Optimal Long-Term Contracting with Learning

Zhiguo He
Booth School of Business, University of Chicago and NBER

Bin Wei
Federal Reserve Bank of Atlanta

Jianfeng Yu
PBC School of Finance, Tsinghua University and Carlson School of
Management, University of Minnesota

Feng Gao
School of Economics and Management, Tsinghua University

We introduce uncertainty into Holmstrom and Milgrom (1987) to study optimal long-term contracting with learning. In a dynamic relationship, the agent's shirking not only reduces current performance, but also increases the agent's information rent due to the persistent belief manipulation effect. We characterize the optimal contract using the dynamic programming technique in which information rent is the unique state variable. In the optimal contract, the optimal effort is front-loaded and stochastically decreases over time. Furthermore, the optimal contract exhibits an option-like feature in that incentives increase after good performance. Implications about managerial incentives and asset management compensations are discussed. (*JEL* D82, D83, E24, J41)

Many long-term contractual relationships feature learning, because uncertainty arises if either project quality or agent ability is unknown when a long-term contract is signed. Dynamic learning is most relevant for venture capital firms

investing in companies with new technologies or firms hiring fresh graduates. Unfortunately, for reasons stated later, the study of long-term contracting with learning is challenging.

We introduce uncertainty and learning into the classic Holmstrom and Milgrom (1987) model with a constant absolute risk aversion (CARA) agent. We choose Holmstrom and Milgrom (1987) for two reasons. First, the Holmstrom and Milgrom (1987) model has a tractable dynamic CARA-normal framework that nicely accommodates learning, and we consider an infinite-horizon variation of Holmstrom and Milgrom (1987) with stationary learning to maintain tractability. Second, against the Holmstrom and Milgrom (1987) benchmark in which the optimal contract is linear, we show that uncertainty and learning make the optimal compensation contract option-like; that is, incentives rise following good performance.

In our model the principal signs a long-term contract with the agent, with commitment by both parties. The observable output each period is the sum of the agent's unobservable effort, the project's unknown profitability (or the agent's unknown ability), and some transitory noise. To focus on learning only (rather than adverse selection), we assume that both the principal and agent share a common prior on the project's profitability when signing the long-term contract.

Unlike Holmstrom and Milgrom (1987), incentive provisions become intertemporally linked over time because of learning. The intertemporal linkage of incentive provisions is rooted in the *hidden information* problem.[1] Along the equilibrium path, the principal knows as much as the agent knows, because both start with the common prior. However, along off-equilibrium paths, the agent strictly knows more, because only the agent knows how much actual efforts deviate from the recommended level of effor. Specifically, imagine that the agent has followed the recommended effort policy in the past; thus both parties share the same correct belief about the project's profitability. If the agent shirks today by exerting some effort below the recommended level, then the lower effort decreases today's output on average. With Bayesian learning, the principal who anticipates a higher effort today would mistakenly attribute today's weak performance to lower profitability. Thus, by shirking today the agent can distort downward the principal's inference about profitability from today onward, which is long-lasting (i.e., persistent hidden information). This *belief manipulation* effect is beneficial to the agent, as the principal will mistakenly reward the agent whenever future performance beats the principal's downward distorted expectations. We refer to this potential benefit due to off-equilibrium private information as the agent's *information rent*.

---

[1] This is in contrast to the standard *hidden action* dynamic agency models in which the agent's unobservable shirking has only a *short-lived* effect. For recent development of dynamic contracting in finance, see DeMarzo and Fishman (2007), Biais et al. (2007), DeMarzo and Sannikov (2006), He (2009), Piskorski and Tchistyi (2010), DeMarzo et al. (2012), and Malenko (2013), among others.

In solving the optimal contract with learning, we need the information rent as the second state variable, in addition to the agent's continuation value. The information rent captures the marginal benefit of the agent's shirking due to the belief manipulation effect, and hence enters the agent's incentive compatibility constraint. The higher the future incentives (i.e., pay–performance sensitivity), the greater the information rent, and the lower the agent's current motivation to expend effort. We show that the information rent can be conveniently expressed as the sum of properly discounted future incentives, and the agent's optimal effort is simply the instantaneous incentive minus the information rent due to the belief manipulation effect. Thanks to the CARA preference that has no wealth effect, the agent's continuation value separates from the problem, and the optimal contract is fully characterized by an ordinary differential equation (ODE) with the information rent as the only state variable. Although we use the first-order approach to solve for the optimal contract, we verify the validity of the first-order approach in Section 3.4 by identifying an upper bound of the agent's deviation value. Section 3.5 discusses how CARA preferences and private savings render the tractability in our model.

Relative to the existing literature of long-term contracting with learning, which focuses on implementing a constant first-best effort (DeMarzo and Sannikov 2017; Prat and Jovanovic 2014), our paper highlights two interesting features in the optimal contract. First, in our model the optimal effort policy, which is always distorted downward relative to the first-best benchmark, has a negative drift, thus exhibiting a front-loaded or time-decreasing pattern. This is somewhat surprising. We have explained that under a given contract the information rent makes the agent want to work less in earlier periods, and casual readers might conclude that in the optimal contract the agent should work *less* earlier. However, the opposite holds in the optimal contract: The principal will purposefully give higher incentives early on so that the agent works *harder* in earlier periods in equilibrium.

In Section 4.1, we solve in closed form the optimal deterministic contract (i.e., the optimal one among the contracts in the subspace that implements deterministic but time-varying incentives only), and show analytically that the optimal deterministic effort policy decreases over time. This pattern holds in the optimal stochastic contract, and the intuition is a result of the belief manipulation effect. As mentioned, later incentives increase the agent's current information rent for shirking. This implies that future pay–performance sensitivities impair the agent's motivation for expending effort in earlier periods, but not the other way around. Given that later incentives are more costly, the optimal contract implements less effort in later periods.[2]

---

[2] Interestingly, the pattern of time-decreasing effort policy in our paper with *post-contracting* information asymmetry is opposite of the dynamic contracting setting with *pre-contracting* asymmetric information in Garrett and Pavan (2012). In that paper, under the assumptions that the agent privately observes his productivity at the time of signing the contract and that the effect of initial productivity on future productivity is declining

Second, the optimal contract is stochastic with higher incentives after good performance, exhibiting an *option-like* feature.[3] The intuition is the result of reducing the agent's belief manipulation in a long-term relationship. For a risk-averse agent, the amount of information rent not only depends on the benefits of belief manipulation that increase with future pay–performance sensitivities, but also the agent's marginal utility at future states when receiving those benefits. Raising incentives after good performance introduces a negative correlation between pay-for-performance and marginal utility. That is, greater future benefits from belief manipulation are associated with the states when the agent cares less. Hence, the option-like compensation contract lowers the agent's information rent standing today.

The combination of long-term contracting and learning that drives front-loaded and option-like incentives. On the one hand, with long-term contracting but no learning, the model is a simple extension of Holmstrom and Milgrom (1987) and a constant effort policy is optimal (Section 3.3). On the other hand, with learning but short-term contracting, the absence of commitment due to the nature of short-term contracting relationships implies that principals at different times will not take the aforementioned belief manipulation effect into account. In that case, similar to Holmstrom (1999), the Gaussian setting with stationary Bayesian learning gives rise to a constant effort process in equilibrium (Section 4.4).

We rely on specific assumptions (i.e., CARA preferences, private savings, stationary Gaussian setting) to fully characterize the optimal long-term contract with learning. However, the economic forces that are driving our main results do not depend on CARA preferences or Gaussian processes. First, in any long-term contracting environment with learning, it is generally true that the agent obtains information rent due to belief manipulation, which captures his desire to shirk so as to distort the principal's future belief downward. This result implies that later incentives enter the agent's forward-looking information rent in earlier periods (but not the other way around). Consequently, later incentives are more costly than earlier ones, giving rise to the time-decreasing effort policy. Second, the option-like feature relies solely on the concavity of the agent's utility function, so that the marginal value of earning future (potential) belief manipulation benefit is lower for the agent after good performance; hence higher

---

over time, the optimal effort policy is time-increasing. Intuitively, in Garrett and Pavan (2012), the downward distortion required for rent extraction is more severe in earlier periods when the major friction is pre-contracting private information. It is intriguing that pre-contracting private information and post-contracting information have opposite predictions for the time-series pattern of effort distortion, but the difference in Garrett and Pavan (2012) also lies in the agent being risk neutral without wealth constraint. Relatedly, Sannikov (2014) allows the agent's current effort to affect future fundamental, and Marinovic and Varas (2016) study the optimal contract when the agent can engage in manipulation to boost short-term performance, but with negative long-term consequences.

[3] That effort policy is history-dependent is surprising given our setting. With a standard CARA-normal setting and learning, as the posterior variance only changes deterministically over time (in our stationary setting, it is a constant), the resultant equilibrium effort profile is usually deterministic (e.g., Gibbons and Murphy 1992; Holmstrom 1999). In contrast, in our model with learning, the optimal long-term contract has an option-like feature in that pay-for-performance rises following good performance.

compensation. Since these economic forces are fairly general, our two main qualitative results–front-loaded effort policy and option-like compensation–are likely robust to other more general settings.

Our model offers some interesting empirical implications. In particular, it provides a mechanism that demonstrates why option-like payoffs are desirable in managerial compensation. In practice the use of option-based compensation is no doubt pervasive.[4] Interestingly, traditional static models typically do not predict option grants.[5] For example, Dittmann and Maug (2007) calibrate a standard static structural model and find that most CEOs should hold more straight equity, hold no stock options, and receive lower salaries. The option-like features of the optimal contract in our paper shed light on the "2-20" and high-water-mark contracts that are widely used in the hedge fund industry. As shown in our paper, that hedge fund contracts exhibit option-like features may well be due to learning about persistent unobservable managerial ability as well as the commitment associated with long-term contracting in the hedge fund industry. In addition, our model also predicts that industries with higher uncertainty should grant more stock options to their managers. The latter cross-sectional prediction is consistent with the evidence in Ittner, Lambert, and Larcker (2003) and Murphy (2003), who document more extensive use of stock options in new-economy firms (e.g., computer-related firms).

Our paper is closest to DeMarzo and Sannikov (2017) and Prat and Jovanovic (2014). As mentioned earlier, both papers deal with long-lasting belief manipulation effect in dynamic agency settings with learning, but restrict attention to the optimal contract that implements a constant first-best level of effort. Prat and Jovanovic (2014) focus on the role of intertemporal commitment in optimal contracting. DeMarzo and Sannikov (2017) impose limited liability constraint on the agent and study the optimal payout and termination policies. In contrast, we solve for the optimal effort policy jointly with the optimal long-term compensation contract and emphasize the general economic mechanisms that shape the optimal effort policy in long-term optimal contracting. As discussed previously, the two main results of our optimal contract, that is, front-loaded effort policy and option-like incentives, cannot hold in the contracting space with constant effort policy.

The long-lasting belief manipulation in dynamic contracting also exists in Bergemann and Hege (1998) and Horner and Samuelson (2013). In Bergemann and Hege (1998), an agent keeps working on a project which may succeed with some probability depending on its quality, and the game ends once the project

---

[4] Hall and Liebman (1998), for example, show a large increase in the use of stock options in CEO compensation for incentive provisions.

[5] There are a few exceptions in a dynamic framework. For instance, Edmans and Gabaix (2011) show that the convexity of the contract depends on the marginal cost of effort. In Ju and Wan (2012), stock options become optimal when the agent has to be paid above a certain subsistence level.

succeeds.[6] The project quality and the agent's effort affect the project success in a multiplicative way, that is, success may occur *only if* the project is good *and* the agent is working. In contrast, our model features an additive production function in which the marginal productivity of effort is independent of the project quality.

The topic of optimal contracting with endogenous learning also relates to the recent literature studying optimal long-term contracts with adverse selection and moral hazard (e.g., Baron and Besanko 1984; Sung 2005; Sannikov 2007; Garrett and Pavan 2012; Gershkov and Perry 2012; Halac, Kartik, and Liu 2016; Cvitanic, Wan, and Yang 2013).[7] In general, when the agent has pre-contracting private information that is persistent, a mechanism design approach naturally arises (e.g., Pavan, Segal, and Toikka 2014; Golosov, Troshkin, and Tsyvinski 2012).[8] However, because our paper focuses on the problem without pre-contracting private information, we do not need to solve for the optimal menu for the agent's truthful reporting when signing the contract.

1. Model

1.1 Setting

Consider a continuous-time infinite-horizon principal-agent model with a common constant discount rate $r > 0$. The project generates a cumulative output Y_t up to time t, which evolves according to

$$dY_t = (\mu_t + \theta_t)dt + \sigma dB_t, \tag{1}$$

where $\{B_t\}$ is a standard Brownian motion on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$, μ_t is the agent's unobservable effort level, θ_t is the project's profitability, and the constant $\sigma > 0$ is the volatility of cash flows. Moral hazard arises from the agent's unobservable effort choice, which affects the instantaneous cash-flow dY_t.

---

[6] This assumption is crucial for the tractability of Bergemann and Hege (1998). It is worth noting that the "real option" mentioned in the abstract of Bergemann and Hege (1998) is different from our result. In our paper, "option-like incentives" refer to the fact that incentives rise after good performance; but in their paper the game ends after any good performance (i.e., project success).

[7] Other papers that are related to learning but do not deal with the belief manipulation effect. Adrian and Westerfield (2009) focus on the disagreement between the principal and the agent about the agent's ability, where the agent is dogmatic about his belief (i.e., the agent never updates his posterior belief about profitability from past performance), which eliminates the belief manipulation effect. In that paper, although the agent could distort the principal's belief by shirking, the dogmatic agent (who does not realize that the firm's profitability is, in fact, higher than that perceived by the principal) will not gain anything from this channel, and as a result there is no belief manipulation effect. More recently, Cosimano, Speight, and Yun (2011) study the long-term contracting problem with binary unobservable productivity states, and show that the optimal contract tends to be sticky. They assume that the agent's effort is observable but not contractible, and hence both the principal and the agent always have the same information set, on both equilibrium and off-equilibrium paths.

[8] Pavan, Segal, and Toikka (2014) and Golosov, Troshkin, and Tsyvinski (2012) use the first-order approach to solve the agent's problem. This is the same approach used in Williams (2009, 2011) and Zhang (2009), who study persistent information in a continuous-time principal-agent setting. We also use the first-order approach to solve the agent's problem and verify the validity of the first-order approach in Section 2.4.

The risk-neutral principal (hereinafter *she*) offers the CARA agent (hereinafter *he*) a contract $\{c_t, \mu_t\}$, so that the agent is recommended to take the effort policy $\mu = \{\mu_t\}$ and is compensated by the wage process $c = \{c_t\}$. Both elements are measurable to $\mathcal{Y}_t \equiv \mathcal{F}\{Y_s : 0 \le s \le t\}$, which is the filtration generated by the output history. Both parties can commit to the long-term relationship at $t = 0$, at which point the agent has no personal wealth and has an exogenous reservation utility of v_0. Without loss of generality, we assume the principal has all the bargaining power.

Relative to Holmstrom and Milgrom (1987), we introduce the project's *unknown* profitability θ_t into the output process in Equation (1). Equivalently, one can interpret θ_t as the agent's unknown ability. We assume that profitability $\{\theta_t\}$ follows a martingale process so that

$$d\theta_t = \phi \sigma \, dB_t^\theta,$$

where the Brownian motion $\{B^\theta\}$ is independent of $\{B\}$, and $\phi > 0$ is a constant. At time 0, the principal and the agent share the common normal prior: $\theta_0 \sim \mathcal{N}(m_0, \Sigma_0^\theta)$. We mainly focus on stationary learning; we discuss nonstationary learning for robustness checks in the Internet Appendix. For learning to be stationary, the prior uncertainty is assumed to satisfy $\Sigma_0^\theta = \sigma^2 \phi$, so that the posterior variance $\Sigma_t^\theta = \Sigma_0^\theta$ for all t and Bayesian updating is time independent. When $\phi = 0$, our model features no uncertainty (or, θ_t is perfectly observable), and thus is reduced to the benchmark model of Holmstrom and Milgrom (1987), as analyzed in Section 3.3.

We further assume that the agent can privately save (i.e., hidden savings, or consumption is not contractible) to smooth his consumption intertemporally, if he wishes. CARA preferences do not have a wealth effect, and the issue of private savings can be easily dealt with (e.g., Fudenberg, Holmstrom, and Milgrom 1990; Williams 2009; He 2011). In Section 3.5, we explain the reason why the agent's ability to smooth his own consumption renders extra tractability for this model.

Private savings imply that the agent's actual consumption can differ from wage c_t. The agent's actual consumption is represented by $\widehat{c}_t$ and actual effort by $\widehat{\mu}_t$; then the agent with a CARA preference (exponential utility) has an instantaneous utility of

$$u(\widehat{c}_t, \widehat{\mu}_t) = -\frac{1}{a} \exp[-a(\widehat{c}_t - g(\widehat{\mu}_t))],$$

where $a > 0$ is the agent's absolute risk-aversion coefficient, and $g(\widehat{\mu}_t) \equiv \frac{1}{2}\widehat{\mu}_t^2$ is the instantaneous quadratic monetary cost of exerting effort $\widehat{\mu}_t$.[9] The quadratic form of $g(\cdot)$ simplifies our results, but our analysis holds as long as $g(\cdot)$ is strictly increasing and strictly convex.

---

[9] In the tradition of Holmstrom and Milgrom (1987), the CARA preference allows for negative consumption; that is, both c_t and $\widehat{c}_t$ can take negative values. In contrast, in DeMarzo and Sannikov (2017) the agent is protected by limited liabilities, and hence the endogenous contract termination arises. It is unclear how the limited-liability restriction affects the qualitative results of our paper.

1.2 Bayesian learning and effort

Recall that at time 0, the principal and the agent share the common normal prior $\theta_0 \sim \mathcal{N}\left(m_0, \Sigma_0^\theta\right)$. From now on we normalize $m_0 = 0$. Both parties update their beliefs based on their own respective information sets. Recall that $\mathcal{Y}_t = \mathcal{F}\{Y_s : 0 \le s \le t\}$ is the augmented filtration generated by output path Y. Given any contract $\{c_t, \mu_t\}$, the principal's information set at time t is $\mathcal{F}\{Y_s, \mu_s : 0 \le s \le t\}$, as the principal knows the recommended effort policy $\mu \equiv \{\mu_t\}$. However, the agent's information set also includes his actual effort policy $\widehat{\mu} \equiv \{\widehat{\mu}_t\}$, that is, $\mathcal{F}\{Y_s, \mu_s, \widehat{\mu}_s : 0 \le s \le t\}$. Intuitively, relative to the principal, the agent knows (weakly) more because he knows his actual past effort choices $\widehat{\mu}$, which may deviate from the recommended policy μ. This distinction is important for our analysis.

