$$

On fixed one-sided neighborhoods of the endpoints, define

$$\bar{v}_L(z) = v_{\text{stop}}(z_L) + C(z - z_L)^\alpha, \quad \bar{v}_R(z) = v_{\text{stop}}(z_R) + C(z_R - z)^\alpha. \quad (\text{OA.3-93})$$

We now verify the supersolution inequalities. Write $x = z - z_L$. For the left barrier,

$$\bar{v}'_L(z) = C\alpha x^{\alpha-1}, \quad \bar{v}''_L(z) = -C\alpha(1 - \alpha)x^{\alpha-2}. \quad (\text{OA.3-94})$$

After optimizing over ϕ for a fixed β , the terms in the Hamiltonian that depend on β are

$$-(r + \rho)\bar{v}'_L(z)\beta - \frac{\sigma^2 \gamma r}{2}\beta^2 - \frac{\sigma^2 \gamma^2 r^2 (\bar{v}'_L(z))^2}{2\bar{v}''_L(z)}\beta^2. \quad (\text{OA.3-95})$$

Since

$$\frac{(\bar{v}'_L(z))^2}{-\bar{v}''_L(z)} = \frac{C\alpha}{1-\alpha}x^\alpha, \quad (\text{OA.3-96})$$

the negative term of order $Cx^{\alpha-1}\beta$ dominates the possible positive quadratic term of order $Cx^\alpha\beta^2$ on a sufficiently small neighborhood of z_L , uniformly in C . Hence the maximizer is $\beta = 0$, followed by $\phi = 0$, and

$$\mathcal{H}(z, \bar{v}'_L(z), \bar{v}''_L(z)) = f(z) + aC\alpha x^\alpha. \quad (\text{OA.3-97})$$

Because $(r + \lambda)v_{stop}(z_L) = f(z_L)$,

$$(r + \lambda)\bar{v}_L(z) - \mathcal{H}(z, \bar{v}'_L(z), \bar{v}''_L(z)) = C(r + \lambda - a\alpha)x^\alpha + f(z_L) - f(z). \quad (\text{OA.3-98})$$

The choice of α implies $r + \lambda - a\alpha > 0$. Since $f(z) - f(z_L)$ is of order x , the right-hand side is nonnegative near z_L for sufficiently large C .

For the right barrier, write $y = z_R - z$. Then

$$\bar{v}'_R(z) = -C\alpha y^{\alpha-1}, \quad \bar{v}''_R(z) = -C\alpha(1-\alpha)y^{\alpha-2}. \quad (\text{OA.3-99})$$

The positive marginal value of increasing β is of order $Cy^{\alpha-1}$, while the curvature adjustment after optimizing over ϕ is of order Cy^α . On a sufficiently small left neighborhood of z_R , the maximizer is therefore $\beta = 1$, and optimization over ϕ gives

$$\mathcal{H}(z, \bar{v}'_R(z), \bar{v}''_R(z)) = f(z) - \frac{\sigma^2\gamma r}{2} + C\alpha \left(a + \frac{k}{2(1-\alpha)} \right) y^\alpha. \quad (\text{OA.3-100})$$

Since $(r + \lambda)v_{stop}(z_R) = f(z_R) - \sigma^2\gamma r/2$,

$$(r + \lambda)\bar{v}_R(z) - \mathcal{H}(z, \bar{v}'_R(z), \bar{v}''_R(z)) = C \left(r + \lambda - \alpha a - \frac{\alpha k}{2(1-\alpha)} \right) y^\alpha + f(z_R) - f(z).$$

Both terms are nonnegative by the choice of α and the monotonicity of f on $[z_L, z_R]$. Thus $\bar{v}_L$ and $\bar{v}_R$ are supersolutions of the unbounded-control equation on their respective neighborhoods. They are consequently supersolutions for every bounded Hamiltonian $\mathcal{H}^n$, with constants independent of n .

Choose C large enough that the barriers also dominate the uniform upper bound at the interior ends of the two neighborhoods. State-constrained comparison for each bounded-control problem yields $v^n(z) \leq \bar{v}_L(z)$ and $v^n(z) \leq \bar{v}_R(z)$ on the corresponding neighborhoods. Passing to the limit $v^n \uparrow v$ yields

$$\limsup_{z \downarrow z_L} v(z) \leq v_{stop}(z_L), \quad \limsup_{z \uparrow z_R} v(z) \leq v_{stop}(z_R). \quad (\text{OA.3-101})$$

Thus the one-sided limits equal the endpoint values, making v continuous and concave on $[z_L, z_R]$. Substituting the endpoint stopped policies into v_{stop} gives the two boundary values.

Finally, a concave function that is not strictly concave must be affine on a nondegenerate

interval. Suppose, toward a contradiction, that $v(z) = A + Bz$ on such an interval. If $B \neq 0$, then for any fixed $\beta > 0$ the Hamiltonian contains the term

$$B(r + \rho)\sigma^2\gamma r\beta\phi. \quad (\text{OA.3-102})$$

Since ϕ is unbounded and the curvature term is zero, the Hamiltonian is infinite, contradicting the viscosity supersolution inequality with the affine function as a lower test. If $B = 0$, the constant function is both an upper and a lower test at every interior point of the interval. The subsolution inequality (OA.3-83) and the supersolution inequality therefore imply

$$(r + \lambda)A = f(z), \quad (\text{OA.3-103})$$

which is impossible because f is strictly concave. Hence v is strictly concave on (z_L, z_R) . Continuity on the compact state space implies that v attains a maximum, and strict concavity makes the maximizer unique. Direct differentiation of v_{stop} gives

$$v'_{stop}(z_L) = \frac{z_R - z_L}{\theta(r + \lambda)} > 0, \quad v'_{stop}(z_R) = -\frac{\sigma^2\gamma r(r + \rho + \lambda)}{(r + \lambda)(r + \rho)} < 0. \quad (\text{OA.3-104})$$

Because $v \geq v_{stop}$ and equality holds at both endpoints, $v(z) > v(z_L)$ immediately to the right of z_L and $v(z) > v(z_R)$ immediately to the left of z_R . Neither endpoint can maximize v , so the unique maximizer satisfies $z^* \in (z_L, z_R)$.

OA.3.2 Policies

At points where v is twice differentiable, the HJB policy correspondence is the set of pointwise maximizers of the Hamiltonian. At nonsmooth points, we use the same term in the local dynamic-programming sense: a control belongs to the correspondence only if it is not strictly dominated by holding an alternative control fixed over an arbitrarily short localized horizon and then using the same ϵ -optimal continuation. This definition permits the policy signs to be characterized directly from concavity without presuming additional smoothness.

Throughout the remainder of this appendix,

$$f(z) = \frac{r + \rho + \lambda\zeta}{\theta(r + \lambda + \rho)}z - \frac{1}{2\theta}z^2 \quad (\text{OA.3-105})$$

denotes the flow payoff in the HJB equation (B-11), gross of the risk cost $\frac{\sigma^2\gamma r}{2}\beta^2$.

We first establish a useful result.

Lemma OA.1 *At the unique maximizer z^* ,*

$$(r + \lambda)v(z^*) = f(z^*), \quad (\text{OA.3-106})$$

$$D^-v(z^*) = 0, \quad (\text{OA.3-107})$$

where $D^-v(z^*)$ denotes the derivative from the left.

Proof. Fix a closed interval $J \subset (z_L, z_R)$ whose interior contains z^* . Strict concavity implies that z^* is the unique maximizer of v on J . Let v^n be the bounded-control value functions

defined in (OA.3-78) of Online Appendix Section OA.3, which also establishes the locally uniform convergence $v^n \rightarrow v$. This convergence implies that, for all sufficiently large n , any maximizer z_n of v^n on J lies in the interior of J and $z_n \rightarrow z^*$. The constant function with value $v^n(z_n)$ is an upper test for v^n at z_n . The viscosity subsolution inequality for the bounded Hamiltonian therefore gives

$$(r + \lambda)v^n(z_n) \leq \mathcal{H}^n(z_n, 0, 0) = f(z_n), \quad (\text{OA.3-108})$$

because the maximizing choice when both derivatives of the test function are zero is $\beta = 0$; the value is then independent of ϕ . Passing to the limit yields

$$(r + \lambda)v(z^*) \leq f(z^*). \quad (\text{OA.3-109})$$

Concavity implies that the finite left derivative $D^-v(z^*)$ exists, and maximality implies $D^-v(z^*) \geq 0$. Fix $z \in (z_L, z^*)$ and apply the feasible controls $\beta_s = \phi_s = 0$ until the state reaches z^* . Writing $a = r + \rho + \lambda$, the state follows

$$dz_s = a(z_s - z_L)ds \quad (\text{OA.3-110})$$

and reaches z^* at time

$$T(z) = \frac{1}{a} \log \left(\frac{z^* - z_L}{z - z_L} \right). \quad (\text{OA.3-111})$$

Concatenating this policy with an ϵ -optimal continuation at z^* , and then letting $\epsilon \downarrow 0$, gives

$$v(z) \geq \int_0^{T(z)} e^{-(r+\lambda)s} f(z_s) ds + e^{-(r+\lambda)T(z)} v(z^*). \quad (\text{OA.3-112})$$

As $z \uparrow z^*$,

$$T(z) = \frac{z^* - z}{a(z^* - z_L)} + o(z^* - z), \quad (\text{OA.3-113})$$

$$\int_0^{T(z)} e^{-(r+\lambda)s} f(z_s) ds = f(z^*)T(z) + o(T(z)). \quad (\text{OA.3-114})$$

Using the definition of the left derivative in the preceding inequality yields

$$(r + \lambda)v(z^*) - f(z^*) \geq a(z^* - z_L)D^-v(z^*). \quad (\text{OA.3-115})$$

The left side of (OA.3-115) is non-positive by (OA.3-109), whereas its right side is non-negative because $z^* > z_L$ and $D^-v(z^*) \geq 0$. Both sides must therefore equal zero. This proves (OA.3-106) and (OA.3-107). ■

Next we derive the policy characterization to the left of z^* .

Lemma OA.2 (Policies below z^*) *For every $z \in (z_L, z^*)$, the HJB Hamiltonian has the unique maximizer $\beta(z) = 0$, $\phi(z) = 0$.*

Proof. By Lemma OA.1, let $\widehat{v}$ solve

$$(r + \lambda)\widehat{v}(z) = f(z) + (r + \rho + \lambda)(z - z_L)\widehat{v}'(z), \quad (\text{OA.3-116})$$

$$\widehat{v}(z^*) = v(z^*) = \frac{f(z^*)}{r + \lambda} \quad (\text{OA.3-117})$$

on (z_L, z^*) . Writing $a = r + \rho + \lambda$ and $p = \widehat{v}'$, these equations imply $p(z^*) = 0$, and differentiation gives

$$0 = \frac{z_R - z}{\theta} + \rho p(z) + a(z - z_L)p'(z). \quad (\text{OA.3-118})$$

This linear equation has the representation

$$p(z) = \frac{1}{a\theta}(z - z_L)^{-\rho/a} \int_z^{z^*} (z_R - y)(y - z_L)^{\rho/a-1} dy. \quad (\text{OA.3-119})$$

Consequently, $p(z) > 0$ and (OA.3-118) gives $p'(z) < 0$ on (z_L, z^*) . The candidate is analytic on this interval and extends continuously to $\widehat{v}(z_L) = f(z_L)/(r + \lambda) = v(z_L)$.