If the agent follows the recommended effort policy μ, the principal's posterior belief about θ_t is correct and fully summarized by the first two moments:

$$m_t^\mu \equiv \mathbb{E}[\theta_t \mid \mathcal{Y}_t, \mu] \text{ and } \Sigma_t^{\theta,\mu} \equiv \mathbb{E}\left[\left(\theta_t - m_t^\mu\right)^2 \mid \mathcal{Y}_t, \mu\right].$$

A standard filtering argument (e.g., Theorem 12.2 in Liptser and Shiryaev 1977) implies that $\Sigma_t^{\theta,\mu} = \sigma^2 \phi$ for all t (due to the stationary assumption $\Sigma_0^\theta = \sigma^2 \phi$), and

$$dm_t^\mu = \Sigma_t^{\theta,\mu} \frac{dY_t - \left(\mu_t + m_t^\mu\right)dt}{\sigma^2} = \sigma \phi \, dB_t^\mu \text{ with } m_0 = 0, \tag{2}$$

where B_t^μ is a standard Brownian motion under the measure induced by the effort policy μ:

$$dB_t^\mu \equiv \frac{dY_t - \left(\mu_t + m_t^\mu\right)dt}{\sigma}. \tag{3}$$

Conditional on the actual effort policy $\{\widehat{\mu}_t\}$, the agent forms his posterior belief as

$$m_t^{\widehat{\mu}} \equiv \mathbb{E}[\theta_t \mid \mathcal{Y}_t, \widehat{\mu}] \text{ and } \Sigma_t^{\theta,\widehat{\mu}} \equiv \mathbb{E}\left[\left(\theta_t - m_t^{\widehat{\mu}}\right)^2 \mid \mathcal{Y}_t, \widehat{\mu}\right].$$

The superscript $\widehat{\mu}$ emphasizes the dependence on the agent's actual effort policy $\widehat{\mu}$ (which the principal does not know). Similarly, $\Sigma_t^{\theta,\widehat{\mu}} = \sigma^2 \phi$ for all t, and

$$dm_t^{\widehat{\mu}} = \Sigma_t^{\theta,\widehat{\mu}} \frac{dY_t - \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right)dt}{\sigma^2} = \sigma \phi \, dB_t^{\widehat{\mu}}, \text{ with } m_0 = 0, \tag{4}$$

where $B_t^{\widehat{\mu}}$ is a standard Brownian motion under the measure induced by the actual effort policy $\widehat{\mu}$:

$$dB_t^{\widehat{\mu}} \equiv \frac{dY_t - \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right)dt}{\sigma}. \tag{5}$$

1.3 Formulating the optimal contracting problem

We first state the agent's problem. S_t denotes the balance of the agent's savings account, which earns interest at the constant rate $r > 0$. Given the contract $\{c_t, \mu_t\}$ the agent's problem is

$$\max_{\{\widehat{c}, \widehat{\mu}\}} \mathbb{E}_0^{\widehat{\mu}}\left[\int_0^\infty e^{-rt} u(\widehat{c}_t, \widehat{\mu}_t) dt\right] \tag{6}$$

$$\text{s.t. } dY_t = \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right) dt + \sigma d B_t^{\widehat{\mu}},$$

$$dS_t = r S_t dt + c_t dt - \widehat{c}_t dt \text{ with } S_0 = 0,$$

with the transversality condition, say the saving balance S_t has to be bounded.[10] Here, $\mathbb{E}^{\widehat{\mu}}[\cdot]$ denotes the expectation under the probability measure induced by the actual effort policy $\{\widehat{\mu}_t\}$, and $\{\widehat{c}_t\}$ is the actual consumption policy. Denote the optimal solution to Problem (6) by $\{c_t^\star, \mu_t^\star\}$.

We call the contract $\{c_t, \mu_t\}$ *incentive-compatible and no-savings* if, given the contract $\{c_t, \mu_t\}$, the solution to the agent's problem in Equation (6) is $c_t^\star = c_t$ and $\mu_t^\star = \mu_t$, which further implies $S_t = 0$ for any t (i.e., no private savings at any time). In other words, the agent finds it optimal to consume his wages and work as recommended. As a standard result in the literature, the following lemma shows that there is no loss of generality by restricting attention to incentive-compatible and no-savings contracts. The idea is similar to the revelation principle. Once the principal knows the agent's actual effort policy, she will perform correct Bayesian updating based on that policy; and since the principal can fully commit to the contract, she can save for the agent. Note, the optimal no-savings contract also can be implemented by some other compensation scheme in which the agent saves for himself.

Lemma 1. It is without loss of generality to focus on contracts that are incentive-compatible and no-savings.

Proof. The Appendix provides all proofs. ∎

The optimal contract solves the principal's problem:

$$\max_{\{c_t, \mu_t\} \text{ is incentive-compatible and no-savings}} \mathbb{E}_0^{\mu}\left[\int_0^\infty e^{-rt} (dY_t - c_t dt)\right], \tag{7}$$

so that $dY_t = \left(\mu_t + m_t^{\mu}\right) dt + \sigma d B_t^{\mu}$, and

$$\mathbb{E}_0^{\mu}\left[\int_0^\infty e^{-rt} u(c_t, \mu_t) dt\right] = v_0. \tag{8}$$

---

[10] In Appendix A.3, we explicitly impose the assumption of private savings being bounded in Assumption 1 in the proof of Proposition 1.

Equation (8) is the agent's participation constraint at $t=0$ for the agent with a reservation value v_0. Since negative transfers are allowed, this participation constraint at $t=0$ must bind.

2. The Agent's Problem

In this section we illustrate heuristically the necessary conditions for a contract $\{c_t, \mu_t\}$ to be incentive-compatible and no-savings.

2.1 Continuation value and incentives

Given the incentive-compatible and no-savings contract $\{c_t, \mu_t\}$, the agent's continuation value, which is his expected payoff from the continuation contract, is defined as:

$$v_t \equiv \mathbb{E}_t^\mu \left[\int_t^\infty e^{-r(s-t)} u(c_s, \mu_s) ds \right]. \tag{9}$$

According to the standard martingale representation argument (e.g., Sannikov, 2008), there exists some progressively measurable process $\{\beta_t\}$ so that

$$dv_t = rv_t dt - u(c_t, \mu_t) dt + \beta_t(-arv_t)\left(dY_t - \mu_t dt - m_t^\mu dt\right) \tag{10}$$

$$= rv_t dt - u(c_t, \mu_t) dt + \beta_t(-arv_t)\sigma dB_t^\mu.$$

We can interpret β_t as the *dollar* incentive on the agent's unexpected performance. From Sannikov (2008), we know that $\beta_t(-arv_t)$ can be interpreted as the incentive loading–measured in the agent's *utilities*–on his unexpected performance $dY_t - m_t^\mu dt$. We show shortly that $(-arv_t) > 0$ is the agent's marginal utility from consumption at time t, that is, $u_c(c_t, \mu_t)$. As a result, dividing utility incentives $\beta_t(-arv_t)$ by the marginal utility yields dollar incentives received by the agent. This is important for model tractability: As we show later in Section 2.4.1, using dollar incentives allows us to cancel $(-arv_t)$ and derive a simple expression for the agent's incentive compatibility condition that is *independent* of his continuation value v_t.

Later, we simply refer to pay–performance sensitivities $\{\beta_t\}$ as incentives. Throughout the paper, we impose a further technical condition for ease of our analysis. Essentially, we restrict the feasible incentive slopes $\{\beta_t\}$ to be bounded, that is, some sufficiently large constant M exists such that $\beta_t \in [-M, M]$. This assumption ensures that the endogenous state variable in the problem, the expected (properly) discounted future incentives, is bounded for any feasible contracts. Later, we will show that, given this restriction, the optimal incentives are independent of the exogenous bound M.[11]

---

[11] This boundedness assumption shares the same spirit as imposing a transversality condition. For instance, in the standard consumption-portfolio problem, to rule out Ponzi schemes, one often imposes the agent's wealth being bounded from below. In that context, the optimal portfolio strategy is also independent of the lower bound for the agent's wealth.

2.2 No savings

Following He (2011), we first show that the no-savings condition under CARA preferences implies that

$$rv_t = u(c_t, \mu_t) = -\frac{1}{a} \exp[-a(c_t - g(\mu_t))]. \tag{11}$$

We have the following lemma for any compensation contract, Π.

Lemma 2. At any time $t \geq 0$, consider a deviating agent who has some arbitrary savings S and faces the continuation contract Π_t. $v_t(S; \Pi)$ denotes the deviation continuation value. We have

$$v_t(S; \Pi) = v_t(0; \Pi) \cdot e^{-arS} = v_t \cdot e^{-arS}, \tag{12}$$

where we have used the fact that $v_t(0; \Pi)$ is the agent's continuation value v_t along the no-savings path defined in Equation (9).

The driving force behind this result is simple. Due to CARA preferences, the agent's problem is translation-invariant with respect to his underlying wealth level, as evident by $u(c_s + rS, \mu_s) = e^{-arS} u(c_s, \mu_s)$. Thus, for a CARA agent, given the extra savings S, his new optimal deviation policy is to take the optimal consumption-effort-learning policy without savings–which explains v_t in Equation (12), and to consume an extra rS more uniformly across all future dates/states–which explains the adjusting factor e^{-arS} in Equation (12).

The optimality of the agent's consumption-savings policy implies that his marginal utility from consumption must equal his marginal value of wealth. Equation (12) then implies that:

$$u_c(c_t, \mu_t) = \left. \frac{\partial v_t(S; \Pi)}{\partial S} \right|_{S=0} \overset{\text{due to (12)}}{=} -arv_t. \tag{13}$$

Equation (11) follows immediately from Equation (13) because under CARA preferences, the agent's utility level is linear in his marginal utility:

$$au(c_t, \mu_t) = -u_c(c_t, \mu_t). \tag{14}$$

Once we have established the key result in (11), we can plug it back into Equation (10), and find that v_t follows an (exponential) martingale:

$$dv_t = \beta_t(-arv_t)\sigma dB_t^\mu \Leftrightarrow v_s$$
$$= v_t \exp\left(-\int_t^s ar\beta_u \sigma dB_u^\mu - \frac{1}{2}\int_t^s a^2 r^2 \beta_u^2 \sigma^2 du\right) \text{ for } s > t.^{[12]} \tag{15}$$

Intuitively, a good performance $dB_u^\mu = \frac{1}{\sigma}(dY_t - \mu_t dt - m_t^\mu dt)$ for $u \in [t, s]$ increases v_s (recall $v_t < 0$ for CARA preferences), all else being equal. That v_s/v_t only depends on incentives $\{\beta_u\}_{s \leq u \leq t}$ is key to tractability for later analysis.

---

[12] Because $|\beta| < M$ is bounded, the local martingale $\{v_t\}$ is indeed a martingale. This result can also be understood by combining two observations: First, the agent can smooth out his consumption intertemporally, and hence his marginal utility has to follow a martingale. Second, his continuation value v_t is linear in his marginal utility u_c because of Equations (13) and (14).

2.3 Effort and belief distortion

The difficulty of introducing learning into the dynamic moral hazard problem is not learning per se. Rather, the challenge is to deal with the issue of belief manipulation: the agent, simply by shirking from the recommended effort today, can distort the principal's future beliefs about project profitability downward.

Consider the following thought experiment. Suppose that at time t the agent exerts an effort level $\widehat{\mu}_t$ below the recommended effort μ_t, and thus output is lower than what is expected by the principal on average. Crucially, however, the principal thinks the agent is exerting an effort of μ_t—thus she (through learning) mistakenly attributes lower output to a lower value of profitability θ_t. In contrast, the agent updates profitability θ_t based on his true effort level $\widehat{\mu}_t$, leading to a positive wedge $m_t^{\widehat{\mu}} - m_t^{\mu} = \mathbb{E}[\theta_t | \mathcal{Y}_t, \widehat{\mu}] - \mathbb{E}[\theta_t | \mathcal{Y}_t, \mu]$ between the beliefs of the agent and principal. In other words, by shirking, the agent makes the principal (mistakenly) underestimate profitability. This belief manipulation is beneficial to the agent in a dynamic setting—when future outputs turn out to be high, the agent gets rewarded for high profitability (based on the agent's correct information set) rather than his effort.

The above logic implies that any current effort deviation has a long-lasting effect in distorting the principal's belief, and we now formalize this effect. When the agent deviates from the recommended effort path $\{\mu\}$ by choosing effort policy $\{\widehat{\mu}\}$, the principal's belief about θ_s for $s > t$ is distorted downward. This distortion, represented by Δ_s, has the following intuitive expression:

$$\Delta_s \equiv m_s^{\widehat{\mu}} - m_s^{\mu} = \phi \int_0^s e^{-\phi(s-u)}(\mu_u - \widehat{\mu}_u)du.^{13} \tag{16}$$

Intuitively, the current belief distortion at time s equals the agent's cumulative effort deviations in the past $u \in [0, s]$, with a discount factor of ϕ. When $\phi = 0$, the zero prior uncertainty $\Sigma_0^{\theta} = \sigma^2 \phi = 0$ eliminates any belief divergence, and the issue of belief manipulation is absent.

Figure 1 heuristically illustrates the long-lasting belief distortion effect for a one-time effort deviation. The left panel shows that the solid line, which is the agent's actual effort $\{\widehat{\mu}\}$, lies below the dashed line, which is the recommended effort $\{\mu\}$, only at time interval $[t, t+dt]$, for some t. Let us say $\widehat{\mu}_t = \mu_t - \epsilon$, and for illustration we have assumed that $\{\mu\}$ takes a constant value. The right panel shows that this one-shot deviation triggers a long-lasting belief distortion with a decaying factor ϕ:

$$\Delta_s = m_s^{\widehat{\mu}} - m_s^{\mu} = \epsilon \cdot \phi e^{-\phi(s-t)}dt \quad \text{for } s > t. \tag{17}$$

---

[13] According to Equations (2) and (4), we have:

$$d\Delta_t = dm_t^{\widehat{\mu}} - dm_t^{\mu} = \phi\left(dY_t - \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right)dt\right) - \phi\left(dY_t - \left(\mu_t + m_t^{\mu}\right)dt\right) = \phi(\mu_t - \widehat{\mu}_t - \Delta_t)dt,$$

which leads to the expression of Δ_t in (13). Here, we have used $\Delta_0 = 0$, as both parties share the common prior when signing the contract.
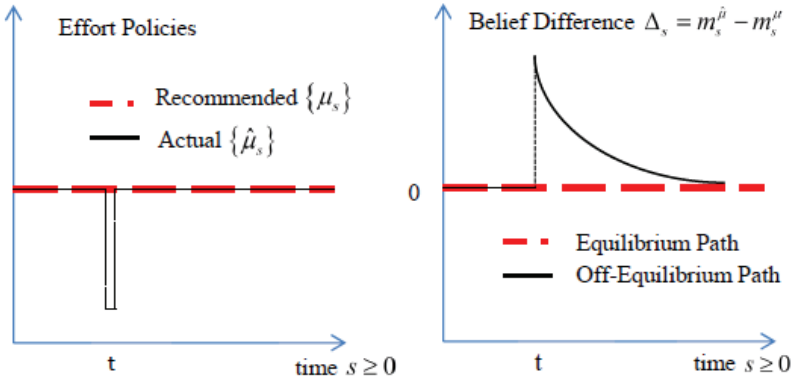
Figure 1
Long-lasting belief distortion (right) due to a one-time effort deviation (left). For illustration, we assume μ_t takes a constant value. The agent shirks at $[s, s+ds]$ so that $\widehat{\mu}_s = \mu_s - \epsilon$; this triggers a long-lasting belief distortion $\Delta_t = m_t^{\widehat{\mu}} - m_t^{\mu} = \epsilon \cdot \phi e^{-\phi(t-s)}$ so that the agent knows that the project is better than the principal thought (in off-equilibrium path).

Intuitively, as new information flows in, this belief divergence persists but decays over time exponentially at the rate of ϕ. As a result, even at time $s > t$, the principal mistakenly thinks the project is of a worse quality than the agent thinks.

As suggested by Equation (10), the contract relies on the agent's "unexpected" performance along the equilibrium path $dY_s - (\mu_s + m_s^{\mu}) ds$. This equals σdB_s^{μ} under the equilibrium measure and has a mean of zero. For the agent who deviates by exerting $\widehat{\mu} \neq \mu$, under his information set the above "unexpected" performance no longer has zero mean. Suppose that the agent has deviated before s so that $\widehat{\mu}_t \neq \mu_t$ where $t < s$. Even if the agent exerts the same effort at time s so that $\mu_s = \widehat{\mu}_s$, Equation (5) implies that

$$dY_s - \left(\widehat{\mu}_s + m_s^{\widehat{\mu}}\right) ds = dY_s - \left(\mu_s + m_s^{\widehat{\mu}}\right) ds \qquad (18)$$

has zero mean under the agent's information set. Hence, the "unexpected" performance $dY_s - (\mu_s + m_s^{\mu}) ds$ displays a positive drift under the agent's information set:

$$dY_s - \left(\mu_s + m_s^{\mu}\right) ds = \underbrace{\left[dY_s - \left(\mu_s + m_s^{\widehat{\mu}}\right) ds\right]}_{\text{zero mean under agent's info. set}} + \underbrace{\Delta_s ds}_{\text{belief divergence}} ,$$

Like in the previous example, a one-shot deviation in the past $\widehat{\mu}_t < \mu_t$ with $t < s$ implies that $\Delta_s > 0$. Intuitively, the principal would mistakenly think the project is worse than it actually is (under the agent's correct measure), and the agent can easily beat the principal's expectation and hence gain by $\Delta_s ds > 0$ for all future $s > t$.

2.4 Incentive compatibility constraint and intuition

Proposition 1 characterizes the agent's incentive compatibility constraint, along with the equilibrium consumption and continuation value heuristically derived above. We provide a rigorous proof for Proposition 1 in Appendix A.3. We also highlight that the agent's incentive compatibility constraint in Proposition 1 is essentially the agent's first-order condition in his effort decision, and we further show that the first-order condition is also sufficient for the agent's global optimality in Section 3.4, given certain conditions imposed on the derived optimal contract.

Proposition 1.   Agent's incentive compatibility constraint. For the contract $\{c_t, \mu_t\}$ to be incentive-compatible and no-savings, $\{\beta_t\}$ must satisfy

$$\mu_t = \underbrace{\beta_t}_{\text{instantaneous incentive}} - \underbrace{\mathbb{E}_t^\mu \left[\int_t^\infty \phi e^{-(\phi+r)(s-t)} \frac{\beta_s v_s}{v_t} ds \right]}_{\text{future information rent } p_t} = \beta_t - p_t \qquad (19)$$

where p_t denotes "*information rent*":

$$p_t \equiv \mathbb{E}_t^\mu \left[\int_t^\infty \phi e^{-(\phi+r)(s-t)} \beta_s \exp\left(-\int_t^s ar\beta_u \sigma dB_u^\mu - \frac{1}{2} \int_t^s a^2 r^2 \beta_u^2 \sigma^2 du \right) ds \right], \qquad (20)$$

as the exp term inside the bracket equals v_s/v_t, using Equation (12). In addition, Equation (11) implies that consumption (or wage) follows

$$c_t = g(\mu_t) - \frac{\ln(-arv_t)}{a}, \qquad (21)$$

and the continuation payoff from the contract is

$$v_t = v_0 \exp\left(-\int_0^t ar\beta_s \sigma dB_s^\mu - \frac{1}{2} \int_0^t a^2 r^2 \beta_s^2 \sigma^2 ds \right). \qquad (22)$$

In a standard dynamic agency problem without profitability uncertainty (e.g., $\phi = 0$), the agent's effort μ_t at time t should depend only on the time-t incentive β_t offered by the contract (i.e., $\mu_t = \beta_t$; recall the quadratic effort cost $g(\mu_t) = \mu_t^2/2$). With learning and associated belief-manipulation, the agent's effort decisions across periods are interlinked, as evident by the forward-looking nature of the second downward adjustment term in Equation (19).

The forward-looking downward adjustment term represents the *information rent* to the agent. Intuitively, this term captures the marginal benefit of manipulating the principal's future belief downward.[14] Also, the expression in (19) implies that the agent's continuation payoffs $\{v\}$ drop out, which allows us to write the agent's incentive compatibility constraint independent of $\{v\}$. This convenient property is crucial for the tractability of our problem.

---

[14] This information rent term captures the *marginal* rent that the agent may enjoy by deviating from the recommended effort slightly, rather than the rent that the agent actually enjoys in equilibrium; in equilibrium the principal knows the agent's actual effort exactly. Nevertheless, like in any typical moral hazard model, the marginal deviation benefit (marginal rent) is important in characterizing the agent's incentive-compatibility condition.

2.4.1 Intuition for the incentive compatibility constraint. The rest of this subsection is devoted to understanding the key incentive compatibility constraint (19). Consider again the example in Section 2.3 in which the agent reduces his effort to slightly below the recommended effort level μ_t, say $\mu_t - \epsilon$, only at the time interval $[t, t+dt]$. In other words, given the recommended policy $\{\mu\}$, the deviation effort policy is

$$\mu^\epsilon \equiv \begin{cases} \mu_s & \text{for } s \notin [t, t+dt]; \\ \mu_s - \epsilon & \text{otherwise.} \end{cases} \tag{23}$$

What is the impact of this deviation effort policy on the agent's total payoff from time t onwards, including his instantaneous utility?

In Appendix A.4, we show that, under the new effort policy μ^ϵ, the agent's continuation payoff together his instantaneous flow payoff at t can be written as

$$u(c_t, \mu_t - \epsilon)dt + v_t + \mathbb{E}_t^{\mu^\epsilon} \left[\int_t^\infty e^{-r(s-t)} dv_s \right], \tag{24}$$

where $\mathbb{E}_t^{\mu^\epsilon}$ emphasizes that the agent forms his expectation based on his information set induced by μ^ϵ. Using the result in Equation (12), we can rewrite (24) heuristically as:

$$u(c_t, \mu_t - \epsilon)dt + v_t + \mathbb{E}_t^{\mu^\epsilon} \left\{ \begin{array}{l} \beta_t(-arv_t)\left(dY_t(\mu_t - \epsilon) - \mu_t dt - m_t^\mu dt\right) + \\ \int_{t+dt}^\infty e^{-r(s-t)} \beta_s(-arv_s)\left[dY_s - \left(\mu_s + m_s^\mu\right)ds\right] \end{array} \right\}$$

$$= \underbrace{u(c_t, \mu_t - \epsilon)dt}_{\text{saving effort cost instantaneously}} + v_t +$$

$$\mathbb{E}_t^{\mu^\epsilon} \left\{ \int_{t+dt}^\infty e^{-r(s-t)} \beta_s(-arv_s) \underbrace{\left[\underbrace{\left(dY_s - \left(\mu_s^\epsilon + m_s^{\mu^\epsilon}\right)dt\right)}_{\text{martingale under info set generated by } \mu} + \underbrace{\Delta_s ds}_{\text{belief divergences}} \right]}_{\text{creating belief divergence persistently}} \right\}.$$

$$\underbrace{\beta_t(-arv_t)\left(dY_t(\mu_t - \epsilon) - \mu_t dt - m_t^\mu dt\right) +}_{\text{hurting performance instantaneously}}$$

$$\tag{25}$$

There should be another correction term in $(\mu_s^\epsilon - \mu_s)ds$ in the second equality, but it is zero because of (23), that is, we consider a one-shot deviation at time t from the equilibrium effort policy.

There are two channels through which shirking at time t affects the agent's continuation value. The first channel captures the instantaneous performance effect, that is, the agent's effort affects instantaneous performance dY_t and, thus, his continuation value. To see this, write performance $dY_t(\mu_t)$ over $[t, t+dt]$ as a function of time-t effort μ_t. Exerting effort $\mu_t - \epsilon$ hurts the short-term performance over $[t, t+dt]$ because

$$dY_t(\mu_t - \epsilon) = (\mu_t - \epsilon)dt + m_t^\mu dt + \sigma dB_t^\mu = dY_t(\mu_t) - \epsilon dt.$$

Modulated by incentives, this leads to a drop in the agent's continuation value by $\beta_t(-arv_t) \cdot \epsilon dt$, via the channel of "hurting performance instantaneously."

The second channel is the persistent effect due to belief manipulation. As discussed in Section 2.3, the agent's shirking at time t shifts the belief divergence path $\{\Delta_s\}$ away from the equilibrium path $\{\Delta_s = 0\}$ for $s > t$, according to Equation (13).

We show that the incentive compatibility constraint in Equation (19) is implied by Equation (25). By "reducing effort cost instantaneously" in Equation (25), the agent's marginal gain from shirking at t is $-u_\mu(c_t, \mu_t) \cdot \epsilon dt$. Since $u_\mu(c_t, \mu_t) = -u_c(c_t, \mu_t)\mu_t = arv_t\mu_t$, this marginal gain is $(-arv_t)\mu_t \cdot \epsilon dt$. On the other hand, shirking "hurts performance instantaneously" in Equation (25), which gives rise to a marginal cost of $\beta_t(-arv_t) \cdot \epsilon dt$. In standard models without belief manipulation, these two forces fully determine the agent's trade-off in choosing his optimal effort at time t.

Next we analyze the novel term "creating belief divergence persistently" in (25). There, because $dY_s - \left(\mu_s^\epsilon + m_s^{\mu^\epsilon}\right)dt$ has zero mean, this term equals

$$\mathbb{E}_t^{\mu^\epsilon}\left[\int_t^\infty e^{-r(s-t)}\beta_s(-arv_s)\Delta_s ds\right]. \tag{26}$$

Recall that Equation (17) says that the belief divergence in any future time $s > t$ is $\Delta_s = \phi e^{-\phi(s-t)}\epsilon dt$. Plugging in to (26), the marginal impact of shirking via the channel of belief manipulation is

$$\mathbb{E}_t^\mu\left[\int_t^\infty \phi e^{-(\phi+r)(s-t)} \underbrace{\beta_s}_{\text{future incentives}} \underbrace{(-arv_s)}_{\text{marginal utility}} ds\right] \cdot \epsilon dt + o(\epsilon dt).^{15}$$

Intuitively, if the principal mistakenly believes that the project is less profitable than it should be, the agent's normal performance will be considered superb. The higher-powered the future incentives $\{\beta_s\}$, the greater the information rent. And, for a risk-averse agent, the information rent depends on the agent's future marginal utility $(-arv_s)$ when receiving these manipulation benefits.

Combining three pieces together (canceling ϵdt and ignoring higher-order terms), and dividing both sides by time-t marginal utility $(-arv_t)$, we arrive at the agent's incentive compatibility constraint as Equation (19).

3. Principal's Problem and Optimal Contracting

From now on we focus on incentive-compatible contracts such that both parties will have the same information set along the equilibrium path. As a result, we write dB_t^μ and $\mathbb{E}^\mu[\cdot]$ as dB_t and $\mathbb{E}[\cdot]$, respectively, for ease of notation.

3.1 Rewriting the principal's problem

In light of Proposition 1, we first rewrite the principal's problem in Equation (7). Proposition 1 establishes an important link between recommended effort $\{\mu_t\}$ and incentives $\{\beta_t\}$ in any incentive-compatible contracts. Moreover,

the principal can choose the optimal $\{\beta_t^*\}$ to maximize her value, and the corresponding optimal consumption process $\{c_t^*\}$ and the optimal effort policy $\{\mu_t^*\}$ are determined by Equations (21) and (19), respectively. Therefore, the principal is as if choosing incentives $\{\beta_t\}$ only:

$$\max_{\{\beta_t\}} \mathbb{E}\left[\int_0^\infty e^{-rt}(dY_t - c_t dt)\right], \tag{27}$$

$$s.t. \qquad dY_t = (\mu_t + m_t)dt + \sigma dB_t \text{ and } dm_t = \phi\sigma dB_t, \tag{28}$$

$$c_t = g(\mu_t) - \frac{\ln(-arv_t)}{a}, \text{ where } g(\mu_t) = \frac{1}{2}\mu_t^2, \tag{29}$$

$$dv_t = \beta_t(-arv_t)\sigma dB_t, \text{ given } v_0, \tag{30}$$

$$\mu_t = \beta_t - p_t. \tag{31}$$

Here, Equation (28) describes the dynamics of output and posterior belief; Equations (29)–(31) are derived from Equations (19)–(22) in Proposition 1; and p_t in Equation (31) is given by Equation (20).

Thanks to the CARA preference, the agent's continuation value v_t separates from the problem and the optimal contracting problem can be rewritten without v_t. Start from the principal's objective in Equation (27). In Appendix A.5, we show that

$$\mathbb{E}\left[\int_0^\infty e^{-rt}(dY_t - c_t dt)\right]$$

$$= \mathbb{E}\left[\underbrace{\int_0^\infty e^{-rt}\mu_t dt}_{\text{expected output}}\right] - \left(\underbrace{-\frac{\ln(-arv_0)}{ar}}_{\text{C.E. of outside option } v_0}\right)$$

$$- \mathbb{E}\left[\int_0^\infty e^{-rt}\left(\underbrace{g(\mu_t)}_{\text{effort cost}} + \underbrace{ar\sigma^2\beta_t^2/2}_{\text{risk comp.}}\right)dt\right], \tag{32}$$

The discounted expected output is driven by the agent's effort (recall that we normalize the project's initial profitability as $m_0 = 0$). The total compensation cost is the certainty equivalent (i.e., $-\ln(-arv_0)/(ar)$) of delivering the agent's outside option v_0, plus the monetary effort cost (i.e., $g(\mu_t) = \mu_t^2/2$), and the discounted risk compensation due to incentive provisions. Thus, the certainty equivalent separates from the problem, and the optimal solution $\{\beta_t^*\}$ will be independent of the agent's initial outside option v_0. This result comes from the lack of wealth effect under CARA preferences.

Combining Equations (31) and (32), the principal's problem is simplified to

$$\max_{\{\beta_t\}} \mathbb{E}\left[\int_0^\infty e^{-rt}\left(\mu_t - \frac{1}{2}\mu_t^2 - \frac{1}{2}ar\sigma^2\beta_t^2\right)dt\right] \tag{33}$$

$$s.t. \ \mu_t = \beta_t - p_t \text{ with } p_t \text{ as given in (20)}.$$

Importantly, only incentives $\{\beta\}$, but not continuation payoffs $\{v\}$, enter the problem (p_t depends on $\{\beta\}$ only).

3.2 Recursive formulation

We now recursively formulate the principal's problem in (33) and solve it by dynamic programming. Let the continuation value in problem (33) be

$$V(p) \equiv \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} \left(\mu_s - \frac{1}{2}\mu_s^2 - \frac{1}{2}ar\sigma^2\beta_s^2 \right) ds \right]. \qquad (34)$$

Consequently, the information rent p_t serves as the only state variable for the principal when designing the optimal contract. The information rent captures the marginal benefit of the agent's shirking due to the belief manipulation effect. Recall the definition of the information rent in Equation (20), which, together with the martingale representation theorem, implies that there exists some progressively measurable process $\left\{\sigma_t^p\right\}$ so that the dynamics of p_t follows (see Appendix A.6):

$$dp_t = \left[(\phi+r)\,p_t + \beta_t \left(ar\sigma\sigma_t^p - \phi \right) \right] dt + \sigma_t^p dB_t. \qquad (35)$$

From now on, we interpret $\left\{\sigma_t^p, \beta_t\right\}$ as our control because the pair determines the drift and diffusion of p_t in (35). As we will derive σ_t^p and β_t as a function of the state variable p_t, the control pair $\left\{\sigma_t^p, \beta_t\right\}$ gives the full history of $\{\beta_t : t \geq 0\}$ that we are after.

Remark 1. Strictly speaking, the value function in Equation (34) is only a part of the principal's full value function. Following the same steps in Equation (32), one can write the principal's full value function $J(m_t, v_t, p_t)$, which depends on project posterior mean m_t, the agent's continuation value v_t, and the agent's information rent p_t, as

$$J(m_t, v_t, p_t) \equiv \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)}(dY_s - c_s ds) \right]$$

$$= \underbrace{\frac{m_t}{r}}_{\text{expected proj. value}} + \underbrace{\frac{\ln(-arv_t)}{ar}}_{\text{C.E. of agent's } v_t} + \underbrace{V(p_t)}_{\text{value function}}. \qquad (36)$$

The additive structure in Equation (1) gives rise to the first term, which captures the expected project value m_t/r without effort; and the CARA preference allows us to separate the agent's certainty equivalent given his continuation value v_t (the second term) from the problem.[16] Maximizing $J(m_t, v_t, p_t)$ is equivalent to maximizing $V(p_t)$. As a result, we refer to $V(p_t)$ simply as the principal's value function wherever no confusion arises.

------

[16] The certainty equivalent is the amount of money that an individual would view as equally desirable as a stream of risky cash flows. Consuming $\frac{\ln(-arv)}{a}$ per period forever delivers a value of v for the CARA agent in our model. For why this separation works, see Section 3.5.

The optimal contract can now be fully characterized by an ordinary differential equation (ODE), which is the Hamilton-Jacobi-Bellman (HJB) equation for the problem (33):

$$rV(p) = \max_{\beta, \sigma^p} (\beta - p) - \frac{1}{2}(\beta - p)^2 - \frac{ar\sigma^2}{2}\beta^2 + V_p[(\phi + r)p$$

$$+ \beta(ar\sigma\sigma^p - \phi)] + \frac{1}{2}V_{pp}(\sigma^p)^2. \tag{37}$$

We will verify in Proposition 2 that $1 + ar\sigma^2 + a^2r^2\sigma^2 (V_p)^2/V_{pp} > 0$ and $V_{pp} < 0$. Then, the first-order optimality conditions for the optimal control $\{\sigma_t^{p*}, \beta_t^*\}$ are given by

$$\beta^* = \frac{1 + p - \phi V_p}{1 + ar\sigma^2 + a^2r^2\sigma^2 \frac{(V_p)^2}{V_{pp}}} \text{ and } \sigma^{p*} = -ar\sigma\beta^* \frac{V_p}{V_{pp}}. \tag{38}$$

Plugging them back into the HJB Equation (37), we have

$$rV(p) = \frac{1}{2} \frac{\left(1 + p - \phi V_p(p)\right)^2}{1 + ar\sigma^2 + a^2r^2\sigma^2 \frac{[V_p(p)]^2}{V_{pp}(p)}} - p - \frac{1}{2}p^2 + (\phi + r)p \cdot V_p(p). \tag{39}$$

We solve the problem in (37) by analyzing the above ODE in (39).

3.3 Optimal contracting
Before we start analyzing the optimal contract, we first consider a (trivial) benchmark case. Suppose that the profitability θ_t is observable. This is essentially the classic Holmstrom and Milgrom (1987) model, except that the optimal contract always benchmarks the agent's performance to θ_t. Using the incentive constraint $\mu_t = \beta_t$, the optimal solution is

$$\mu_t^{HM} = \beta_t^{HM} = \frac{1}{1 + ar\sigma^2}, \tag{40}$$

and the principal's value is $V^{HM} = 1/(2r(1 + ar\sigma^2))$. The optimal contract can be implemented by a constant equity share $1/(1 + ar\sigma^2)$ (with proper benchmarking). What is more, the value V^{HM} serves as an upper bound for our value function $V(p)$ when profitability is unobservable:

$$V(p) \leq V^{HM} = \frac{1}{2r(1 + ar\sigma^2)}. \tag{41}$$

This is because V^{HM} will be the principal's value in our model but after seeing the additional (precise) information about θ_t (and she can dispose this information freely).

To solve for the optimal contract, we analyze the ODE (39) with the boundary condition in Equation (A13) using the technique of dynamic

programming. The following proposition is our main result, which characterizes the properties of the value function, and hence the optimal policy $\{\beta^*, \sigma^{P,*}\}$ like in Equation (38). We impose the following parametric condition throughout the paper, which restricts ϕ to be relatively small:

$$\frac{(r+\phi)^3}{\phi} < \frac{r^2}{2a}\left(1+ar\sigma^2\right). \tag{42}$$

The sole purpose of this condition is to ensure the concavity of $V(p)$ in the proof method we employ.[17]

Proposition 2. Property of value function of optimal contracting. Suppose that Equation (42) holds. We have the following properties for $V(p) \in \mathbb{C}^2$ which characterize the optimal contract.

1. $V(0)=0$ and $V_p(0)=1/\phi$.
2. $V(p)$ is strictly concave over a compact interval, and $1+ar\sigma^2+ a^2r^2\sigma^2\frac{(V_p)^2}{V_{pp}} > 0$.
3. There exists a unique $\overline{p} \in (0, \overline{p}^d)$ such that $V_p(\overline{p})=0$, where the constant

$$\overline{p}^d \equiv \frac{2\phi}{(2\phi+r)ar\sigma^2+r+\sqrt{(2\phi+r)^2 a^2r^2\sigma^4+2ar\sigma^2\left[(\phi+r)^2+\phi^2\right]+r^2}} > 0. \tag{43}$$

    Under the optimal policy, $\overline{p}$ is an upper entrance-no-exit boundary, and 0 is a lower absorbing boundary 0 with $V(0)=0$. This implies that under the optimal policy the endogenous state variable p_t^* never exits the interval $[0, \overline{p}]$.

In the optimal contract, the principal sets the initial information rent p_0^* to be $\overline{p}$. Afterwards, the state variable p_t^* evolves according to Equation (35), and the optimal control is characterized by Equation (38). Interestingly, property 3 in Proposition 2 states that the information rent p_t^* will never wander out of an endogenous interval $[0, \overline{p}]$, which suggests that it is suboptimal to promise too much future incentives (recall information rent p_t^* is the discounted promised future incentives). This result is related to Holmstrom and Milgrom (1987) in which the optimal incentives $\left\{\beta_t^{HM} = \frac{1}{1+ar\sigma^2}\right\}$ remain constant over time. In our model with learning, the optimal incentives $\{\beta_t^*\}$ become stochastic, but the information rent $\{p_t^*\}$ and hence incentives $\{\beta_t^*\}$ remain endogenously bounded due to stationary model primitives (CARA-normal setting, additive technology in Equation (28), stationary learning, and the effort cost becomes prohibitive for unbounded μ).

---

[17] We require this condition for our *particular proof* for the concavity of the value function $V(p)$. When condition (42) fails, other proof methods might exist to show the concavity of $V(p)$.

Remark 2. Though we are able to theoretically analyze the ODE (39) in Proposition 2, numerically solving (39) and then investigating the properties of optimal contracting are far from easy tasks. This is because the ODE in (39) has two singular points at both $p=0$ and $p=\overline{p}$: the coefficient in front of the second-order derivative becomes zero at both end points (i.e., $\sigma^p(0)=\sigma^p(\overline{p})=0$), and as a result the ODE collapses to one of the first order. In addition, the singular point $\overline{p}$ is a free boundary itself that we need to pin down. We are able to develop a numerical algorithm to solve the ODE based on the approach of numerical integration with a desirable degree of accuracy and numerical stability. The Internet Appendix provides the details about the algorithm, as well as the Matlab programs.[18]

3.4 Validity of the first-order approach

In deriving the optimal contract in Proposition 2, we rely on the agent's incentive compatibility constraint (19), which is the agent's first-order condition in his effort decision. This is the so-called "first-order approach", and in the dynamic agency literature it is challenging to show that the necessary local first-order condition for the agent's problem is indeed sufficient for the agent's global optimality.

We have shown that in the optimal contract, the optimal policy $\{\beta_t^*, \sigma_t^{p*}\}$ are bounded. In this section, we show that we are able to guarantee the validity of the first-order approach, after imposing certain sufficient conditions on the volatility of information rent p_t, that is, σ_t^{p*}, in the optimal contract. More specifically, we show that the first-order conditions in Proposition 1 are sufficient to ensure the agent's global optimality by following an upper-bound approach employed in Sannikov (2014).

Proposition 3. Validity of the first-order approach. Suppose that in the optimal contract $\left|\sigma_t^{p,*}\right|$ is not too large, so that either (A22) or (A23) in the proof in Appendix A.8 holds. Then under the usual transversality condition, given the optimal contract the policy in Proposition 1 solves the agent's problem in (6).