We next verify that the candidate solves the full HJB equation. Differentiating (OA.3-118) and adding the result to that equation gives

$$0 = \frac{z_R - z}{\theta} - \frac{1}{\theta} + ap'(z) + \rho(p(z) + p'(z)) + a(z - z_L)(p'(z) + p''(z)). \quad (\text{OA.3-120})$$

At any zero of $p + p'$, the sum of the first two terms in (OA.3-120) is negative because $z_R - z < z_R - z_L = \frac{r+\rho}{r+\rho+\lambda} \leq 1$, and $ap' < 0$; hence $(p + p')' > 0$ at every such zero. But $p(z^*) + p'(z^*) = p'(z^*) < 0$. Therefore $p + p'$ cannot vanish on (z_L, z^*) , and

$$p(z) + p'(z) < 0. \quad (\text{OA.3-121})$$

For a fixed $\beta > 0$, maximization of the candidate Hamiltonian over ϕ gives

$$\phi^*(z, \beta) = \frac{p(z)}{-p'(z)} \frac{\gamma r \beta}{r + \rho}. \quad (\text{OA.3-122})$$

By (OA.3-121) and $\beta \leq 1$,

$$0 < \phi^*(z, \beta) < \frac{\gamma r}{r + \rho}. \quad (\text{OA.3-123})$$

Assumption 2 therefore implies

$$\sigma^2 \gamma r \phi^*(z, \beta) < \frac{\sigma^2 \gamma^2 r^2}{r + \rho} < 1. \quad (\text{OA.3-124})$$

After maximizing over ϕ , the envelope derivative of the Hamiltonian with respect to β is

$$-\sigma^2 \gamma r \beta + p(z)(r + \rho)(\sigma^2 \gamma r \phi^*(z, \beta) - 1) < 0 \quad (\text{OA.3-125})$$

for every $\beta > 0$. Thus $\beta = 0$ is the unique maximizer. Given $\beta = 0$ and $p'(z) < 0$, the Hamiltonian is uniquely maximized over ϕ at $\phi = 0$. Hence (OA.3-116) is the full HJB equation evaluated at its unique maximizing controls.

It remains to identify the candidate with the economic value function. Stop the state at its first exit from $(z_L + \epsilon, z^*)$. Itô's Lemma, the HJB inequality, and the boundary equalities imply $\hat{v}(z) \geq v(z)$ after letting $\epsilon \downarrow 0$. Conversely, under $\beta = \phi = 0$, the state follows $dz_s = a(z_s - z_L)ds$ and reaches z^* in finite time; integrating (OA.3-116) along this path and using an ϵ -optimal continuation at z^* gives the reverse inequality. Thus, $v = \hat{v}$ on $[z_L, z^*]$, so v is analytic on (z_L, z^*) and the displayed controls uniquely maximize its HJB Hamiltonian. ■

Finally we derive the policy characterization to the right of z^* .

Lemma OA.3 (Policies above z^*) *For every $z \in (z^*, z_R)$,*

1. *Every control in the HJB policy correspondence satisfies $\beta(z) > 0$ and $\phi(z) \leq 0$.*
2. *At every point where v is twice differentiable, which is Lebesgue-almost every $z \in (z^*, z_R)$, $v''(z) < 0$, and*

$$\phi(z) = -\frac{\gamma r v'(z) \beta(z)}{(r + \rho) v''(z)} < 0. \quad (\text{OA.3-126})$$

3. *There exists $z^{**} \in (z^*, z_R)$ such that $\beta(z) = 1$ for all $z \in (z^{**}, z_R)$.*

Proof. Write $R = r + \rho$ and $C = \sigma^2 \gamma r$. Strict concavity and the uniqueness of z^* imply that, for every $z \in (z^*, z_R)$,

$$D^+ v(z) \leq D^- v(z) < 0. \quad (\text{OA.3-127})$$

We first establish the policy signs without requiring differentiability. Fix $z \in (z^*, z_R)$, localize to a compact neighborhood, hold each compared control constant until a short horizon h or the first exit, and then use the same ϵ -optimal continuation. If $\beta = 0$, a nonzero ϕ adds a mean-preserving spread without changing the drift, so concavity of v and f makes $\phi = 0$ weakly preferable. Starting from $(\beta, \phi) = (0, 0)$, raising β to a small value $\epsilon > 0$ lowers the drift by $R\epsilon$. By (OA.3-127), the resulting continuation gain is of order $h\epsilon$, with a strictly positive coefficient, whereas the additional risk cost is $C h \epsilon^2 / 2$ and the change in the accumulated flow payoff caused by the state displacement is of smaller order in h . Thus $\beta = 0$ is strictly dominated for sufficiently small ϵ and h , so every control in the correspondence has $\beta(z) > 0$. If instead $\phi > 0$, replacing it by $-\phi$ leaves the risk cost and diffusion distribution unchanged but lowers the drift by $2R\sigma^2 \gamma r \beta \phi > 0$. Since both one-sided slopes of v are strictly negative, this raises the continuation value to first order in h , whereas the change in the accumulated flow payoff is of smaller order. Thus every control in the correspondence has $\phi(z) \leq 0$.