To illustrate the basic idea, suppose that the agent facing the employment contract has deviated in the past, by having saved a bit and/or shirked a bit. For private savings, the agent's deviation state is his saving balance $S_t = \int_0^t e^{r(t-s)}(c_s - \widehat{c}_s)ds$; while for shirking that distorts the principal's current and future beliefs, the relevant deviation state is the belief distortion $\Delta_t = \phi \int_0^t e^{\phi(s-t)}(\mu_s - \widehat{\mu}_s)ds$. Given these two deviation states, we define a function

---

[18] As an alternative approach, we have also conducted an asymptotical analysis that is tractable but may lead to inaccurate approximation results when the agent is not sufficiently risk tolerant. The Internet Appendix provides the details about the asymptotical analysis.

$W(v_t, p_t; S_t, \Delta_t)$ which is constructed to be the upper bound of the agent's deviation value given the optimal contract and these two deviation states:

$$
W\left(\underbrace{v_t, p_t}_{\text{eqbm contract}}; \underbrace{S_t, \Delta_t}_{\text{dev. states}}\right) \equiv \underbrace{v_t}_{\text{eqbm cont. payoff}} \cdot \underbrace{\exp(-ar\,S_t)}_{\text{dev. value from savings}}
$$

$$
\cdot \underbrace{\exp\left(-ar\left(\frac{1}{\phi}\Delta_t\,p_t + 0.5k\Delta_t^2\right)\right)}_{\text{dev. value from belief distortions}}. \tag{44}
$$

In (44), two deviation states–private savings S_t and belief manipulation Δ_t– enter the proposed upper bound of the agent's deviation value in a multiplicative way, capturing the potential interdependence between the agent's deviating incentives of consumption and effort.

The functional form of $W(v_t, p_t; S_t, \Delta_t)$ is intuitive. When the agent never deviates, that is, $S_t = \Delta_t = 0$, then $W(v_t, p_t; S_t, \Delta_t) = v_t$ is the agent's equilibrium continuation payoff achieved by the equilibrium strategy satisfying the first-order conditions. The second term $\exp(-ar\,S_t)$ in (44) is the extra value that the agent gains by having a private saving of S and hence always consuming rS extra in all future states. The third term is about the gain from belief distortion due to past effort deviations. We know that the first-order gain from belief manipulation is the information rent p_t, which explains the linear coefficient p_t/ϕ in front of the belief distortion Δ_t inside the parentheses. The quadratic coefficient k is an appropriately chosen constant (see the proof of Proposition 3 in Appendix A.8) to ensure $W(v_t, p_t; S_t, \Delta_t)$ being the *upper bound* of the agent's deviation value, given his current deviation state-pair (S_t, Δ_t).[19] Because this upper bound satisfies the property of $W(v_0, p_0 = \overline{p}; S_0 = 0, \Delta_0 = 0) = v_0$, the strategy satisfying first-order conditions achieves this upper bound, and hence is indeed optimal for the agent who is endowed with zero savings and zero belief distortion.

The proof of Proposition 3 goes through if the volatility of information rent σ_t^{p*} in the optimal contract is not excessively high. For instance, in the proof in Appendix A.8, one sufficient condition (A23) requires that $(\sigma_t^p)^2 \le \sigma^2 \phi^2 (r + 2\phi - \phi^2)$, and (A22) is a bit weaker; both conditions are easily satisfied in our numerical examples. A similar condition for the volatility of the endogenous state is required in Sannikov (2014). Intuitively, all else equal, the agent's global deviation value tends to be increasing in the volatility σ_t^{p*} of the state, because the agent has the "option" to adjust his optimal strategy swiftly following a sequence of deviations and performance shocks.

---

[19] It is worth noting that $W(v_t, p_t; S_t, \Delta_t)$ is not *exactly* the deviation value of the agent; it just provides an upper bound for the agent's deviation value. This result is established by showing that the auxiliary gain process $\int_0^t e^{-rs} u(\widehat{c}_s, \widehat{\mu}_s) ds + e^{-rt} W(v_t, p_t; S_t, \Delta_t)$ follows a supermartingale for any feasible policy $\{\widehat{c}_t, \widehat{\mu}_t\}$. For more details, see the proof for Proposition 3 in Appendix A.8.

3.5 Discussion of assumptions

We make two simplifying assumptions in this paper: one is the assumption of CARA utility function and the other is private savings. We now discuss the roles played by these two assumptions in making the model tractable. In short, CARA preference without wealth effect is the key to reducing the dimensionality into a unidimensional problem; while private savings, under CARA preferences, helps simplify the solution greatly. As will be discussed in the Concluding Remarks, both assumptions are not responsible for our key qualitative results.

3.5.1 First-order approach and state variables. We first briefly outline the general first-order approach that is widely used in the literature in solving this class of problems. For any general utility function $u(c, \mu)$, following the same steps in Section 2.4.1, we can derive the first-order incentive-compatibility condition for the (interior) optimal effort policy as

$$-u_\mu(c_t, \mu_t) = \widetilde{\beta}_t - \widetilde{p}_t, \tag{45}$$

where $\widetilde{\beta}_t$ the diffusion term, expressed in utilities, in the process of continuation value v_t (see Equation (10)):

$$\widetilde{\beta}_t \equiv (-arv_t) \cdot \beta_t; \tag{46}$$

and $\widetilde{p}_t$ the information rent that captures the additional value of shirking due to belief manipulation:

$$\widetilde{p}_t \equiv \mathbb{E}_t\left[\int_t^\infty \phi e^{-(r+\phi)(s-t)} \widetilde{\beta}_s ds\right]. \tag{47}$$

In the case of private savings, there is an additional incentive constraint for the agent's optimal consumption policy, as the agent can privately save:

$$u_c(c_t, \mu_t) = \widetilde{q}_t, \tag{48}$$

where $\widetilde{q}_t$ is a new state variable capturing the marginal value of private savings (or consumption).

There are two major obstacles in solving the general problem using the first-order approach outlined above. The first issue is dimensionality: In general, besides m_t which captures the project's quality, the principal's value function depends on the state variables v_t and $\widetilde{p}_t$, and also $\widetilde{q}_t$ if private savings are further allowed. Given the additive cash-flow technology in (1), m_t enters the principal's value additively with m_t/r, and we will focus on the function $\widetilde{J}(v_t, \widetilde{p}_t, \widetilde{q}_t)$ from now on.

Oftentimes, the solution $\widetilde{J}$ can only be obtained by numerical methods. This leads to the second—and more important—concern: This first-order approach might not be valid. In other words, numerical solutions typically make it harder to rigorously verify that, facing the proposed optimal contract, the agent cannot have strictly profitable (global) deviations. In contrast, in our model, the combination of the CARA preference and private savings allows us to give a full characterization of the solution and to further verify the validity of the first-order approach in Section 3.4.

3.5.2 CARA preferences. For CARA preferences, the state variable, v_t (i.e., the agent's continuation payoff) always separates out from the problem, regardless whether we allow for private savings or not. More specifically, when the agent has CARA preferences, the principal's value function is in the form of

$$\widetilde{J}(v_t, \widetilde{p}_t, \widetilde{q}_t) = \widetilde{J}(-1, \widetilde{p}_t, \widetilde{q}_t) + \frac{\ln(-v_t)}{ar}.^{20} \tag{49}$$

The intuition is as follows. Since $u(c_t + \delta_c, \mu_t) = e^{-a\delta_c} u(c_t, \mu_t)$, shifting the CARA agent's consumption by a constant δ_c in all states multiplies the agent's utility by the *same* factor $e^{-a\delta_c}$ under both the recommended strategy and all deviations. As a result, shifting consumption by $\delta_c = -\frac{\ln(-v)}{a}$, which shifts the agent's continuation payoff multiplicatively by a factor of $-v > 0$, does not change the incentive compatibility of the contract. Applying this argument to the optimal contract, (20) simply says that the principal is as if facing an agent with a normalized continuation value of -1, but then shifting the agent's consumption all the states by $-\frac{\ln(-v)}{a}$ at the cost of $\frac{-\ln(-v)}{ar}$ in present value. This argument holds regardless whether the no-saving constraint is present or not.

For general utility functions, we typically need to solve a partial differential equation (PDE) with v being one of the state variables. For simplicity, suppose that the agent cannot privately save, so that the principal's value function can be written as $\widetilde{J}(v_t, \widetilde{p}_t)$. Standard argument implies that $\widetilde{J}(\cdot, \cdot)$ satisfies the following PDE:

$$r\widetilde{J}(v, \widetilde{p}) = \max_{c, \widetilde{\beta}, \widetilde{\gamma}} \mu\left(c; \widetilde{\beta}, \widetilde{p}\right) - c + \widetilde{J}_v\left(rv - u\left(c, \mu\left(c, \widetilde{\beta}; \widetilde{p}\right)\right)\right)$$

$$+ \widetilde{J}_{\widetilde{p}}\left((r+\phi)\widetilde{p} - \phi\widetilde{\beta}\right) + \frac{\sigma^2}{2}\left[\widetilde{J}_{vv}\widetilde{\beta}^2 + \widetilde{J}_{\widetilde{p}\widetilde{p}}\widetilde{\gamma}^2 + 2\widetilde{J}_{v\widetilde{p}}\widetilde{\beta}\widetilde{\gamma}\right]. \tag{50}$$

Given the optimal consumption c^*, $\mu\left(c^*; \widetilde{\beta}, \widetilde{p}\right)$ denotes the agent's optimal effort satisfying the first-order condition $-u_\mu\left(c^*, \mu\left(c^*; \widetilde{\beta}, \widetilde{p}\right)\right) = \widetilde{\beta} - \widetilde{p}$ in Equation (45), and $\widetilde{\gamma}$ is the diffusion term associated with $\widetilde{p}$. Solving (50) is a daunting task in general.[21]

3.5.3 What if the agent cannot privately save?. We have also assumed that the agent can privately save to smooth his consumption over his life time. Although for general utility functions allowing for private savings demands another state variable q_t, for CARA preferences it does not. To see this, Lemma 2 and Equation (13) imply that $\widetilde{q}_t = -arv_t$ always, rendering $\widetilde{q}_t$ to be redundant

---

[20] The certainty equivalent term $\frac{\ln(-arv_t)}{ar}$ in (36) differs from $\frac{\ln(-v_t)}{ar}$ by a constant $\frac{\ln(ar)}{ar}$ which is absorbed in $\widetilde{J}(-1, \widetilde{p}_t, \widetilde{q}_t)$.

[21] For papers studying dynamic contracting problems with private savings in which the agent has non-CARA preferences, see Kocherlakota (2004), He (2012), and, more recently, Di Tella and Sannikov (2016).

given v_t. The intuition is simple: Private savings imply that the marginal value of saving equals the marginal value of consumption, which is proportional to the level of utility under the CARA utility.

As a result, under CARA preferences, we only need to keep track of the agent's information rent as the single state variable, whether or not private savings are allowed. This paper fully solves the case with private savings, and Appendix A.9 outlines the derivations for the case without private savings. There, we show that similar to the private saving case, the key state variable for the optimal contract is again (recall (46) and (47)) $p_t = \widetilde{p}_t/(-arv_t)$, and we derive the ODE for the principal's value function $V^{ns}(p)$.

We emphasize that allowing for private savings under CARA preferences greatly simplifies our problem, which facilitates our rigorous characterization of the optimal contract in Proposition 2 and hence the verification of the first-order approach in Proposition 3. We illustrate this point by comparing the agent's key incentive compatibility condition for the setting with private savings to that without. Under both settings, the agent's incentive compatibility condition is $-u_\mu(c_t,\mu_t) = \widetilde{\beta}_t - \widetilde{p}_t$ in (45), which, after rewriting in terms of β and p, is

$$arv_t(\beta_t - p_t) = u_\mu(c_t,\mu_t) \overset{\text{CARA}}{=} a\mu_t \cdot u_t \Rightarrow \mu_t = (\beta_t - p_t) \cdot \frac{rv_t}{u_t}.$$

Here, we use the property of CARA utility in the second equation. With private savings, the agent's consumption smoothing implies $u_t = rv_t$ as in (11), rendering the simple and intuitive Incentive-Compatibility condition

$$\mu_t = \beta_t - p_t. \tag{51}$$

In contrast, when the agent cannot smooth his consumption, the principal optimally chosen the agent's current consumption c_t to control the ratio between the agent's instantaneous utility u_t and his continuation payoff v_t, and as shown in Appendix A.9, the resultant optimal effort in the optimal contract satisfies

$$\mu_t \left(1 + a\mu_t^2 - a\mu_t\right) = \left(1 - arp_t \cdot V_p^{ns}(p_t)\right)(\beta_t - p_t).$$

Comparing to (51), this is a cubic equation in μ_t, with right-hand-side involving the first-order derivative of the value function $V_p^{ns}(p)$.[22] What is more, the final ODE (A30) for $V^{ns}(p)$ derived in Appendix A.9 seems dauntingly complicated for rigorous analytical analysis on its key properties like in Proposition 2, and we await future research to make progress on this front.

4. Model Implications

To better understand our results, we first analyze the case in which we restrict the incentives $\{\beta_t\}$ to be deterministic. We then turn to the general case in

---

[22] When the principal chooses the agent's effort μ_t (while fixing the agent's consumption c_t), this choice affects the agent's instantaneous utility $u_t(c_t,\mu_t)$ and hence the drift of v_t, that is, $u_t - rv_t$ like in (10). Because the drift of v_t enters the drift of $p_t = \widetilde{p}_t/(-arv_t)$, the principal takes into account the first-order impact $V_p^{ns}(p_t)$ in choosing μ_t. In contrast, when the agent can control his own consumption, $u_t - rv_t = 0$ always holds thanks to consumption smoothing by the agent.

which the optimal policies are stochastic and compare it to both the Holmstrom and Milgrom (1987) benchmark and the contract with optimal deterministic incentives. The discussion focuses on two qualitative features of optimal contracting: front-loaded incentives and option-like incentives.

4.1 Contract with deterministic incentives

We will show that the optimal incentives are front-loaded (or, time decreasing) in dynamic contracting with learning. This result is best illustrated when we constrain the incentives $\{\beta\}$ to be deterministic (but can vary over time), a case in which we can analytically derive the time-decreasing incentives. This case also provides an important benchmark for the fully stochastic optimal contract, because deterministic contracts do rule out the option-like feature (i.e., raising incentives following good performance).

The reason that $\{\beta\}$ being deterministic helps is that we can move the conditional expectation in Equation (20) inside the integral,[23] so that the information rent $p_t = \phi \int_t^\infty e^{-(\phi+r)(s-t)} \beta_s \, ds$ is a deterministic process with $\sigma^p = 0$. $V^d(p)$ denotes the value function with deterministic policies, where the superscript "d" stands for "deterministic." Plugging $\sigma^p = 0$ into (37), we have $\beta^d(p) = (1 + p - \phi V_p^d)/(1 + ar\sigma^2)$, with the resultant HJB equation as:

$$r V^d(p) = \frac{1}{2} \frac{(1 + p - \phi V_p^d(p))^2}{1 + ar\sigma^2} - p - \frac{1}{2} p^2 + V_p^d(p)(\phi + r) p.$$

The following proposition solves the above ODE in closed form.

Proposition 4. Optimal deterministic contracts. Within the class of deterministic contracts, the value function $V^d(p)$ is quadratic

$$V^d(p) = -\frac{1}{2} A^d p^2 + B^d p. \tag{52}$$

The evolution of information rent, incentive, and effort are given by:

$$p_t^d = \frac{B^d}{A^d} e^{-\lambda t}, \quad \beta_t^d = \frac{1 + A^d \phi}{1 + ar\sigma^2} p_t^d, \text{ and } \mu_t^d = \beta_t^d - p_t^d = \frac{A^d \phi - ar\sigma^2}{1 + ar\sigma^2} p_t^d. \tag{53}$$

where $\lambda \equiv -\phi - r + \frac{1 + A^d \phi}{1 + ar\sigma^2} \phi > 0$, $B^d \equiv 1/\phi$ and

$$A^d \equiv \frac{(2\phi + r)ar\sigma^2 + r + \sqrt{(2\phi + r)^2 a^2 r^2 \sigma^4 + 2ar\sigma^2 [(\phi + r)^2 + \phi^2] + r^2}}{2\phi^2}. \tag{54}$$

Note that $\overline{p}^d$ in Equation (43) equals p_0^d, which maximizes the time-0 principal's value under deterministic contracts.

---

[23] This is because of the property of exponential martingale (recall $\{\beta\}$ being bounded): $\mathbb{E}_t \left[\exp\left(-ar\sigma \int_t^s \beta_u \, dB_u - \frac{(ar\sigma)^2}{2} \int_t^s \beta_u^2 \, du \right) \right] = 1$.

The above proposition shows that in the optimal deterministic contract, the information rent p_t^d, the incentive β_t^d, and the optimal effort μ_t^d all follow certain exponentially decaying paths (toward zero). Moreover, at $t=0$, from Equation (53), we have

$$\mu_0^d = \frac{A^d\phi - ar\sigma^2}{1+ar\sigma^2} \quad p_0^d = \frac{1-ar\sigma^2 p_0^d}{1+ar\sigma^2} < \frac{1}{1+ar\sigma^2} = \mu_0^{HM}.$$

Thus, the entire optimal effort path is below the Holmstrom and Milgrom (1987) benchmark.

The optimality of the front-loaded effort policies comes from the forward-looking nature of information rent. From the agent's incentive-compatibility condition in Equation (19), the belief manipulation effect implies that giving incentives later tends to make the agent shirk earlier, but not the other way around. This implies that later incentives are more costly than early ones, and, consequently, the optimal contract implements higher effort in earlier periods. Clearly, this result relies on the commitment ability in long-term contracting. Indeed, in Section 4.4 we show that equilibrium incentives and effort policies are constant over time when relationships are short term.

Both Prat and Jovanovic (2014) and our model find front-loaded incentives to be optimal. Because Prat and Jovanovic (2014) implements a constant effort, the forward-looking nature of information rent implies that the compensation contract has to offer front-loaded incentives.[24] Our model allows the optimal contract to adjust on the effort margin (not just incentives), and cheaper incentive provisions in earlier periods naturally push the optimal contract to implement a front-loaded effort profile.

The front-loaded effort policies also arise in models with career concerns (e.g., Gibbons and Murphy 1992; Holmstrom 1999), but through a distinct mechanism. There, agents in their early careers face *higher uncertainty* in their abilities, and thus work harder to impress the market (but the market will not be fooled in equilibrium, a standard signal-jamming problem). This force is not present in our stationary model, as the uncertainty of the profitability/ability (i.e., the posterior variance of θ_t) stays constant over time.