At a point where v is twice differentiable, write $p = v'(z) < 0$ and $X = v''(z) \leq 0$. Because $\beta > 0$, $X = 0$ would leave the nonzero linear term $p R \sigma^2 \gamma r \beta \phi$ in the HJB and make

its supremum over ϕ infinite. Thus $X < 0$, and the first-order condition gives

$$\phi(z) = -\frac{\gamma r v'(z) \beta(z)}{R v''(z)} < 0. \quad (\text{OA.3-128})$$

Concavity implies twice differentiability Lebesgue-almost everywhere, proving the second statement.

It remains to establish full PPS near the right boundary. Since $v \geq v_{stop}$ and the two functions agree at z_R , concavity gives

$$D^-v(z_R) \leq v'_{stop}(z_R) = -\frac{C(R + \lambda)}{(r + \lambda)R} < -\frac{C}{R}. \quad (\text{OA.3-129})$$

The one-sided derivatives of a concave function are decreasing, so both are less than $-C/R$ on $(z_R - \delta, z_R)$ for some $\delta > 0$. Fix a state in this interval and use the same short-horizon comparison as above. If $\beta < 1$, increasing it by a small $\epsilon > 0$ while holding ϕ fixed leaves the diffusion unchanged and lowers the drift by $R\epsilon(1 - \sigma^2\gamma r\phi) > 0$ because $\phi \leq 0$. By (OA.3-129), the continuation gain per unit of h exceeds $C\epsilon$, whereas the risk cost rises by $C\beta\epsilon + C\epsilon^2/2 < C\epsilon$ for sufficiently small ϵ ; the change in the accumulated flow payoff caused by the state displacement is of smaller order in h . Hence every control in the correspondence has $\beta = 1$ throughout this interval. Taking $z^{**} = z_R - \delta$ completes the proof. ■

Remark 2 (The full-PPS region) Write $R = r + \rho$, $C = \sigma^2\gamma r$, and $a = r + \rho + \lambda$. At every $z \in (z^*, z_R)$ at which v is twice differentiable, $v''(z) < 0$ by the second statement of Lemma OA.3; substituting the maximizing ϕ into the HJB equation (B-11) and using the first-order condition for β then eliminates $v''(z)$ and yields

$$\beta(z) = \min \left\{ 1, \frac{2((r + \lambda)v(z) - f(z) + |v'(z)|(az - \lambda\zeta))}{R|v'(z)|} \right\}, \quad (\text{OA.3-130})$$

with $\beta(z) = 1$ if and only if

$$\Psi(z) = f(z) - (r + \lambda)v(z) - |v'(z)| \left(az - \lambda\zeta - \frac{R}{2} \right) \leq 0. \quad (\text{OA.3-131})$$

Lemma OA.1 and the strict monotonicity of v and f on (z^*, z_R) imply $f(z) - (r + \lambda)v(z) > 0$ there. Two bounds on the full-PPS region follow. First, $\Psi(z) > 0$ whenever $az - \lambda\zeta \leq R/2$, so $\beta(z) < 1$ at every point of twice differentiability with $z \leq (\lambda\zeta + R/2)/a$. Second, because $1 + \gamma r(v'(z))^2/v''(z) < 1$, the corner condition for $\beta(z) = 1$ holds whenever $|v'(z)| \geq C/R$; by strict concavity, the argument establishing the third statement of Lemma OA.3 therefore applies on the entire interval where the one-sided derivatives of v lie below $-C/R$, so that z^{**} may be taken to be $\tilde{z} = \inf\{z \in (z^*, z_R) : D^+v(z) < -C/R\}$. $\{z : \beta(z) = 1\}$ is a single interval adjacent to z_R if and only if Ψ changes sign exactly once.

Numerically, we have so far found Ψ to change sign exactly once every time, meaning that the region in which $\beta(z) = 1$ is always a single interval. Proving such a result *ex ante* for any given set of parameters is infeasible given that $\Psi(z)$ is a function of endogenous properties of $v(z)$. Nevertheless, whether such a region is connected or not does not affect

the existence of the “serial founder” regime in Section 6.1. Such a regime is obtained with κ close to 1 and Lemma OA.3 guarantees a nonempty region with $\beta(z) = 1$ adjacent to z_R .

OA.3.3 A Long-Lived Principal Facing Agency Problems

We solve the model with a single long-lived principal. The agent’s lifetime incentives are those given in Proposition 2 with $\lambda = 0$. We will additionally set $\zeta = 0$ and $\bar{V} = 0$ so that the agent’s continuation utility is held at $W_t = -\frac{1}{\gamma r} \exp[-\gamma r S_t]$, which is analogous to the baseline problem’s (11). Then, (28) becomes

$$dM_t = \left(\frac{\theta}{2} \hat{h}_t^2 + \frac{\sigma^2 \gamma r}{2} \beta_t^2 \right) dt + \sigma \beta_t d\hat{Z}_t \quad (\text{OA.3-132})$$

The principal maximizes the expected discounted value of output subject to the results of Proposition 2. Starting with the principal’s objective function (8) for $\tau = \infty$ and re-deriving Proposition 4, we have

$$V^{\lambda=0}(H, z) = v^{\lambda=0}(z) + H \quad (\text{OA.3-133})$$

where

$$v^{\lambda=0}(z_t) = \max \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} \left(\frac{1}{\theta} z_s - \frac{1}{2\theta} z_s^2 - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right], \quad (\text{OA.3-134})$$

which is characterized for $z \in [0, 1]$ by the solution to (B-11), (B-9), and (B-10) with $\lambda = \zeta = \bar{V} = 0$.

OA.4 Additional Regularity

This section provides supplementary properties of the value function $v(z)$. In all our numerical simulations, the value function is strictly concave and smooth in the interior, with policy functions of the shapes described in the main text. The main analytical challenges arise from two points, z^* and z_R , because the controlled diffusion degenerates as $\phi \rightarrow 0$. At z^* , the policy dynamics transition from a deterministic region to a locally nondegenerate diffusion, while z_R is a degenerate state-constraint boundary. Consequently, standard uniformly elliptic interior-regularity and diffusion-existence results do not apply directly at these points. To develop stronger regularity and a classical feedback verification, throughout this section we impose the following technical assumption on the control problem. There exists $M < \infty$ such that, for every compact interval $J \subset (z^*, z_R)$, there exists $m_J > 0$ with the following property: for every initial state in J , restricting the admissible controls to satisfy

$$m_J \leq |\phi_s| \leq M \quad (\text{OA.4-135})$$

until the state first exits J does not change the value function v . The lower bound m_J may depend on J and may converge to zero as J approaches either z^* or z_R . Thus the assumption imposes nondegeneracy only on compact subsets of (z^*, z_R) and permits $\phi(z^*) = \phi(z_R) = 0$, which are part of the policy characterization. The common upper bound M rules

out unboundedly large $|\phi|$ without changing the localized value functions. The conclusions stated in Proposition 4 are established in the preceding sections. The condition above is used only for the stronger regularity and classical feedback-verification results developed in this section. Throughout this section, v remains the value of the original problem in which $\beta \in [0, 1]$ and $\phi \in \mathbb{R}$.

Below, we establish interior $C^{2,\alpha}$ regularity on (z^*, z_R) and differentiability at z^* , and use these properties to provide a classical feedback implementation and verification of the policy characterization.

Proposition OA.4 (Interior regularity of the unrestricted problem) *v is C^2 and locally $C^{2,\alpha}$ on (z^*, z_R) , where the exponent may depend on the compact subinterval. In particular, $v'(z) < 0$ and $v''(z) < 0$ throughout this interval, the unrestricted HJB holds pointwise there, and the policy conclusions of Lemma OA.3 hold everywhere rather than only Lebesgue-almost everywhere.*

Proof. Fix a compact interval $J \subset (z^*, z_R)$. The unrestricted HJB is

$$0 = \max_{\beta \in [0,1], \phi \in \mathbb{R}} \left\{ f(z) - \frac{\sigma^2 \gamma r}{2} \beta^2 - (r + \lambda)v(z) + v'(z)\mu(z, \beta, \phi) + \frac{1}{2}(r + \rho)^2 \sigma^2 \phi^2 v''(z) \right\}. \quad (\text{OA.4-136})$$

By Condition (OA.4-135), the exit-time problem on J with the compact control set

$$[0, 1] \times \{\phi : m_J \leq |\phi| \leq M\}, \quad (\text{OA.4-137})$$

running payoff $f(z) - \frac{\sigma^2 \gamma r}{2} \beta^2$, and terminal payoff v at the exit state has value function equal to v on J . For every control in this set, the coefficient multiplying v'' is bounded below by

$$\frac{1}{2}(r + \rho)^2 \sigma^2 m_J^2 > 0. \quad (\text{OA.4-138})$$

The localized Bellman equation is therefore uniformly elliptic. Its coefficients are smooth in the state and controls, its control set is compact, and its terminal data are continuous. The dynamic programming principle for this restricted exit-time problem implies that v is a continuous viscosity solution of the localized equation on the interior of J . The standard one-dimensional interior regularity theorem for uniformly elliptic Bellman equations with compact controls and smooth bounded coefficients then implies that $v \in C^{2,\alpha_J}(J)$ for some $\alpha_J \in (0, 1)$; see, for example, Fleming and Soner (2006). Since J was arbitrary, $v \in C^2((z^*, z_R))$.

Strict concavity and the uniqueness of z^* imply $v'(z) < 0$ on (z^*, z_R) , while concavity gives $v''(z) \leq 0$. Suppose $v''(z_0) = 0$ at an interior point. Since the HJB now holds pointwise, fix any $\beta > 0$. The terms involving ϕ at z_0 include

$$v'(z_0)(r + \rho)\sigma^2 \gamma r \beta \phi. \quad (\text{OA.4-139})$$

Because $v'(z_0) < 0$ and the original control set leaves ϕ unrestricted, the supremum over ϕ is infinite, contradicting the finite HJB value. Hence $v''(z) < 0$ throughout (z^*, z_R) .

The Hamiltonian can consequently be maximized directly at every interior state. Write $p = v'(z) < 0$ and $X = v''(z) < 0$. For fixed β , its unique maximizer over ϕ is

$$\phi(z, \beta) = -\frac{\gamma r p \beta}{(r + \rho)X}. \quad (\text{OA.4-140})$$

The right derivative of the Hamiltonian with respect to β at $\beta = 0$ is $-(r + \rho)p > 0$, so the maximizing β is strictly positive and the maximizing ϕ is strictly negative. Substituting the maximizing ϕ leaves the following terms that depend on β :

$$-(r + \rho)p\beta - \frac{\sigma^2 \gamma r}{2} \left(1 + \frac{\gamma r p^2}{X}\right) \beta^2. \quad (\text{OA.4-141})$$

Define

$$D(z) = 1 + \frac{\gamma r (v'(z))^2}{v''(z)}. \quad (\text{OA.4-142})$$

If $D(z) > 0$, the objective in (OA.4-141) is strictly concave and its unique maximizer on $[0, 1]$ is

$$\beta(z) = \min \left\{ 1, \frac{-(r + \rho)v'(z)}{\sigma^2 \gamma r D(z)} \right\}. \quad (\text{OA.4-143})$$