4.2 Value function and optimal policies
Now we return to the contracting space of fully stochastic incentives, and illustrate two qualitative properties of our optimal contract. First, similar to the case of deterministic contracts studied in Section 4.1, the fully stochastic optimal contract features front-loaded incentives. Second, the

---

[24] Both DeMarzo and Sannikov (2017) and Prat and Jovanovic (2014) assume that the effort cost is linear over the feasible interval $[0,1]$ and focus on implementing the highest effort level 1. In addition, Prat and Jovanovic (2014) study the nonstationary case in which the underlying profitability θ (as a parameter) never changes, and as time passes, both parties eventually learn the true profitability. In the Internet Appendix, we show that the pattern of time-decreasing effort pattern is robust to this assumption.
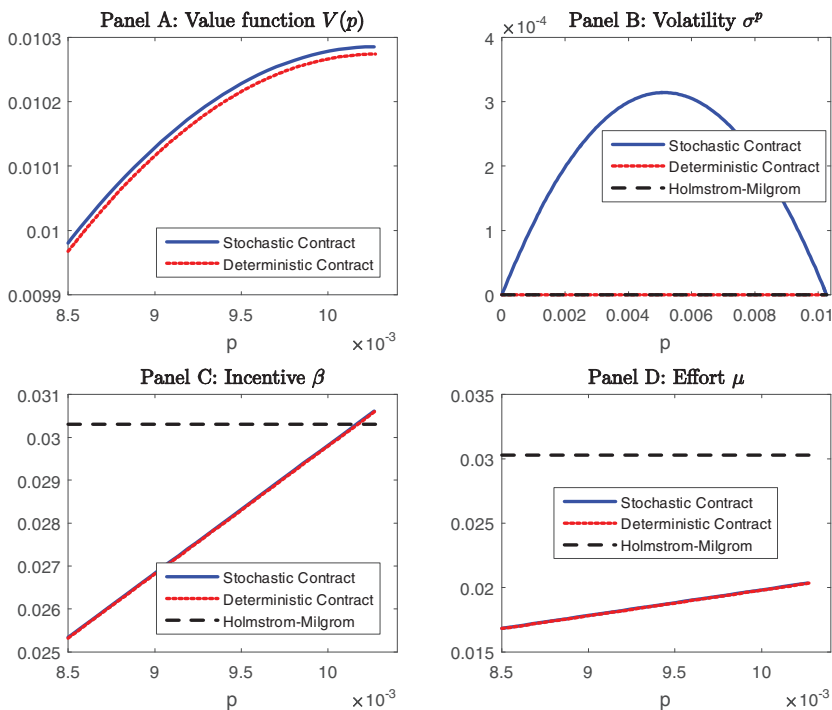
Figure 2
Value function and optimal policies in the optimal contract. Solid lines correspond to the optimal stochastic contract, and dashed lines correspond to the optimal deterministic contract. The parameters are $r = 0.5, a = 1, \sigma = 8$, and $\phi = 0.5$. The Holmstrom-Milgrom (1987) benchmark has $V^{HM} = \dfrac{1}{2r\left(1+ar\sigma^2\right)} = 0.03$ and $\beta^{HM} = \mu^{HM} = \dfrac{1}{1+ar\sigma^2} = 0.03$ under the parameter specification.

optimal management of the agent's information rent leads to an option-like feature in the optimal contract, that is, incentives rise after good performance. As explained, this option-like feature is explicitly ruled out in deterministic contracts.

4.2.1 How does the optimal stochastic contract help?. From now on we always refer to optimal policies, and without risk of confusion we omit the superscript asterisk. Figure 2 plots the value function $V(p)$, the optimal control $\{\beta(p), \sigma^P(p)\}$, and the associated optimal policy $\mu_t(p) = \beta_t(p) - p$ in solid lines. For comparison, in each panel we also plot the corresponding deterministic counterparts in dashed lines, and the Holmstrom and Milgrom (1987) benchmark in dotted lines.

The value delivered by the optimal stochastic contract must exceed the one under the deterministic counterpart, as shown in panel A in Figure 2. Panel B plots the volatility of the agent's information rent, σ^P, which is zero when the

contract is restricted to be deterministic. A positive σ^p in the optimal stochastic contract implies that the information rent rises after good performance shocks, an interesting property which will be discussed shortly.

What drives the stochastic contract to be superior to the deterministic one? It is because the stochastic contract implements a more efficient effort policy, closer to the higher Holmstrom and Milgrom (1987) effort benchmark. Panel C shows that incentives $\beta(p)$ sit above the deterministic counterparts for almost the entire range (and, thus, gets closer to the Holmstrom and Milgrom 1987 benchmark level), except for low p's, which are close to zero. A similar pattern holds for the implemented effort $\mu(p) = \beta(p) - p$ in panel D of Figure 2.

Interestingly, though not evident in Figure 2, when p is close to zero both the incentive $\beta(p)$ and effort $\mu(p)$ lie below their deterministic counterparts. This result is a robust feature of the model. Indeed, with the aid of asymptotic analysis (see Proposition B.1 in the companion Internet Appendix), one can analytically verify that the difference between the deterministic and stochastic contracts is negative by setting $p \simeq 0$. The seemingly counterintuitive result is rooted in the "option-like" feature in the optimal contract, to which we turn next.

4.2.2 Option-like incentives. In our model it is optimal to implement a history-dependent effort policy. This is surprising: As the posterior variance only changes over time deterministically in a standard CARA-normal setting with learning (in our stationary setting, the posterior variance is a constant in particular), usually the resultant equilibrium effort profile is a deterministic process as well (e.g., Holmstrom 1999).

To understand the economic mechanism that drives this result, we study how history-dependent effort policies improve over deterministic policies. To this end, we investigate the response of incentive β (or, effort μ) to unexpected shocks. This is captured by the diffusion term of $d\beta(p_t)$ (or, $d\mu(p_t)$), that is, $\beta'(p_t)\sigma^p dB_t$ (or $(\beta'(p_t)-1)\sigma^p dB_t$), and, as shown in the top panels in Figure 3, these diffusion terms are positive. There, we also plot the drift and diffusion for the key state variable p_t, that is, information rent. This interesting result implies that incentive (or, effort) rises following good performance, suggesting that the optimal contract is "convex" in output. In conclusion, in contrast to the Holmstrom and Milgrom (1987) benchmark where the optimal contract features a constant equity share, with learning the optimal contract has an option-like feature.

The optimality of this option-like feature is a result of reducing the agent's information rent in a long-term relation. As explained in Section 2.4, the thrust of endogenous learning in dynamic contracting is that the agent can (marginally) manipulate the principal's future belief downward by shirking today, and thus enjoy the potential information rent:

$$p_t = \frac{1}{u_c(c_t, \mu_t)} \mathbb{E}_t \left[\int_t^\infty \phi e^{-(\phi+r)(s-t)} \beta_s u_c(c_s, \mu_s) ds \right].$$
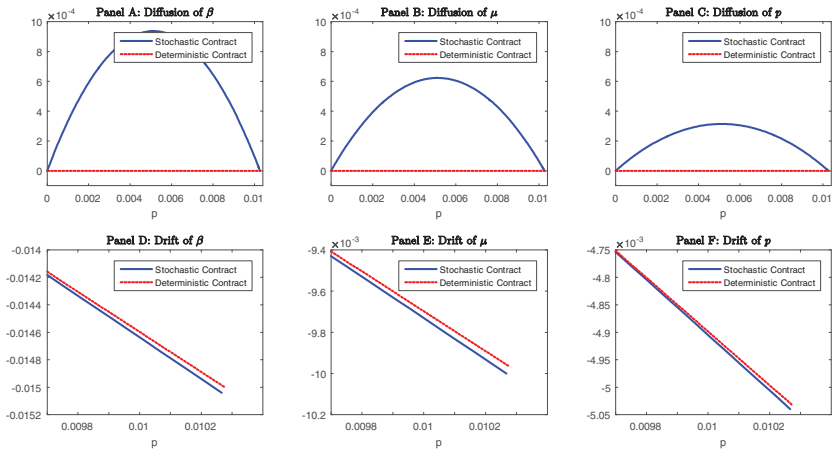
Figure 3
Diffusion and drift for incentives β_t, effort μ_t, and information rent p_t in the optimal contract. Solid lines correspond to the optimal stochastic contract, and dotted lines correspond to the optimal deterministic contract. The parameters are $r=0.5, a=1, \sigma=8$, and $\phi=0.5$.

The information rent captures the agent's additional future rewards when the principal mistakenly attributes the higher-profitability-driven good performance to the agent's effort, and this is why future incentives $\{\beta_s\}$ matter. Equally important, for a risk-averse agent, the amount of information rent also depends on his marginal utilities $u_c(c_s, \mu_s)$ when receiving manipulation benefits in those future states.

Because future incentives β_s and future marginal utilities $u_c(c_s, \mu_s)$ enter the information rent p_t multiplicatively, a negative correlation between β_s and $u_c(c_s, \mu_s)$ lowers p_t today. Intuitively, information rent can be reduced if the contract allocates greater belief manipulation benefits in states where the agent cares less. Interestingly, the option-like feature achieves this negative correlation. To see this, following a positive output shock, the agent becomes wealthier, implying a lower marginal utility $u_c(c_s, \mu_s)=-ar v_s$.[25] By making the optimal contract option-like, the principal raises incentives after good performance and thus imposes a negative correlation between incentives and the agent's marginal utility.

The option-like feature explains the intriguing result that the agent works less in the optimal stochastic contract than the deterministic one when p is close to zero, as discussed toward the end of Section 4.2.1. A positive diffusion of incentive β (effort μ) implies that the optimal contract allocates lower incentives in states with poor historical performance (and hence a high marginal utility). Because the information rent p is positively correlated with

---

[25] Formally, we have the evolution of marginal utility as $d(-ar v_t)=-ar\beta_t(-ar v_t)dB_t$, which has a negative diffusion coefficient in front of the performance shock dB_t.
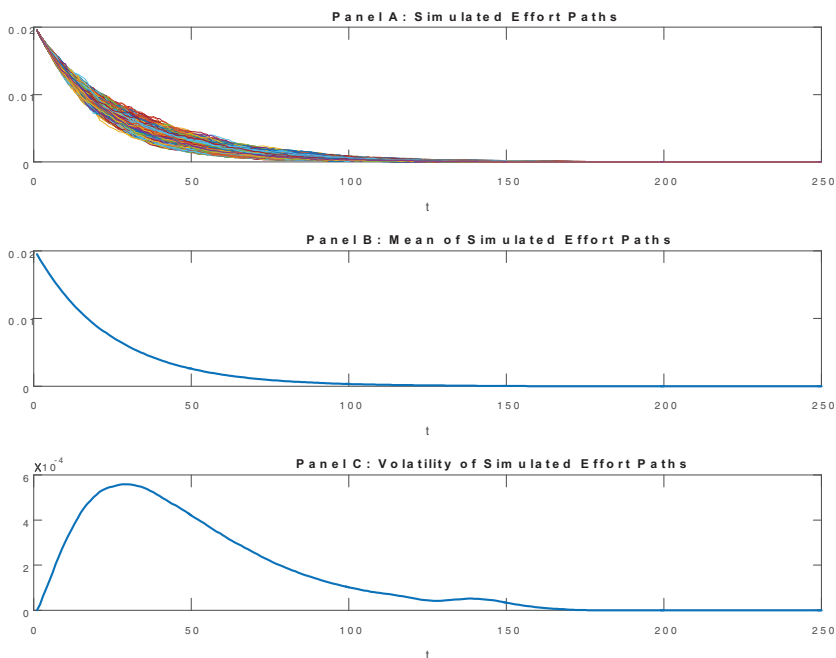
Figure 4
Simulated paths of effort policies in optimal contract. We simulate the model under the baseline parameter
specification: $r=0.5, a=1, \sigma=8, \phi=0.5$ for 10,000 rounds. In each round, we simulate a path of 250 months of
equilibrium efforts starting at the initial optimal value $\overline{p}$.

performance as indicated by panel B in Figure 2, the stochastic optimal contract
implements lower incentives in states with $p \simeq 0$.

4.3 Time-dcreasing effort policies
We have shown analytically in Section 4.1 that, because of forward-looking
information rent, the effort policy is decreasing over time in the optimal
deterministic contract. Not surprisingly, this pattern persists in the fully
stochastic optimal contract. Graphically, the front-loaded effort policy is
reflected by the negative drifts of incentives β and effort μ in the bottom
panels in Figure 2.

The feature of declining effort under the optimal contract is further confirmed
by Monte Carlo simulations. Specifically, we simulate the model under the
baseline parameter specification for 10,000 rounds. In each round we simulate
an equilibrium path of 250 months starting at the initial optimal state $\overline{p}$.
As shown in panel A, effort policies in all simulated paths tend to decrease
over time. The average equilibrium effort decreases to close to zero after 100
months (panel B). One interesting observation is that although the average
effort monotonically decreases over time, its volatility increases initially, and
then subsequently decreases (panel C). This is because the diffusion of p_t equals

Figure 5
Comparative statics with respect to the uncertainty parameter ϕ. Solid lines correspond to the optimal stochastic contract under the parameter specification: $r=0.5$, $a=1$, $\sigma=8$, and $\phi=0.5$. The dotted lines correspond to the optimal stochastic contract under the same parameter specification, except $\phi=0.75$.

zero at both boundaries 0 and $\overline{p}$, and achieves its maximum at an intermediate value in between.

4.3.1 Comparative statics on uncertainty. We study comparative static results with respect to the uncertainty parameter ϕ. Economically, ϕ measures the degree of informational uncertainty relative to cash flow risk in the model, which highlights our contribution to the literature ($\phi=0$ corresponds to the classic Holmstrom and Milgrom model).[26] Figure 5 plots the equilibrium outcomes under the baseline parameter specification $\phi=0.5$ (the solid line) as well as the one with $\phi=0.75$ (the dashed line), while keeping other parameters unchanged. A higher informational uncertainty can be considered as the situation with more severe agency conflicts, since all else being equal the agent enjoys greater information rents. This explains that, as a result of raising $\phi=0.5$ to $\phi=0.75$, the principal has a lower value function (panel A, lower $V(\cdot)$) and the agent works less (panel B, lower μ). To mitigate the agent's excessive

---

[26] Recall that $\phi = \Sigma_0^\theta / \sigma^2$ where Σ_0^θ is the prior uncertainty of θ and σ^2 is the volatility of cash flows.

information rents and curb his rent-seeking behavior, the optimal contract sets lower incentives (panel C, lower β) and becomes more option-like (panel D, higher σ^P).

4.4 Short-term contractual relationship

We want to emphasize that the two foregoing features, that is, front-loaded and option-like incentives, are due to the interaction between long-term contracting and learning. The case of observable θ_t shuts down learning, and, like in Holmstrom and Milgrom (1987), the optimal effort and incentives are constant over time. What if learning is present, but short-term (say, due to lack of commitment) contractual relationships are required?

Imagine the following setting with short-term contracting, in which a long-lived agent with unknown ability θ_t is working for a continuum of principals. At any time $t > 0$, there is one principal who signs a short-term incentive contract with the agent. The relationship, however, only lasts for the interval $[t, t+dt]$. The short-term contract consists of a fixed wage α_t, an incentive β_t, and the recommended effort μ_t, so that given date t belief $\mathbb{E}_t[\theta_t] = m_t$ the agent receives a compensation flow of

$$\alpha_t dt + \beta_t (dY_t - \mu_t dt - m_t dt)$$

at the end of period $t+dt$. Afterwards, the relationship breaks and the agent signs another contract $\{\alpha_{t+dt}, \beta_{t+dt}\}$ with another principal indexed by $t+dt$. Importantly, short-term relationships rule out inter-period commitment, implying each principal takes other principals' equilibrium offers as given.

For simplicity, to determine the history of fixed wages $\{\alpha_t\}$, we assign all the bargaining power to principals (as we have assumed in Section 1.1). We have the following proposition:

Proposition 5. Short-term relationships. Suppose that contractual relationships are short-term and principals have all the bargaining power. Then the equilibrium incentive β_t^{ST} is constant over time:

$$\beta_t^{ST} = \frac{\phi + r}{r + ar\sigma^2(\phi + r)} \quad \text{for all } t,$$

and the equilibrium effort μ_t^{ST} is constant over time as well

$$\mu_t^{ST} = \frac{r}{\phi + r} \beta_t^{ST} = \frac{r}{r + ar\sigma^2(\phi + r)} \quad \text{for all } t.$$

When the principals have all the bargaining power, Proposition 1 still applies to the agent's problem.[27] Thus, given today's incentive β_t^{ST} and future

---

[27] When the agent does not have any bargaining power, the proof in Proposition 5 shows that for the agent's problem, the short-term incentives $\left\{\beta_t^{ST}\right\}$ here play the same role as the incentives $\{\beta_t\}$ in long-term contracts analyzed in Proposition 1.

incentives $\{\beta_{t+s}^{ST} : s > 0\}$, the agent exerts $\mu_t^{ST} = \beta_t^{ST} - p_t^{ST}$, where p_t^{ST}, the properly discounted future incentives $\{\beta_{t+s}^{ST} : s > 0\}$, is defined analogously like in Equation (20). The time-*t* principal takes p_t^{ST} as given and maximizes the expected output $\beta_t^{ST} - p_t^{ST} + m_t$, minus the total compensation which is the sum of the effort cost $(\beta_t^{ST} - p_t^{ST})^2/2$ and the risk compensation $\frac{1}{2}ar\sigma^2(\beta_t^{ST})^2$. Ignoring the given project quality m_t, the time-*t* principal maximizes the flow payoff in Equation (33) only:

$$\max_{\beta_t^{ST}} (\beta_t^{ST} - p_t^{ST}) - \frac{1}{2}(\beta_t^{ST} - p_t^{ST})^2 - \frac{ar\sigma^2}{2}(\beta_t^{ST})^2 \Rightarrow \beta_t^{ST} = \frac{1+p_t^{ST}}{1+ar\sigma^2}. \quad (55)$$

Stationarity implies that both β_t^{ST} and p_t^{ST} are constants, and the result in Proposition 5 follows.

Intuitively, without commitment, in short-term contracting each principal at different points of time solves her individual myopic optimization problem in (55). In contrast, with long-term contracting, a long-lived single principal not only maximizes the flow payoff in (55) but also takes into account the effect of β_{t+s} on the forward-looking information rent p_t.[28] This forward-looking force in the full commitment environment, combined with learning, makes the optimal effort policy time decreasing and stochastic.

4.5 Empirical implications

4.5.1 Labor and CEO compensation. Our model has a few key empirical implications. First, the optimal long-term contract with learning implements front-loaded effort policies. This is consistent with the findings in Medoff and Abraham (1981), who measure the productivity of different age groups and find that young people are more productive, controlling for job categories. Their findings support the prediction that young workers supply more labor if workers in the same job category have roughly similar abilities.

Our model suggests that it is more efficient to assign higher incentives after good performance, because the agent has a lower marginal utility at that time (and hence less information rent). This option-like feature of the optimal contract lends support to the pervasive use of option-based compensation in practice (e.g., Hall and Liebman 1998). There is further empirical support for this prediction: Core and Guay (1999) find that the annual grant of options and stocks to a CEO is increasing in past stock returns, and Bergman and Jenter (2006) document that option and stock grants per manager are increasing in past

---

[28] This result is in contrast to that of Fudenberg, Holmstrom, and Milgrom (1990), whose model does not contain learning. They show that, with dynamic moral hazard only, the optimal long-term contract can be implemented by short-term ones under CARA preferences. In a way, their result suggests that commitment itself–when learning is absent–is not that important. In contrast, our model shows that the commitment in long-term contracting is important because of the long-lasting belief manipulation effect with endogenous learning.

stock returns. More recently, He et al. (2014) find that managerial incentives increase with past firm-level profitability.[29]

Pushing this point a bit further, our model also implies that managerial incentives should be procyclical at the aggregate level. The idea is simple: Although aggregate economic conditions should be indexed out in the optimal contract, the fact that the agent tends to have a marginal utility in good times implies that it is relatively cheaper to assign incentives there. This prediction is consistent with the empirical finding in Eisfeldt and Rampini (2008), who show that the Hodrick-Prescott-filtered executive compensation is remarkably procyclical.

Last, as suggested by the comparative static result in panel D in Figure 5, industries or firms with higher uncertainty should have more option-based contracts for managerial compensation. There is no doubt that, compared to traditional industries, new-economy firms (such as computer, software, the Internet, or telecommunication companies) tend to be associated with higher uncertainty. Consistent with our model predictions, both Ittner, Lambert, and Larcker (2003) and Murphy (2003) find that new-economy firms indeed grant more stock options to their managers.

There is one caveat in linking our optimal contracting results to compensation contracts in practice. As emphasized, we focus on long-term contracting with full commitment, which is theoretically appealing because it gives the upper bound of other long-term relations with partial commitment. In practice, without full commitment career concerns (Gibbons and Murphy 1992; Holmstrom 1999) are another theoretically important and empirically relevant force, especially when the labor market is mobile and agents/workers can easily move.[30] Therefore, our model applies more to the situation where human capital is more firm specific, and thus long-term job security is a primary concern.