If $D(z) \leq 0$, the derivative of (OA.4-141) is strictly positive throughout $[0, 1]$, so its unique maximizer is $\beta(z) = 1$. In either case,

$$\phi(z) = -\frac{\gamma r v'(z) \beta(z)}{(r + \rho)v''(z)} < 0. \quad (\text{OA.4-144})$$

It remains to establish the policy near the right boundary. Since $v \geq v_{stop}$ and the two functions agree at z_R , concavity gives

$$\begin{aligned} D^-v(z_R) &\leq v'_{stop}(z_R) \\ &= -\frac{\sigma^2 \gamma r (r + \rho + \lambda)}{(r + \lambda)(r + \rho)} < -\frac{\sigma^2 \gamma r}{r + \rho}. \end{aligned} \quad (\text{OA.4-145})$$

The derivative of a concave function converges to its left boundary derivative. Hence there is $\delta > 0$ such that

$$v'(z) < -\frac{\sigma^2 \gamma r}{r + \rho} \quad (\text{OA.4-146})$$

for every $z \in (z_R - \delta, z_R)$. For any $\beta \in [0, 1]$, the maximizing $\phi(z, \beta)$ is non-positive. The envelope derivative of the Hamiltonian after maximizing over ϕ is therefore

$$-\sigma^2 \gamma r \beta + v'(z)(r + \rho) (\sigma^2 \gamma r \phi(z, \beta) - 1) > 0. \quad (\text{OA.4-147})$$

Thus $\beta(z) = 1$ and $\phi(z) < 0$ throughout $(z_R - \delta, z_R)$. This establishes the policy conclusions

pointwise without invoking Lemma OA.3. ■

OA.4.1 Feedback Construction

The policies in Lemma OA.3 were derived from the original unrestricted HJB. Because the restriction in (OA.4-135) leaves the value unchanged, the unique maximizing feedback of the unrestricted HJB lies in the restricted set and is therefore bounded by the common constant M .

We next combine the policy results into a single feedback rule. This step only constructs the coefficients of the candidate state process. The existence and verification of that process are established separately below.

Define

$$(\hat{\beta}(z), \hat{\phi}(z)) = \begin{cases} (0, 0), & z \in [z_L, z^*], \\ (\beta(z), \phi(z)), & z \in (z^*, z_R), \\ (1, 0), & z = z_R, \end{cases} \quad (\text{OA.4-148})$$

where $(\beta(z), \phi(z))$ is the unique pointwise maximizer characterized in Proposition OA.4.

Lemma OA.5 (Bounded candidate feedback) *The feedback in (OA.4-148) is Borel measurable on $[z_L, z_R]$ and continuous on every compact interval contained in (z^*, z_R) . It satisfies*

$$0 \leq \hat{\beta}(z) \leq 1, \quad |\hat{\phi}(z)| \leq M \quad (\text{OA.4-149})$$

for every $z \in [z_L, z_R]$. Consequently, the feedback drift $\mu(z, \hat{\beta}(z), \hat{\phi}(z))$ and diffusion coefficient $(r + \rho)\sigma\hat{\phi}(z)$ are bounded Borel functions on the state space and are continuous on (z^*, z_R) .

Proof. Proposition OA.4 implies that v' and v'' are continuous on (z^*, z_R) and that $v''(z) < 0$ there. The formula for $\phi(z)$ in (OA.4-144) is therefore continuous wherever $\beta(z)$ is continuous. After maximizing over ϕ , the objective in β is a continuous quadratic with a unique maximizer on the compact interval $[0, 1]$. The maximum theorem then implies that $\beta(z)$, and hence $\phi(z)$, is continuous on (z^*, z_R) . The piecewise definition in (OA.4-148) is consequently Borel measurable on the full state space.

The bound on $\hat{\beta}$ follows from its definition. In the interior, v solves the restricted and the unrestricted Bellman equations pointwise with the same value, so the compact restricted control set contains a maximizer of the unrestricted Hamiltonian; because that maximizer is unique, it equals $(\beta(z), \phi(z))$, and Condition (OA.4-135) gives $|\hat{\phi}(z)| \leq M$ in the interior. The assigned values are zero elsewhere. Since the state space is compact and μ is continuous in (z, β, ϕ) , substituting the bounded Borel feedback gives bounded Borel drift and diffusion coefficients. Their continuity on (z^*, z_R) follows from the interior continuity of the feedback. ■

OA.4.2 Admissibility

The only junction not covered by interior regularity is z^* . We first establish smooth fit there and then verify that the feedback coefficients generate a viable state process.

Lemma OA.6 (Right smooth fit) *The value function is differentiable at z^* and*

$$D^-v(z^*) = D^+v(z^*) = 0. \quad (\text{OA.4-150})$$

Proof. Lemma OA.1 gives

$$(r + \lambda)v(z^*) = f(z^*), \quad D^-v(z^*) = 0. \quad (\text{OA.4-151})$$

Concavity and maximality imply $D^+v(z^*) \leq 0$. Suppose instead that $q := D^+v(z^*) < 0$. Write $b^* = \beta_{stop}(z^*) \in (0, 1)$ and fix $p \in (q, 0)$ sufficiently close to zero that

$$\beta_p = -\frac{p(r + \rho)}{\sigma^2 \gamma r} \in (0, 1), \quad p > -\frac{2\sigma^2 \gamma r}{r + \rho} b^*. \quad (\text{OA.4-152})$$

For every sufficiently large K , the function

$$\psi_K(z) = v(z^*) + p(z - z^*) - \frac{K}{2}(z - z^*)^2 \quad (\text{OA.4-153})$$

touches v from above at z^* on a sufficiently small neighborhood. The viscosity subsolution inequality, value matching, and maximization over ϕ give

$$0 \leq p(r + \rho)\beta_{stop}(z^*) + \max_{\beta \in [0, 1]} \left\{ -p(r + \rho)\beta - \frac{\sigma^2 \gamma r}{2}\beta^2 + \frac{\sigma^2 \gamma^2 r^2 p^2}{2K}\beta^2 \right\}. \quad (\text{OA.4-154})$$

Letting $K \rightarrow \infty$ yields

$$0 \leq p(r + \rho)b^* + \max_{\beta \in [0, 1]} \left\{ -p(r + \rho)\beta - \frac{\sigma^2 \gamma r}{2}\beta^2 \right\}, \quad (\text{OA.4-155})$$

By the choice of p , the maximum in (OA.4-155) is attained at β_p , and its right side equals

$$p(r + \rho) \left(b^* + \frac{p(r + \rho)}{2\sigma^2 \gamma r} \right) < 0, \quad (\text{OA.4-156})$$

a contradiction. Thus $D^+v(z^*) = 0$. ■

Lemma OA.7 (Behavior at the junctions) *The candidate feedback satisfies*

$$\lim_{z \downarrow z^*} \widehat{\beta}(z) = 0, \quad \lim_{z \uparrow z_R} \widehat{\phi}(z) = 0. \quad (\text{OA.4-157})$$

Moreover,

$$\lim_{z \downarrow z^*} \mu(z, \widehat{\beta}(z), \widehat{\phi}(z)) = (r + \rho)\beta_{stop}(z^*) > 0, \quad (\text{OA.4-158})$$

$$\lim_{z \uparrow z_R} \mu(z, \widehat{\beta}(z), \widehat{\phi}(z)) = 0. \quad (\text{OA.4-159})$$

Proof. Write $p = v'(z) < 0$ and $X = v''(z) < 0$ on (z^*, z_R) . By Lemma OA.6, $p \rightarrow 0$ as $z \downarrow z^*$. The feedback and its drift are bounded by Lemma OA.5. The first-order condition

for ϕ gives

$$X(r + \rho)^2 \sigma^2 \widehat{\phi}(z)^2 = -p(r + \rho) \sigma^2 \gamma r \widehat{\beta}(z) \widehat{\phi}(z). \quad (\text{OA.4-160})$$

Both the drift term $p\mu$ and the curvature term in the HJB therefore converge to zero as $z \downarrow z^*$. Since $(r + \lambda)v(z^*) = f(z^*)$, the pointwise HJB then implies

$$\frac{\sigma^2 \gamma r}{2} \widehat{\beta}(z)^2 \rightarrow 0. \quad (\text{OA.4-161})$$

This proves the first limit in (OA.4-157). Boundedness of $\widehat{\phi}$ also gives $\widehat{\beta}\widehat{\phi} \rightarrow 0$, which proves (OA.4-158).

For the right boundary, Proposition OA.4 gives $\widehat{\beta}(z) = 1$ on a left neighborhood of z_R . Set $u = -v' > 0$, $R = r + \rho$, $C = \sigma^2 \gamma r$, and $k = \sigma^2 \gamma^2 r^2$. The HJB and the first-order condition for ϕ give

$$(r + \lambda)v(z) - f(z) + \frac{C}{2} = Ru(z)(1 - \beta_{\text{stop}}(z)) + \frac{k u(z)^2}{2|v''(z)|}. \quad (\text{OA.4-162})$$

Both terms on the right are non-negative. The boundary condition and $\beta_{\text{stop}}(z_R) = 1$ imply that the left side converges to zero. Hence $u(z)^2/|v''(z)| \rightarrow 0$. Inequality (OA.3-129) bounds u away from zero near z_R , and therefore

$$|\widehat{\phi}(z)| = \frac{\gamma r}{R} \frac{u(z)}{|v''(z)|} \rightarrow 0. \quad (\text{OA.4-163})$$

The definition of μ now gives (OA.4-159). ■

Proposition OA.8 (Viable candidate process) *For every initial state $z \in [z_L, z_R]$ there is a weak solution to*

$$d\widehat{z}_s = \mu(\widehat{z}_s, \widehat{\beta}(\widehat{z}_s), \widehat{\phi}(\widehat{z}_s))ds + (r + \rho)\sigma\widehat{\phi}(\widehat{z}_s)dZ_s \quad (\text{OA.4-164})$$

that remains in $[z_L, z_R]$. The associated feedback control belongs to $\mathcal{A}(z)$.

Proof. On every compact interval contained in (z^*, z_R) , Lemma OA.5 gives continuous coefficients and (OA.4-135) bounds the diffusion coefficient away from zero. Standard one-dimensional weak-existence results therefore give a solution up to the first exit from the interval. On $[z_L, z^*]$, the feedback is $(0, 0)$ and the state follows the ordinary differential equation

$$d\widehat{z}_s = (r + \rho + \lambda)(\widehat{z}_s - z_L)ds. \quad (\text{OA.4-165})$$

Thus z_L is absorbing and every state in $(z_L, z^*]$ moves deterministically toward the interior.

At z_R , both state coefficients vanish. Lemma OA.7 shows that they converge to zero from the left, so a solution that reaches z_R can be continued there as the constant process. The state therefore cannot cross the right boundary.