4.5.2 Incentive contracts in asset management. Our analysis also sheds light on the difference between compensation contracts observed in the hedge fund and mutual fund industries. Hedge funds tend to compensate their managers based on long-term "explicit incentives," and the option feature is embedded in the widely used "2-20" and high-water-mark contracts. That hedge fund contracts exhibit option-like features may well be related to learning about the manager's persistent unobservable ability in asset trading. In contrast,

---

[29] He et al. (2014) emphasize the different roles of uncertainty and risk in shaping optimal executive compensation, which potentially helps distinguish our uncertainty-based mechanism from the leading explanation that the purpose of option-based compensation is to provide CEOs with incentives for taking risk (e.g., Stulz and Smith 1985).

[30] Although both Gibbons and Murphy (1992) and our paper feature an optimal front-loaded effort policy, the predictions regarding optimal incentive profiles are different. Due to career concerns, in Gibbons and Murphy (1992) the agent works hard even without high-powered in-job incentives. In contrast, all incentives in our model are from the long-term contract, and the front-loaded effort profile requires a front-loaded incentive contract. This is a common feature in a dynamic contracting model with full commitment and learning, such as in Prat and Jovanovic (2014) and DeMarzo and Sannikov (2017).

mutual fund managers are often compensated by "implicit incentives," which are management fees proportional to assets under management. This is in line with the result that linear compensation remains optimal for short-term contracting, a case analyzed in Section 4.4 in which both parties cannot commit to a long-term relationship.[31] The difference between these two industries seems to be consistent with the casual observation that, compared with the hedge fund industry, there is greater job mobility in the mutual fund industry because the human capital of mutual fund managers is more fungible across different funds.[32]

Within the hedge fund industry, some interesting empirical predictions can be made based on our main theoretical findings that (a) relative to shorter-term contracting, the agent in a longer-term contracting relationship shirks more for the purpose of information rent extraction (Equation (19) in Proposition 1), and (b) to mitigate such motives, incentives tend to increase following good performance. The former finding (a) predicts that hedge fund managers in a longer-term contracting relationship (e.g., a longer lock-in period) tend to work less, all else being equal (e.g., fixing incentives). Translating to observable measures, our model predicts that hedge fund managers with a longer-term contracting relationship tend to be associated with worse fund returns, all else being equal. The latter point (b) speaks to the relationship between profit sharing and high-water-mark. Because the high-water mark increases following superior performance, it is optimal to assign managers a greater share of profits whenever his high-water mark rises. This suggests that the current practice of fixing the profit share, say at the 20% level, could be improved. Of course, this conclusion might not be robust to other first-order factors (say, limited liability) that are missing in our analysis.

5. Concluding Remarks

We introduce profitability uncertainty into the model of Holmstrom and Milgrom (1987) and study optimal long-term contracting with endogenous learning. Although the principal and the agent hold the same belief about project profitability along the equilibrium path, the agent's potential deviation by exerting effort below the recommended level leads to potential long-lasting belief divergence between both parties and thus a "hidden information" problem. By utilizing the convenient property of CARA preferences, we show that optimal contracting can be reformulated to a dynamic programming problem with only one state variable, and we characterize the optimal contract

---

[31] It is worth noting that compared with the canonical Berk and Green (2004) model for mutual funds, our paper features symmetric learning but abstracts away the endogenous fund flows (and hence the endogenous fund size or assets under management).

[32] In practice, although it is common for hedge fund managers to sign a so-called "noncompete clause" when hired, mutual fund managers rarely sign these types of clauses, especially for funds offering passive investment products (e.g., dimensional).

by the solution to an ODE. We show that optimal effort decreases with tenure, and the optimal contract exhibits an option-like feature in the sense that incentives/effort rise after positive performance shocks. These two properties rely on a combination of learning and long-term contracting, as we show the resultant equilibrium effort profile is constant over time in the case of either long-term contracting without learning (Holmstrom and Milgrom 1987) or short-term contracting with learning (Holmstrom 1999).

Although we are able to give a full characterization of the optimal long-term contract only under a specific setting (e.g., CARA preferences, Gaussian processes), the foregoing qualitative results, that is, front-loaded effort policy and option-like feature, likely are robust to more general settings. The main reason we think the results extend is that the economic force behind these results do not depend on CARA preferences or Gaussian processes. The agent's information rent due to belief manipulation, that is, the agent's inclination to shirk to distort the principal's future belief downward, is general in any long-term contracting environment with learning. Because later incentives enter the agent's forward-looking information rent in earlier periods, but not the other way around, it is more efficient to provide incentives early on, rendering the optimality of front-loaded effort policy. Additionally, the option-like feature comes from the fact that the agent is risk averse, so that the marginal value of belief manipulation benefits is lower after good performance.

Appendix A: Proofs

A.1 Proof for Lemma 1
The argument is similar to that of He (2011). Consider any contract $\Pi = \{c_t, \mu_t\}$ that induces an optimal policy $\{c_t^\star, \mu_t^\star\}$ from the agent with a value $v_0^\star$, so that

$$v_0^\star = \mathbb{E}^{\mu^\star}\left[\int_0^\infty e^{-rt} u\left(c_t^\star, \mu_t^\star\right) dt\right]$$

$$\text{s.t. } dY_t = \left(\mu_t^\star + m_t^{\mu^\star}\right) dt + \sigma \, dB_t^{\mu^\star},$$

$$dS_t = r S_t dt + c_t dt - c_t^\star dt \text{ with } S_0 = 0.$$

The principal knows the resultant optimal effort policy $\{\mu_t^\star\}$, and she updates her belief according to $\{\mu_t^\star\}$, rather than the recommended effort policy $\{\mu_t\}$. From the agent's budget equation, we have

$$S_t = \int_0^t e^{r(t-s)}\left(c_s - c_s^\star\right) ds,$$

which gives the agent's optimal savings path. Note that if S_t is bounded, then the transversality condition holds for all measures induced by any feasible effort policies.

By invoking the replication argument similar to revelation principle, we consider giving the agent a direct contract $\Pi^\star = \{c_t^\star, \mu_t^\star\}$. Clearly, taking consumption-effort policy $\{c_t^\star, \mu_t^\star\}$ is feasible for the agent with no private-savings. Now we show that $\{c_t^\star, \mu_t^\star\}$ is optimal for the agent given this contract.

Suppose, counter-factually, that given the contract $\Pi^\star$, the agent finds that $\{c_t', \mu_t'\}$ yields a strictly higher payoff $v_0' > v_0^\star$ in her problem, with associated savings path

$$S_t' = \int_0^t e^{r(t-s)}\left(c_s^\star - c_s'\right) ds,$$

which satisfies the transversality condition. Formally, we have

$$v_0' = \mathbb{E}^{\mu'}\left[\int_0^\infty e^{-rt} u\left(c_t', \mu_t'\right) dt\right] > v_0^\star$$

$$\text{s.t. } dY_t' = \left(\mu_t' + m_t^{\mu'}\right) dt + \sigma dB_t^{\mu'},$$

$$dS_t' = r S_t' dt + c_t^\star dt - c_t' dt \text{ with } S_0 = 0.$$

Now we construct an contradiction to the preassumption that "given $\Pi = \{c_t, \mu_t\}$ the agent's optimal policy is $\{c_t^\star, \mu_t^\star\}$ with a value of $v_0^\star$." Suppose that given $\Pi = \{c_t, \mu_t\}$, the agent takes the policy $\{c_t', \mu_t'\}$ instead of $\{c_t^\star, \mu_t^\star\}$ which is claimed to be optimal. Because $v_0' > v_0^\star$ this alternative policy strictly dominates $\{c_t^\star, \mu_t^\star\}$; the only thing left is to verify whether the consumption plan is feasible given some saving policy. But, the saving policy $S_t'' = S_t + S_t' = \int_0^t e^{r(t-s)}\left(c_s - c_s'\right) ds$ achieves $\{c_t'\}$ given the income process $\{c_t\}$, because

$$dS_t'' = dS_t + dS_t' = r S_t dt + c_t dt - c_t^\star dt + \left[r S_t' dt + c_t^\star dt - c_t' dt\right],$$

$$= r\left(S_t + S_t'\right) dt + c_t dt - c_t' dt,$$

$$= r S_t'' + c_t dt - c_t' dt,$$

which also satisfies the transversality condition $\lim_{T \to \infty} \mathbb{E}\left[e^{-rT} S_T''\right] = 0$ if both S_t and S_t' satisfy the transversality condition. Thus, given the original contract Π, the saving rule $\{S_t''\}$ supports $\{c_t', \mu_t'\}$ but delivers a strictly higher payoff v_0'. This contradicts with the optimality of $\{c_t^\star, \mu_t^\star\}$ under the contract Π.

Finally, because the principal knows that $\{c_t^\star, \mu_t^\star\}$ is optimal for the agent, the principal still correctly knows the agent's actual optimal effort policy $\{\mu_t^\star\}$ and thus perform the correct Bayesian updating, and her payoff is the same as that under the contract $\Pi = \{c_t, \mu_t\}$. Hence it is without loss of generality to focus on contracts that are incentive-compatible and no-savings.

A.2 Proof for Lemma 2

Fix any constant S. Given any savings $S_t = S$ and a contract $\Pi = \{c\}$, from time-t on the agent's problem is

$$\max_{\{\widehat{c}_s\}, \{\widehat{\mu}_s\}} \quad \mathbb{E}^{\widehat{\mu}}\left[\int_t^\infty -\frac{1}{a} e^{-a\left(\widehat{c}_s - \frac{1}{2}\widehat{\mu}_s^2\right) - r(s-t)} ds\right], \tag{A1}$$

$$s.t. \qquad dS_s = r S_s ds + c_s ds - \widehat{c}_s ds, \ S_t = S, \ s > t,$$

$$dY_s = \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right) dt + \sigma dB_s^{\widehat{\mu}},$$

given his information set. Note that the agent will learn actively. $\{c_s^\star, \mu_s^\star\}$ is the solution to the above problem, and $v_t(S; \Pi)$ is the agent's value.

Now consider the problem with $S = 0$, which is the continuation payoff along the equilibrium path:

$$\max_{\{\widehat{c}_s\}, \{\widehat{\mu}_s\}} \quad \mathbb{E}^{\widehat{\mu}}\left[\int_t^\infty -\frac{1}{a} e^{-a\left(\widehat{c}_s - \frac{1}{2}\widehat{\mu}_s^2\right) - r(s-t)} ds\right],$$

$$s.t. \qquad dS_s = r S_s ds + c_s ds - \widehat{c}_s ds, \ S_t = 0, \ s > t,$$

$$dY_s = \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right) dt + \sigma dB_s^{\widehat{\mu}},$$

We claim that the solution to this problem is $\{c_s^\star - rS, \mu_s^\star\}$, and therefore the value is $v_t(0; \Pi) = e^{arS} v_t(S; \Pi)$. There are two steps to show this. First, this solution is feasible. Second, suppose that

there exists another policy $\{\widehat{c}'_s, \widehat{\mu}'_s\}$ that is superior to $\{c^\star_s - rS, \mu^\star_s\}$, so that the associated value $v'_t(0; \Pi) > e^{-arS} v_t(S; \Pi)$. Consider $\{\widehat{c}'_s + rS, \widehat{\mu}'_s\}$, which is feasible to the problem in Equation (A1). Under this plan, however, the agent's objective is

$$e^{-arS} \cdot \max \mathbb{E}^{\widehat{\mu}}_t \left[\int_t^\infty -\frac{1}{a} e^{-\gamma\left(\widehat{c}'_s - \frac{1}{2}\widehat{\mu}'^2_s\right) - r(s-t)} ds \right] = e^{-arS} v'_t(0; \Pi) > v_t(S; \Pi),$$

which contradicts with the optimality of $\{c^\star_s, \mu^\star_s\}$. As a result, $v_t(S; \Pi) = e^{-arS} v_t(0; \Pi)$.

A.3 Proof for Proposition 1

$\{c, \mu\}$ denotes the agent's (proposed) optimal consumption-effort policy given the compensation contract that satisfies the first-order condition stated in the proposition. The agent's continuation payoff v_t follows $dv_t = (-arv_t)\beta_t \sigma dB^\mu_t$ where $\{\beta\}$ are incentives specified by the contract. We will use the following property of $\{v\}$ later:

$$v_t = v_0 \exp\left(\int_0^t ar\beta_u \sigma dB^\mu_u - \int_0^t 0.5a^2 r^2 \beta^2_u \sigma^2 du \right) = v_0 - \int_0^t arv_s \beta_s \sigma dB^\mu_s. \qquad (A2)$$

It is to show that when $|\beta| < M$ is bounded, v_t follows a martingale (Revuz and Yor 1999, 139). This also verifies that v_t is the agent's equilibrium continuation payoff following the equilibrium consumption-effort policy.

We now establish the necessary conditions stated in the proposition by considering deviation strategies on effort and consumption polices respectively. First consider the deviation policy in effort, that is, $\{\widehat{c}_t, \widehat{\mu}_t\} = \{c_t, \mu_t + \varepsilon\delta_t\}$, where the deviation policy $\{\delta_t \neq 0\}$ is arbitrary. Due to CARA preference, we have

$$u(\widehat{c}_t, \widehat{\mu}_t) = u(c_t, \mu_t) e^{a\mu_t \varepsilon\delta_t + 0.5a\varepsilon^2\delta^2_t}.$$

The agent's value under the deviation policy indexed by ε is simply

$$\widehat{v}_0(\varepsilon) \equiv \mathbb{E}^{\widehat{\mu}}_0 \left[\int_0^\infty e^{-rt} rv_t e^{a\mu_t \varepsilon\delta_t + 0.5a\varepsilon^2\delta^2_t} dt \right];$$

note that the expectation is under the measure induced by deviating effort profile $\widehat{\mu}$. Equations (3), (5), and (13) imply that the effort profile $\widehat{\mu}$, which depends on ε, induces a change of measure relative to μ by

$$\sigma dB^\mu_t - \sigma dB^{\widehat{\mu}}_t = \left(\widehat{\mu}_t + m^{\widehat{\mu}}_t - \mu_t - m^\mu_t \right) dt = \left(\varepsilon\delta_t + m^{\widehat{\mu}}_t - m^\mu_t \right) dt, \qquad (A3)$$

$$= \left(\varepsilon\delta_t - \phi \int_0^t e^{-\phi(t-\tau)}(\varepsilon\delta_\tau) d\tau \right) = (\varepsilon\delta_t - \varepsilon\Delta_t) dt,$$

where, like in (13), we write

$$\Delta_s \equiv \phi \int_0^s e^{-\phi(s-u)} \varepsilon\delta_u du. \qquad (A4)$$

Hence we introduce the exponential martingale N_t, indexed by ε:

$$N_t(\varepsilon) \equiv \exp\left(\int_0^t \frac{\varepsilon\delta_s - \varepsilon\Delta_s}{\sigma} dB^\mu_s - \int_0^t \frac{(\varepsilon\delta_s - \varepsilon\Delta_s)^2}{2\sigma^2} ds \right), \text{ with } N_0(\varepsilon) = 1,$$

so that according to Girsanov theorem, we have

$$\widehat{v}_0(\varepsilon) = \mathbb{E}^{\widehat{\mu}}_0 \left[\int_0^\infty e^{-rt} rv_t e^{a\mu_t \varepsilon\delta_t + 0.5a\varepsilon^2\delta^2_t} dt \right] = \mathbb{E}^\mu_0 \left[\int_0^\infty N_t(\varepsilon) e^{-rt} rv_t e^{a\mu_t \varepsilon\delta_t + 0.5a\varepsilon^2\delta^2_t} dt \right].$$

Now, we take derivative of $\widehat{v}_0(\varepsilon)$ with respect to ε, and evaluate it at $\varepsilon = 0$. Because

$$\frac{dN_t(\varepsilon)}{d\varepsilon}\bigg|_{\varepsilon=0} = N_t(\varepsilon)\cdot\left[\int_0^t \frac{\delta_s - \Delta_s}{\sigma}dB_s^\mu - \int_0^t \frac{\varepsilon(\delta_s-\Delta_s)^2}{\sigma^2}ds\right]\bigg|_{\varepsilon=0} = \int_0^t \frac{\delta_s-\Delta_s}{\sigma}dB_s^\mu,$$

we have

$$\frac{d\widehat{v}_0(\varepsilon)}{d\varepsilon}\bigg|_{\varepsilon=0} = \left\{\mathbb{E}_0^\mu\left[\int_0^\infty \frac{dN_t(\varepsilon)}{d\varepsilon}\cdot e^{-rt}rv_t e^{a\mu_t\varepsilon\delta_t + 0.5a\varepsilon^2\delta_t^2}dt\right]\right.$$

$$\left.+\mathbb{E}_0^\mu\left[\int_0^\infty N_t(\varepsilon)e^{-rt}rv_t\frac{d\left(e^{a\mu_t\varepsilon\delta_t + 0.5a\varepsilon^2\delta_t^2}\right)}{d\varepsilon}dt\right]\right\}\bigg|_{\varepsilon=0},$$

$$=\underbrace{\mathbb{E}_0^\mu\left[\int_0^\infty\left(\int_0^t \frac{\delta_s-\Delta_s}{\sigma}dB_s^\mu\right)\cdot e^{-rt}rv_t dt\right]}_{AA} + \mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt}arv_t\mu_t\delta_t dt\right]. \qquad (A5)$$

The first term AA equals to (using (A2)):

$$AA = \mathbb{E}_0^\mu\left[\int_0^\infty re^{-rt}\left(\int_0^t \frac{\delta_s-\Delta_s}{\sigma}dB_s^\mu\right)\cdot\left(v_0 - \int_0^t arv_s\beta_s\sigma dB_s^\mu\right)dt\right],$$

$$=\mathbb{E}_0^\mu\left[\int_0^\infty re^{-rt}\left(\int_0^t \frac{\delta_s-\Delta_s}{\sigma}dB_s^\mu\right)\left(-\int_0^t arv_s\beta_s\sigma dB_s^\mu\right)dt\right]$$

$$=-\mathbb{E}_0^\mu\left[\int_0^\infty re^{-rt}\int_s^\infty ar\beta_s v_s(\delta_s-\Delta_s)ds\right],$$

$$=-\mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt}ar\beta_t v_t(\delta_t-\Delta_t)dt\right],$$

where the last line uses change of order of integration. Plugging this result back into (A5), and using (A2), we have

$$\frac{d\widehat{v}_0(\varepsilon)}{d\varepsilon}\bigg|_{\varepsilon=0} = arv_0\mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt-\int_0^t ar\beta_u\sigma dB_u^\mu - \int_0^t 0.5a^2r^2\beta_u^2\sigma^2 du}[(\mu_t-\beta_t)\delta_t + \beta_t\Delta_t]dt\right],$$

$$=arv_0\mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt-\int_0^t ar\beta_u\sigma dB_u^\mu - \int_0^t 0.5a^2r^2\beta_u^2\sigma^2 du}[(\mu_t-\beta_t)\delta_t]dt\right],$$

$$+arv_0\underbrace{\mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt-\int_0^t ar\beta_u\sigma dB_u^\mu - \int_0^t 0.5a^2r^2\beta_u^2\sigma^2 du}\beta_t\Delta_t dt\right]}_{BB}. \qquad (A6)$$

Let us furhter simplify the term BB in (A6). $Z_t \equiv -\int_0^t ar\beta_u\sigma dB_u^\mu - \int_0^t 0.5a^2r^2\beta_u^2\sigma^2 du$ so that $\mathbb{E}_s^\mu[\exp(Z_t)] = \exp(Z_s)$ for $s < t$ under the condition of $|\beta| < M$ being bounded (see, e.g., Revuz and Yor 1999, 139). Then by changing the order of integration, we obtain (recall (A4)):

$$BB = \mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt+Z_t}\beta_t\Delta_t dt\right] = \mathbb{E}_0^\mu\left[\int_0^\infty e^{-rt+Z_t}\beta_t\phi\left(\int_0^t e^{-\phi(t-s)}\delta_s ds\right)dt\right],$$

$$=\mathbb{E}_0^\mu\left[\int_0^\infty \delta_s\phi\mathbb{E}_s^\mu\left[\int_s^\infty e^{-rt+Z_t}e^{-\phi(t-s)}\beta_t dt\right]ds\right],$$

$$=\mathbb{E}_0^\mu\left[\int_0^\infty e^{-rs+Z_s}\delta_s\phi\mathbb{E}_s^\mu\left[\int_s^\infty e^{Z_t-Z_s}e^{-(r+\phi)(t-s)}\beta_t dt\right]ds\right].$$

As a result, we have

$$\frac{d\widehat{v}_0(\varepsilon)}{\varepsilon}\Big|_{\varepsilon=0}=arv_0\mathbb{E}_0^{\mu}\left[\int_0^{\infty}e^{-rt+Z_t}\delta_t\left(\mu_t-\beta_t+\phi\mathbb{E}_t^{\mu}\left[\int_t^{\infty}\beta_s e^{-(r+\phi)(s-t)}e^{Z_s-Z_t}ds\right]\right)dt\right].$$

This implies that any incentive-compatible and no-saving policy must satisfy

$$\mu_t=\beta_t-\phi\mathbb{E}_t^{\mu}\left[\int_t^{\infty}\beta_s e^{-(r+\phi)(s-t)}e^{Z_s-Z_t}ds\right],a.s.$$

Otherwise, we can choose negative δ_t when $\mu_t>\beta_t-\phi\mathbb{E}_t^{\mu}\left[\int_t^{\infty}\beta_s e^{-(r+\phi)(s-t)}e^{Z_s-Z_t}ds\right]$, and positive δ_t when $\mu_t<\beta_t-\phi\mathbb{E}_t^{\mu}\left[\int_t^{\infty}\beta_s e^{-(r+\phi)(s-t)}e^{Z_s-Z_t}ds\right]$ (note that arv_0 is negative). Then a deviation strategy $\{c_t,\mu_t+\varepsilon\delta_t\}$ for sufficiently small ε will be profitable, leading to a contradiction.

The necessary conditions for the equilibrium consumption plan are much more standard. Fixing μ, it is easy to show that the necessary first-order condition for the agent's consumption-saving problem is that his marginal utility from consumption, that is, $u_c(c_t,\mu_t)$, follows a martingale. Because $u_c(c_t,\mu_t)=-au(c_t,\mu_t)$ for exponential utility, and $v_t=\mathbb{E}_t^{\mu}\left[\int_t^{\infty}e^{-r(s-t)}u(c_t,\mu_t)dt\right]$, the result follows easily. Q.E.D.

A.4 Appendix for Section 2.4.1

To see that the first line is the agent's total payoff from time t onwards given any effort policy $\widehat{\mu}$ and $\widehat{c}$, define $G(t)\equiv\int_0^t e^{-rs}u_s ds+e^{-rt}v_t$ and $G(\infty)$ is the agent's total payoff. Due to private savings, $u_s=rv_s$, and we have $dG(t)=e^{-rt}dv_t$. Therefore the total payoff (inflated by e^{rt}) is $\mathbb{E}_t^{\widehat{\mu}}\left[e^{rt}G(\infty)\right]=e^{rt}G(t)+\mathbb{E}_t^{\widehat{\mu}}\left[\int_t^{\infty}e^{-r(s-t)}dv_s\right]$, which is $u(\widehat{c}_t,\widehat{\mu}_t)dt+v_t+\mathbb{E}_t^{\mu}\left[\int_t^{\infty}e^{-r(s-t)}dv_s\right]$ by ignoring utilities occurring before t. Under equilibrium effort, v_s is martingale, and thus $\mathbb{E}_t^{\mu}\left[\int_t^{\infty}e^{-r(s-t)}dv_s\right]=0$.

A.5 Appendix for Section 3.1

First of all, as m_t follows a martingale with $m_0=0$, we have

$$\mathbb{E}\left[\int_0^{\infty}e^{-rt}dY_t\right]=\mathbb{E}\left[\int_0^{\infty}e^{-rt}(\mu_t+m_t)dt\right]=\mathbb{E}\left[\int_0^{\infty}e^{-rt}\mu_t dt\right].$$

And, for wage cost, we have,

$$\mathbb{E}\left[\int_0^{\infty}-e^{-rt}\frac{1}{a}\ln(-arv_t)dt\right]=\mathbb{E}\left[\int_0^{\infty}\frac{\ln(-arv_t)}{ar}d\left(e^{-rt}\right)\right]$$

$$=-\frac{\ln(-arv_0)}{ar}-\mathbb{E}\left[\int_0^{\infty}e^{-rt}\frac{d\ln(-v_t)}{ar}\right]$$

$$=-\frac{\ln(-arv_0)}{ar}+\mathbb{E}\left[\int_0^{\infty}\frac{1}{2}e^{-rt}ar\sigma^2\beta_t^2 dt\right],$$

where we have used Equation (30) in the last equality and the fact that β_t's are bounded (so $\mathbb{E}\left[\int_0^{\infty}e^{-rt}\frac{dv_t}{v_t}dt\right]=-\mathbb{E}\left[\int_0^{\infty}e^{-rt}ar\beta_t\sigma dB_t\right]=0$).

A.6 Appendix for Section 3.2

Now we derive the evolution of p_t. Define $\widetilde{p}_t\equiv(-arv_t)p_t$. The definition of p_t (see Equation (20)) implies $\widetilde{p}_t=\mathbb{E}_t\left[\int_t^{\infty}\phi(-arv_s\beta_s)e^{-(\phi+r)(s-t)}ds\right]$. Based on the martingale representation theorem applied to $\widetilde{p}_t e^{-(\phi+r)t}$, there exists some progressively measurable process $\widetilde{\gamma}_t$ so that

$$d\widetilde{p}_t=((r+\phi)\widetilde{p}_t-\phi(-arv_t\beta_t))dt+\widetilde{\gamma}_t\sigma dB_t.$$

Note that $dv_t = (-ar\beta_t v_t)\sigma dB_t$ and define $\gamma_t \equiv \widetilde{\gamma}_t/(-arv_t)$, we have

$$dp_t = \frac{d\widetilde{p}_t}{(-arv_t)} + \frac{\widetilde{p}_t}{(arv_t^2)}dv_t + \frac{1}{2}\left(-\frac{2\widetilde{p}_t}{arv_t^3}(dv_t)^2 + \frac{2}{arv_t^2}\langle d\widetilde{p}_t, dv_t\rangle\right)$$

$$= \frac{((r+\phi)\widetilde{p}_t - \phi(-arv_t\beta_t))dt + \widetilde{\gamma}_t\sigma dB_t}{(-arv_t)} + \frac{\widetilde{p}_t}{(arv_t^2)}(-arv_t\beta_t)\sigma dB_t$$

$$\quad - \frac{\widetilde{p}_t}{arv_t^3}(-arv_t\beta_t)^2\sigma^2 dt + \frac{1}{arv_t^2}\sigma^2(-arv_t\beta_t)\widetilde{\gamma}_t dt$$

$$= \left[(r+\phi)p_t - \phi\beta_t + ar\sigma^2\beta_t(\gamma_t + arp_t\beta_t)\right]dt + \underbrace{\sigma(\gamma_t + arp_t\beta_t)}_{\equiv \sigma_t^p}dB_t$$

$$\equiv \left[(r+\phi)p_t + \beta_t\left(ar\sigma\sigma_t^p - \phi\right)\right]dt + \sigma_t^p dB_t.$$

A.7 Proof for Proposition 2

We will prove the key Proposition 2 in several steps. As a preparation, we will first show that for any feasible contract, p_t is bounded in $[-M_p, -M_p]$. In other words, in Proposition 2, $V(p)$ is strictly concave over an compact interval $[-M_p, -M_p]$.

A.7.1 Step 0: Relaxed problem and parameter restrictions.

Recall that we restrict the feasible incentive slopes $\{\beta_t\}$ to be bounded, that is, some sufficiently large constant M exists such that $\beta_t \in [-M, M]$. This is in the same spirit as imposing the transversality condition, because given bounded incentives $\{\beta_t\}$, the promised information rent p–as expected future discounted incentives—is also bounded. This boundedness result holds for any feasible contract, not just the optimal one.

Define $\pm M_p \equiv \pm\frac{\phi M}{\phi+r}$. The following lemma shows that the information rent p_t is bounded in $[-M_p, +M_p]$ if incentives $\{\beta_t\}$ are bounded in $[-M, M]$. Further, the boundaries for p_t, if ever reached (might be off equilibrium), are absorbing.

Lemma 3. Suppose that $\beta_t \in [-M, M]$ where M is a given constant. Then the state variable p_t reaches $\pm M_p$ if, and only if, $\beta_s = \pm M, \forall s \geq t$, which implies that $\pm M_p$ are absorbing states for p. As a result, when $p = \pm M_p$, $V(\pm M_p)$ is quadratic in M like in (A13).

Proof. Suppose that the control variable is constrained such that $\beta_t \in [-M, M]$, where $M > 0$ is an arbitrarily large, but fixed constant. Recall the definition of p_t, and we have

$$p_t = \mathbb{E}\left[\int_0^\infty \phi\beta_t e^{-(\phi+r)t} e^{\left(-\int_0^t ar\beta_v\sigma dB_v - \frac{1}{2}\int_0^t a^2 r^2\beta_v^2\sigma^2 dv\right)}dt\right]$$

$$\leq \mathbb{E}\left[\int_0^\infty \phi M e^{-(\phi+r)t} e^{\left(-\int_0^t ar\beta_v\sigma dB_v - \frac{1}{2}\int_0^t a^2 r^2\beta_v^2\sigma^2 dv\right)}dt\right] = \frac{\phi M}{\phi+r} = M_p,$$

where the equality is obtained only if $\beta_t = M$ for all t. Thus, at any time the feasible state variable p is bounded. Similarly, we can show that $p \geq -\frac{\phi M}{\phi+r} = -M_p$. Moreover, this result implies that whenever $p_t = \pm M_p$, we must have that for $\forall s \geq t$,

$$\beta_s = \pm M, p_s = \pm M_p, \text{and } \mu_s = \beta_s - p_s = \pm\frac{rM}{\phi+r}.$$

Therefore, once p_t hits $\pm M_p$, the state p_s will stay there from then on. In this sense $\pm M_p$ are absorbing boundaries. ∎

Now we can write problem (33) in the standard dynamic programming language. Substituting $\mu_t = \beta_t - p_t$ in the principal's objective, we have

$$V(p) = \max_{\left\{\beta_t, \sigma_t^P\right\}} \mathbb{E}\left\{\int_0^\infty e^{-rt}\left[(\beta_t - p_t) - \frac{1}{2}(\beta_t - p_t)^2 - \frac{1}{2}ar\sigma^2\beta_t^2\right]dt\right\} \qquad (A7)$$

$$s.t. \ dp_t = \left[(\phi + r)p_t + \beta_t\left(ar\sigma\sigma_t^P - \phi\right)\right]dt + \sigma_t^P dB_t \text{ for all } t > 0, \text{ and } p_0 = p$$

$$\beta_s \in [-M, M], \ p_t \in \left[-M_p, +M_p\right], \text{ and } p_t = \pm M_p \text{ are absorbing.}$$

We are after the optimal policy $\left\{\beta_t^*, \sigma_t^{P,*}\right\}$ as functions of the state variable p_t. And Lemma 3 implies that we know the value function at the boundaries $\pm M_p$:

$$V\left(\pm M_p\right) = \int_0^\infty e^{-rt}\left(\pm\frac{rM}{\phi + r} - \frac{1}{2}\left(\frac{rM}{\phi + r}\right)^2 - \frac{1}{2}a^2 r^2\sigma^2 M^2\right)ds$$

$$= \pm\frac{M}{\phi + r} - \frac{r}{2}\left[\frac{1}{(\phi + r)^2} + a^2\sigma^2\right]M^2. \qquad (A8)$$

For ease of argument, we first consider the principal's relaxed maximization problem given M (and $M_p = \frac{\phi M}{\phi + r}$):

$$V(p; M) \equiv \max_{\left\{\beta_t, \sigma_t^P\right\}} \mathbb{E}\left[\int_0^\infty e^{-rt}\left(\mu_t - \frac{1}{2}\mu_t^2 - \frac{1}{2}ar\sigma^2\beta_t^2\right)dt\right] \qquad (A9)$$

$$s.t. \ dp_t = \left[(\phi + r)p_t + \beta_t\left(ar\sigma\sigma_t^P - \phi\right)\right]dt + \sigma_t^P dB_t \text{ for all } t > 0, \text{ and } p_0 = p$$

$$p_t \text{ is absorbing at} \pm M_p,$$

$$\beta_t \text{ can exceed } M \text{ (but remains finite) when } p_t \in \left(-M_p, M_p\right). \qquad (A10)$$

Problem (A9) is a relaxed version of the principal's original problem (33), due to Equation (A10). Essentially, given M, in the original problem (33) we require $\beta_t \in [-M, M]$ for any time t; while in problem (A9) we only require $\beta_t \in [-M, M]$ whenever p_t hits $\pm M_p$ in light of Lemma 3. Thus, when $p = \pm M_p$, the boundary conditions are the same between these two problems. However, this relaxation helps because the relaxed problem (A9) allows us to use the interior first-order condition of β when $p \in \left(-M_p, M_p\right)$. We will show that for sufficiently large M the value achieved in the relaxed problem is the same as that in the original problem, which implies that the solution to the relaxed problem is also that to the original problem.

Let us define two functions, which prove to be useful in saving notations:

$$H_1(p) \equiv \frac{1}{\phi}p - \frac{1}{2}\frac{r}{\phi^2}p^2 - \frac{a\sigma^2(\phi + r)^2}{2\phi^2}p^2,$$

$$H_2(p) \equiv \frac{1}{r}\left[\frac{1}{2}\frac{(1 + p)^2}{1 + ar\sigma^2} - p - \frac{1}{2}p^2\right].$$

Recall the constant $\overline{p}^d$ in Proposition 2 is given in (43).

We state assumptions required for the proof next.

Assumption 2: The parameters satisfy the following condition, which is equivalent to (42)

$$a\left[2\left(\frac{\phi}{r} + 1\right)^3 - \phi\sigma^2\right] < \frac{\phi}{r}. \qquad (A11)$$

Assumption 3: The feasible policy space for incentives is bounded given state p, that is, $\beta(p) < \infty$ (though $\beta(p)$ may exceed M given M).

To solve the relaxed problem, we first focus on the following key ODE, with the boundary condition (A13) given in Equation (A8):

$$rV(p) = \frac{1}{2} \frac{(1+p-\phi V_p(p))^2}{1+ar\sigma^2 + a^2r^2\sigma^2 \frac{(V_p(p))^2}{V_{pp}}} - p - \frac{1}{2}p^2 + V_p(p)(\phi+r)p, \tag{A12}$$

$$s.t. V(\pm M_p) = H_1(\pm M_p). \tag{A13}$$

We proceed with the following four steps:

Step 1: We first focus on the ODE in Equation (A12). We show that at $p=0$, we have $V(0)=0$ and $V_p(0)=\frac{1}{\phi}$.

Step 2: Under assumption 2, this ODE in Equation (A12) satisfies concavity and positivity of the denominator condition.

- (a), Prove the concavity and positivity of the denominator at $p=0$.
- (b), Prove the concavity and positivity of the denominator for general $p>0$.
- (c), Prove the concavity and positivity of the denominator for $p<0$.
- (d), Prove that the lower boundary 0 is absorbing and the upper boundary $\bar{p}$ is entrance-no-exit (heuristically, at $\bar{p}$ the volatility of $\sigma^p(\bar{p})=0$ but its drift is strictly negative).

Step 3: From Steps 2.a-2.d, it follows from a standard verification theorem that the solution to the ODE in Equation (A12) is the value function of the principal's relaxed problem. Furthermore, it is never optimal to run outside the region $[0, \bar{p}]$ for the relaxed problem.

Step 4: Show that the value function for the relaxed problem is also the value function for the principal's original problem and $\bar{p}$ is independent of M for sufficiently large M.

A.7.2 Step 1: $V(0)=0$ **and** $V_p(0)=\frac{1}{\phi}$. We show that $V(0)=0$ and $V_p(0)=\frac{1}{\phi}$ in three steps:

Step 1.a: We first show that there must exist a solution to the following equation:

$$T(p) \equiv 1 + p - \phi V_p = 0.$$

We know that $V(0) \geq 0$, which is the value of deterministic policy. We also know that

$$V(-M_p) = H_1(-M_p) < 0 \text{ and } V(M_p) = H_1(M_p) < 0.$$

Then according to the intermediate value theorem, there exists $p_1 > 0$ so that

$$T(p_1) = 1 + p_1 - \phi V_p(p_1) = 1 + p_1 - \phi \frac{V(M_p) - V(0)}{M_p} > 1,$$

and there also exists $p_2 < 0$ such that

$$T(p_2) = 1 + p_2 - \phi V_p(p_2) = 1 + p_2 - \phi \frac{V(0) - V(-M_p)}{M_p} < 0,$$

for sufficiently high M_p. Therefore, we can find a point $\underline{p}$ such that

$$1 + \underline{p} - \phi V_p(\underline{p}) = 0. \tag{A14}$$

Step 1.b: Suppose that $1 + ar\sigma^2 + a^2r^2\sigma^2 \frac{V_p^2(\underline{p})}{V_{pp}(\underline{p})} \neq 0$. We aim to show that $\underline{p}=0$ so that equations (A14) and (A12) imply $V(0)=0$ and $V_p(0)=\frac{1}{\phi}$. Differentiating the HJB in Equation (A12) with

respect to p at $\underline{p}$, we have

$$r V_p \left(\underline{p} \right) = -1 - \underline{p} + V_{pp} \left(\underline{p} \right) (\phi + r) \underline{p} + V_p \left(\underline{p} \right) (\phi + r),$$

which, together with Equation (A14), imply that

$$V_{pp} \left(\underline{p} \right) \underline{p} = 0.$$

Therefore, either $\underline{p} = 0$ or $V_{pp} \left(\underline{p} \right) = 0$. We first rule out the case of $V_{pp} \left(\underline{p} \right) = 0$. Note that this case implies the denominator of the first term in the right-hand side of Equation (A12) is infinite, so that V must satisfy

$$r V \left(\underline{p} \right) = -\underline{p} - \frac{1}{2} \underline{p}^2 + \left(1 + \underline{p} \right) \left(1 + \frac{r}{\phi} \right) \underline{p}, \tag{A15}$$

Further, it follows from Taylor expansion that

$$V \left(\underline{p} + \epsilon \right) = V \left(\underline{p} \right) + \frac{1}{\phi} \left(1 + \underline{p} \right) \epsilon + o \left(\epsilon^2 \right) \tag{A16}$$

$$V_p \left(\underline{p} + \epsilon \right) = \frac{1}{\phi} \left(1 + \underline{p} \right) + o(\epsilon).$$

Thus, evaluating the HJB Equation (A12) at $\underline{p} + \epsilon$, we have

$$r V \left(\underline{p} + \epsilon \right) = \frac{\frac{1}{2} (\epsilon - o(\epsilon))^2}{1 + a r \sigma^2 + a^2 r^2 \sigma^2 \frac{V_p^2 \left(\underline{p} + \epsilon \right)}{V_{pp} \left(\underline{p} + \epsilon \right)}} - \underline{p} - \epsilon - \frac{\left(\underline{p} + \epsilon \right)^2}{2} + \frac{1}{\phi} \left(1 + \underline{p} \right) (\phi + r) \left(\underline{p} + \epsilon \right) + o \left(\epsilon^2 \right)$$

$$= -\underline{p} - \frac{1}{2} \underline{p}^2 + \left(1 + \underline{p} \right) \left(1 + \frac{r}{\phi} \right) \underline{p} - \epsilon - \underline{p}\epsilon - \frac{1}{2} \epsilon^2 + \left(1 + \underline{p} \right) \left(1 + \frac{r}{\phi} \right) \epsilon + o \left(\epsilon^2 \right)$$

$$= r V \left(\underline{p} \right) + \frac{r}{\phi} \left(1 + \underline{p} \right) \epsilon - \frac{1}{2} \epsilon^2 + o \left(\epsilon^2 \right),$$

where the second equality uses the fact that $\dfrac{V_p^2 \left(\underline{p} + \epsilon \right)}{V_{pp} \left(\underline{p} + \epsilon \right)}$ goes to infinity as ϵ goes to zero because the continuity of $V_{pp} \left(\underline{p} \right)$ implies that $V_{pp} \left(\underline{p} + \epsilon \right)$ is at the order of $o(1)$, and the last equality follows from Equation (A15). But this contradicts Equation (A16), since they do not match at the second order ϵ^2. As a result, $\underline{p} = 0$ and thus $V_p(0) = \frac{1}{\phi}$.