It remains to join the deterministic and stochastic solutions at z^* . The coefficients are bounded, the drift converges from both sides to the strictly positive value in (OA.4-158), and

the diffusion is locally nonzero immediately to the right of z^* . Approximate the coefficients on an interval of width $1/n$ around z^* by bounded continuous coefficients that agree with the candidate coefficients outside that interval and preserve the inward drift. The corresponding weak solutions exist and their laws are tight on every finite horizon. Any weak limit solves (OA.4-164) away from z^* . It remains to rule out a sticky term at the junction. On a fixed neighborhood of z^* , the approximating drifts are bounded below by some $b_0 > 0$, and the diffusion coefficients are bounded. Hence, before any fixed horizon, the expected time the approximating process spends in $(z^* - \epsilon, z^* + \epsilon)$ is $O(\epsilon)$, uniformly in n . Letting $\epsilon \downarrow 0$ implies that every weak limit spends zero Lebesgue time at z^* . Consequently, the value assigned to the coefficients at this single state does not alter the limiting martingale problem, and the limit solves (OA.4-164) on the entire state space.

Combining the interior solutions with the absorbing continuations at the endpoints proves viability. Finally, the feedback is progressively measurable, $0 \leq \hat{\beta} \leq 1$, $|\hat{\phi}| \leq M$, and the state remains in $[z_L, z_R]$. The feedback control therefore belongs to $\mathcal{A}^n$ for every integer $n \geq M$, and hence to $\mathcal{A}(z)$. ■

OA.4.3 Verification

Theorem OA.9 (Feedback Verification) *The feedback control in (OA.4-148) attains the principal's value for every initial state $z \in [z_L, z_R]$. In particular,*

$$v(z) = \max_{(\beta, \phi) \in \mathcal{A}(z)} \mathbb{E}_z \left[\int_0^\infty e^{-(r+\lambda)s} \left(f(z_s) - \frac{\sigma^2 \gamma r}{2} \beta_s^2 \right) ds \right]. \quad (\text{OA.4-166})$$

Proof. Write $\delta = r + \lambda$, $C = \sigma^2 \gamma r$, and

$$a(z) = (r + \rho)\sigma\hat{\phi}(z), \quad m(z) = \mu(z, \hat{\beta}(z), \hat{\phi}(z)). \quad (\text{OA.4-167})$$

The value function is analytic on (z_L, z^*) by Lemma OA.2 and belongs to $C^2((z^*, z_R))$ by Proposition OA.4. Concavity implies that the distributional second derivative of v is a locally finite measure. The stated regularity eliminates any singular component on either side of z^* , while Lemma OA.6 eliminates an atom at z^* . Hence v' is locally absolutely continuous across z^* , and the generalized Itô formula applies without a local-time term there.

Fix a finite horizon T . First stop the candidate process on compact subintervals of (z_L, z_R) and localize its stochastic integral. Applying the generalized Itô formula to $e^{-\delta s}v(\hat{z}_s)$ gives

$$\mathbb{E}_z [e^{-\delta T}v(\hat{z}_T)] = v(z) + \mathbb{E}_z \left[\int_0^T e^{-\delta s} \left(-\delta v(\hat{z}_s) + v'(\hat{z}_s)m(\hat{z}_s) + \frac{1}{2}v''(\hat{z}_s)a(\hat{z}_s)^2 \right) ds \right], \quad (\text{OA.4-168})$$

after removing the localization. At z^* the process spends zero Lebesgue time by the proof of Proposition OA.8, so the value of the integrand at that single state is immaterial.

The candidate feedback attains the pointwise HJB maximum throughout (z_L, z^*) by Lemma OA.2 and throughout (z^*, z_R) by Proposition OA.4. Hence, for almost every time

before an endpoint is reached,

$$\delta v(\widehat{z}_s) = f(\widehat{z}_s) - \frac{C}{2}\widehat{\beta}(\widehat{z}_s)^2 + v'(\widehat{z}_s)m(\widehat{z}_s) + \frac{1}{2}v''(\widehat{z}_s)a(\widehat{z}_s)^2. \quad (\text{OA.4-169})$$

At z_L the state is absorbed under $(\widehat{\beta}, \widehat{\phi}) = (0, 0)$ and the boundary condition gives $\delta v(z_L) = f(z_L)$. At z_R it is absorbed under $(\widehat{\beta}, \widehat{\phi}) = (1, 0)$ and the boundary condition gives $\delta v(z_R) = f(z_R) - C/2$. Thus the same finite-horizon identity continues after either endpoint is reached.

Substituting (OA.4-169) into (OA.4-168) yields

$$v(z) = \mathbb{E}_z \left[\int_0^T e^{-\delta s} \left(f(\widehat{z}_s) - \frac{C}{2}\widehat{\beta}(\widehat{z}_s)^2 \right) ds + e^{-\delta T} v(\widehat{z}_T) \right]. \quad (\text{OA.4-170})$$

To justify removal of the compact-state localization, let the compact intervals increase to (z_L, z_R) . Continuity and boundedness of v control the terminal terms, while boundedness of f and $\widehat{\beta}$ permits dominated convergence for the running payoff. Lemma OA.7 and endpoint absorption give the stated continuation when the limiting process reaches a boundary.

Because v is bounded on the compact state space and $\delta > 0$,

$$\lim_{T \rightarrow \infty} \mathbb{E}_z [e^{-\delta T} v(\widehat{z}_T)] = 0. \quad (\text{OA.4-171})$$

Letting $T \rightarrow \infty$ in (OA.4-170) gives

$$v(z) = \mathbb{E}_z \left[\int_0^\infty e^{-\delta s} \left(f(\widehat{z}_s) - \frac{C}{2}\widehat{\beta}(\widehat{z}_s)^2 \right) ds \right]. \quad (\text{OA.4-172})$$

Proposition OA.8 shows that this feedback belongs to $\mathcal{A}(z)$. Equation (OA.4-172) therefore shows that it attains the unrestricted supremum. This proves (OA.4-166). ■