Step 1.c: Now suppose that $1 + a r \sigma^2 + a^2 r^2 \sigma^2 \dfrac{V_p^2 \left(\underline{p} \right)}{V_{pp} \left(\underline{p} \right)} = 0$, that is,

$$V_{pp} \left(\underline{p} \right) = -\frac{a^2 r^2 \sigma^2 V_p^2 \left(\underline{p} \right)}{1 + a r \sigma^2} = -\frac{a^2 r^2 \sigma^2}{1 + a r \sigma^2} \frac{1}{\phi^2} \left(1 + \underline{p} \right)^2, \tag{A17}$$

which implies that we cannot ignore the term with $1 + p - \phi V_p$. Due to L'Hospital's rule,

$$\frac{\left(1 + p - \phi V_p \right)^2}{1 + a r \sigma^2 + a^2 r^2 \sigma^2 \frac{V_p^2}{V_{pp}}} = \frac{2 \left(1 + p - \phi V_p \right) \left(1 - \phi V_{pp} \right)}{a^2 r^2 \sigma^2 \frac{2 V_p V_{pp}^2 - V_p^2 V_{ppp}}{V_{pp}^2}}. \tag{A18}$$

Differentiating the HJB Equation (A12), we have,

$$rV_p = \frac{\left(1+\underline{p}-\phi V_p\right)\left(1-\phi V_{pp}\right)}{1+ar\sigma^2+a^2r^2\sigma^2\frac{V_p^2}{V_{pp}}} - \frac{1}{2}\frac{\left(1+\underline{p}-\phi V_p\right)^2}{\left(1+ar\sigma^2+a^2r^2\sigma^2\frac{V_p^2}{V_{pp}}\right)^2}\left[a^2r^2\sigma^2\frac{2V_pV_{pp}^2-V_p^2V_{ppp}}{V_{pp}^2}\right]$$

$$-1-\underline{p}+V_p(\phi+r)+V_{pp}(\phi+r)\underline{p}.$$

Plugging Equation (A18) into the above equation, we find that the two terms in the first line cancel each other, and

$$0=-1-\underline{p}+V_p\phi+V_{pp}(\phi+r)\underline{p}=V_{pp}(\phi+r)\underline{p},$$

which is the same as before. Therefore, either we have

$$\underline{p}=0, \text{ and } V_p(0)=\frac{1}{\phi},$$

or we have $V_{pp}\left(\underline{p}\right)=0$ and $\underline{p}=-1$, due to Equation (A17). Furthermore, in the second case, we have

$$V_{pp}(-1)=0, V_p(-1)=0, \text{ and } \frac{V_p^2(-1)}{V_{pp}(-1)}=-\frac{1+ar\sigma^2}{a^2r^2\sigma^2}$$

In Lemma A.4 in the Internet Appendix, we will show that for any p, if

$$V_{pp}(p)=V_p(p)=0 \text{ and } \lim\frac{V_p^2(p)}{V_{pp}(p)} \to -q,$$

then $q=0$ or $\frac{1}{ar}$.[33] Therefore, the second alternative results in a contradiction and we must have $\underline{p}=0$, and $V_p(0)=\frac{1}{\phi}$.

We still need to show that $V(0)=0$ under the assumption $1+ar\sigma^2+a^2r^2\sigma^2\frac{V_p^2\left(\underline{p}\right)}{V_{pp}\left(\underline{p}\right)}=0$. Suppose that $V(0)=v>0$. First, the HJB in Equation (A12) implies that

$$\lim_{p\to 0}\frac{\left(1+p-\phi V_p\right)^2}{1+ar\sigma^2+a^2r^2\sigma^2\frac{V_p^2}{V_{pp}}}=2rv>0.$$

Thus, it follows from Equation (38) that the policy function $\beta(p)$ at $p=0$ has the value

$$\beta(0)=\lim_{p\to 0}\frac{\left(1+p-\phi V_p\right)^2}{1+ar\sigma^2+a^2r^2\sigma^2\frac{V_p^2}{V_{pp}}}\cdot\frac{1}{1+p-\phi V_p}$$

$$=\lim_{p\to 0}2rv\frac{1}{1+p-\phi V_p}=\infty.$$

This contradicts with Assumption 3 that our policy space is restricted to be bounded at $p=0$.

A.7.3 Step 2: Prove the concavity and positivity.
We delegate the proof of this step to the companion Internet Appendix.

---

[33] It is important to point out that the proof for Lemma A4 in the companion Internet Appendix does not use any results from step 1 here. Thus, there is no circular argument.

A.7.4 Step 3: Verification. Define the auxiliary gain process, as a function of the contract Π, as

$$G_t(\Pi) = \int_0^t e^{-rs}\left((\beta_s - p_s) - \frac{1}{2}(\beta_s - p_s)^2 - \frac{1}{2}a^2 r^2 \sigma^2 \beta_s^2\right)ds + e^{-rt}V(p).$$

Define τ as the hitting time when p reaches $\pm M_p$, which could be infinite. Obviously, $G_\tau(\Pi)$ is the actual payoff from the contract Π. For given t, it is easy to show that

$$\mathbb{E}_t\left[e^{rt}dG_t\right] = \left[\begin{array}{c} -rV(p) + (\beta_t - p_t) - \frac{1}{2}(\beta_t - p_t)^2 - \frac{ar\sigma^2}{2}\beta_t^2 \\ + V_p\left[(\phi+r)p_t + \beta_t\left(ar\sigma\sigma^P - \phi\right)\right] + \frac{1}{2}V_{pp}\left(\sigma_t^P\right)^2 \end{array}\right]dt + V_p \sigma_t^P dB_t.$$

Therefore,

$$dG_t = \mu_G(p)dt + e^{-rt}V_p\sigma_t^P dB_t.$$

Due to construction of the ODE in the HJB equation, under the optimal policy Π^* we have $\mu_G(p) = 0$, whereas for other policies we have $\mu_G(p) \leq 0$. Also, since V_p is bounded, and we restrict the policy $\{\sigma_t^P\}$ to be well-behaved (square integrable in the usual sense), $\int_0^t e^{-rt}V_p \sigma_s^P dB_s$ is a martingale. Therefore, under the optimal contract $\mathbb{E}[G_\tau(\Pi^*)] = G_0(\Pi^*) = V(p_0)$.

Given any $T > 0$, we have

$$\mathbb{E}[G_\tau(\Pi)] = \mathbb{E}\left[G_{T\wedge\tau}(\Pi) + \mathbf{1}_{T\leq\tau}\left(\int_T^\tau e^{-rs}\left((\beta_s - p_s) - \frac{1}{2}(\beta_s - p_s)^2 - \frac{1}{2}a^2 r^2 \sigma^2 \beta_s^2\right)ds\right.\right.$$

$$\left.\left. + e^{-r\tau}V(p_\tau)\right)\right]$$

$$\leq G_0 + e^{-rT}\mathbb{E}\left[\int_T^\tau e^{-r(s-T)}\frac{1}{2}ds\right]$$

where $\mathbb{E}\left[\int_T^\tau e^{-r(s-t)}\frac{1}{2}ds\right]$ is the first-best project value. Therefore, let $T \to \infty$, then we have $\mathbb{E}[G_\tau(\Pi)] \leq G_0 = V(p)$. This implies that the proposed contract solves the relaxed problem.

A.7.5 Step 4: $\bar{p}$ is independent of M. We now show that $\bar{p}$ is independent of M. Take some sufficiently large M_1, and consider the solution obtained with the upper entry-no-exit boundary $\bar{p}_1$. Note that $\bar{p}_1 < M_1$ strictly, because $V(\bar{p}_1; M_1) > 0$ while at $M_{p,1} = \left(1 + \frac{r}{\phi}\right)M_1$ the value is strictly negative. And, for $p \in [0, \bar{p}_1]$, we have

$$V(p; M_1) = \mathbb{E}\left[\int_0^\infty e^{-rt}\left((\beta_t - p_t) - \frac{(\beta_t - p_t)^2}{2} - \frac{1}{2}a^2 r^2 \sigma^2 \beta_t^2\right)dt \,\Big|\, p_0 = p\right].$$

It is clear that given $\bar{p}_1$, this function is independent of M_1, because under the optimal policy $p \in [0, \bar{p}_1]$, and M does not affect the flow payoff per se (note that β_t is assumed to be unconstrained in the relaxed problem).

Now consider $M_2 \in (\bar{p}_1, M_1)$. The next lemma follows.

Lemma 4. $V(p; M_1) \geq V(p; M_2)$ for $p \in [0, \bar{p}_1]$.

Proof. $M_{p,i}$ denotes the corresponding M_p's. Since $M_1 > M_2$ and $M_{p,1} > M_{p,2}$, the policy space for the semi-constrained problem with M_1 is strictly larger than the policy space in the problem with M_2. To see this, note that for the problem 1 (with M_1), the principal can choose $\beta_s = \left(1 + \frac{r}{\phi}p\right)$ for $s > t$ once $p_t \in [M_{p,2}, M_{p,1}]$, which is exactly the constraint for the policy space of problem 2 with M_2. As a result, $V(p; M_2) \leq V(p; M_1)$ for $p \in [0, \bar{p}_1]$. ∎

However, given M_2, consider the exact same policy under M_1 with endogenous upper entry-no-exit boundary $\bar{p}_1$, which generates the same value as $V(p; M_1)$. As a result, the policy under M_1 also solves the problem with M_2. Therefore we must have the same solution for both M_i's, and $\bar{p}_1 = \bar{p}_2$.

A.7.6 Step 5: Relaxed problem solves the original problem. We now show that the relaxed problem (A9) solves the original problem (A7). As explained before, our original problem in Equation (A7) has more stringent constraints than the relaxed problem (A9): In the original problem we require $\beta_t \leq M$ always, while for the relaxed problem we only require that $\beta_t \leq M$ whenever p_t hits $\pm M_p$. As a result, we have $V(p) \geq V^C(p)$ always. Here, $V^C(p)$ denotes the value for the principal's original problem.

To prove our theorem, it suffices to show that in the region $[0, \bar{p}]$, we have $V(p) = V^C(p)$, that is, the relaxed problem and the original problem achieve the same value for sufficiently high M. Take the solution $V(p)$ and its corresponding incentive policy $\beta^M(\cdot)$; and define

$$\Xi(M) \equiv \max_{0 \leq p \leq \bar{p}} \left| \beta^M(p) \right|,$$

where $\beta^M(p)$ emphasizes the possibility of the dependence of the optimal policy on the parameter value M. If we can show that we can choose sufficiently high M so that $\Xi(M) \leq M$ holds, then the additional constraints are never binding in the original problem, and both problems share the same solution obtained in Proposition 2.

We show that $\Xi(M)$ is independent of M for sufficiently high M, which immediately implies our result. Clearly, it is sufficient to show that both the relaxed value function $V(p)$ and $\bar{p}$ are independent of M. We have shown that $\bar{p}$ is independent of M when M is sufficiently high. Moreover, since the endogenous state p never goes outside the region $[0, \bar{p}]$, we have

$$V(p) = \mathbb{E}_0 \left[\int_0^\infty e^{-rt} \left((\beta_t - p_t) - \frac{(\beta_t - p_t)^2}{2} - \frac{1}{2} a^2 r^2 \sigma^2 \beta_t^2 \right) dt \, \middle| \, p_0 = p \right]$$

to be independent of M. As a result, $\max_{0 \leq p \leq \bar{p}} \left| \beta^M(p) \right|$ is independent of M, and our result follows.

A.8 Proof for Proposition 3

Consider any alternative policy $\{\hat{c}, \hat{\mu}\}$ deviating from the original policy $\{c, \mu\}$, with an expected payoff $\mathbb{E}_0^{\hat{\mu}} \left[\int_0^\infty e^{-rs} u(\hat{c}_s, \hat{\mu}_s) ds \right]$. To prove that this deviation payoff cannot exceed the equilibrium payoff v, we are follow Sannikov (2014) and construct an upper bound for the deviation policies. We take the optimal contract as given, which gives $\{\beta_t^*, \sigma_t^{P*}\}$. In this proof, we omit "*" on $\{\beta_t^*, \sigma_t^{P*}\}$ without any risk of confusion.

To construct the upper bound for deviation payoffs, we will show that it is sufficient to keep track of two deviation state variables that matter to the agent's potential deviation value. The first deviating state variable captures the agent's private saving:

$$S_t \equiv \int_0^t e^{r(t-s)} (c_s - \hat{c}_s) ds.$$

The second deviating state variable is the belief wedge as defined in (13) captures the persistent belief manipulation effect:

$$\Delta_t = \phi \int_0^t e^{\phi(s-t)} (\mu_s - \hat{\mu}_s) ds.$$

Given these two deviating variables, we propose a candidate of an upper bound for the agent's deviation value, which is defined as

$$W(v_t, p_t; S_t, \Delta_t) \equiv \underbrace{v_t}_{\text{equilibrium contract}} \cdot \underbrace{\exp(-ar S_t)}_{\text{deviation value from savings}} \cdot \underbrace{\exp\left(-ar\left(\frac{1}{\phi}\Delta_t p_t + 0.5k\Delta_t^2\right)\right)}_{\text{deviation value from belief distortions}}.$$

$$\text{(A19)}$$

The linear coefficient p_t in front of Δ_t reflects the first-order gain of the information rent p_t from belief-manipulation; the quadratic coefficient k will be chosen shortly to ensure $W(v_t, S_t, \Delta_t)$ being the upper bound of the agent's deviation value.

We further require the following assumptions on the agent's deviation strategies for the usual transervasality conditions, which are standard in infinite-horizon consumption/saving problems. More specifically, there exist (however large) positive constants L_s and L_Δ, so that the agent's deviation strategies satisfy

$$|S_t| < L_s, \text{ and } |\Delta_t| < L_\Delta.$$

We have the following key lemma showing that $W(v_t, p_t; S_t, \Delta_t)$ is the upper bound of the agent's deviation value.

Lemma 5. Facing the contract $\{c_t, \mu_t\}$, suppose that the agent's deviation history leads to a pair of deviation states as (S_t, Δ_t) at time t. Then the agent's deviation value from time t onwards is bounded above by $W(v_t, p_t; S_t, \Delta_t)$, if either (A22) or (A23) holds.

Proof. We first give the outline of the argument. To prove $W(v_t, S_t, \Delta_t)$ is an upper bound for the agent's deviation value, define the auxiliary gain process G_t associated with any feasible policies $\{\widehat{c}, \widehat{\mu}\}$ as

$$G_t\left(\{\widehat{c}_s, \widehat{\mu}_s\}_{s=0}^\infty\right) \equiv \int_0^t e^{-rs} u(\widehat{c}_s, \widehat{\mu}_s) ds + e^{-rt} W(v_t, p_t; S_t, \Delta_t).$$

Clearly, $\mathbb{E}_0^{\widehat{\mu}}[G_\infty] = \mathbb{E}_0^{\widehat{\mu}}\left[\int_0^\infty e^{-rs} u(\widehat{c}_s, \widehat{\mu}_s) ds\right]$ is the expected payoff under the feasible policy, given the transversality condition $\lim_{s \to \infty} \mathbb{E}_0^{\widehat{\mu}}\left[e^{-rt} W_t\right] = 0$ (which is implied by Assumption 1 for transversality conditions). On the other hand, $G_0 = W(v_0, S_0, \Delta_0)$ is the proposed upper bound of the agent's deviation value given the current relevant deviation states (S_0, Δ_0). Obviously, one sufficient condition for $\mathbb{E}_0^{\widehat{\mu}}[G_\infty] \leq G_0 = W(v_0, S_0, \Delta_0)$, that is, the upper bound (A19) is valid, is that the auxiliary gain process G_t is a supermartingale for any deviation policy under the agent's information set.

Now we start the proof. For ease of notation,

$$\delta_t = \widehat{\mu}_t - \mu_t, \widehat{u}_t \equiv u(\widehat{c}_t, \widehat{\mu}_t), u_t \equiv u(c_t, \mu_t), \text{ with } u_t = r v_t.$$

Under the measure $\widehat{\mu}$, we have evolutions

$$dp_t = \left[(\phi + r)p_t + \beta_t\left(ar\sigma\sigma_t^P - \phi\right) + \frac{\sigma_t^P}{\sigma}(\delta_t + \Delta_t)\right]dt + \sigma_t^P dB_t^{\widehat{\mu}},$$

$$dv_t = -v_t ar\beta_t(\sigma dB_t^{\widehat{\mu}} + [\widehat{\delta}_t + \Delta_t]dt),$$

$$dS_t = (rS_t + c_t - \widehat{c}_t)dt, \text{ and } d\Delta_t = -\phi(\delta_t + \Delta_t)dt.$$

It follows that

$$e^{rt}dG_t = [\widehat{u}_t - rW_t]dt + W_t dv_t - arW_t\left[dS_t + d\left(\frac{1}{\phi}\Delta_t p_t + 0.5k\Delta_t^2\right)\right] + \left(-\frac{ar}{\phi}\Delta_t\right)\frac{W_t}{v_t}\langle dv_t, dp_t\rangle$$

$$= [\widehat{u}_t - rW_t]dt + W_t dv_t - arW_t\left[dS_t + d\left(\frac{1}{\phi}\Delta_t p_t + 0.5k\Delta_t^2\right) - \Delta_t\frac{ar}{\phi}\beta_t\sigma\sigma_t^P dt\right]$$

$$= -ar\beta_t W_t \sigma dB_t^{\widehat{\mu}} + [\widehat{u}_t - rW_t]dt - arW_t \frac{1}{\phi}\Delta_t dp_t$$

$$- arW_t\left(\beta_t[\delta_t + \Delta_t]dt + dS_t + \left(\frac{1}{\phi}p_t + k\Delta_t\right)d\Delta_t - \Delta_t \frac{ar}{\phi}\beta_t\sigma\sigma_t^P dt\right)$$

$$= [-ar\beta_t W_t\sigma - arW_t\frac{1}{\phi}\Delta_t\sigma_t^P]dB_t^{\widehat{\mu}} - rW_t\frac{a}{\phi}\Delta_t\left[(\phi+r)p_t + \beta_t\left(ar\sigma\sigma_t^P - \phi\right) + \frac{\sigma_t^P}{\sigma}(\delta_t + \Delta_t)\right]dt$$

$$+ rW_t\left\{e^{-a\left(\widehat{c}_t - \frac{1}{2}\widehat{\mu}_t^2\right) + a\left(c_t - \frac{1}{2}\mu_t^2\right) + ar\left(S_t + \frac{1}{\phi}p_t\Delta_t + 0.5k\Delta_t^2\right)} - 1\right\}dt$$

$$+ rW_t\left\{-a\beta_t[\delta_t + \Delta_t] - a(rS_t + c_t - \widehat{c}_t) + a(p_t + \phi k\Delta_t)(\delta_t + \Delta_t) + \Delta_t\frac{a^2r}{\phi}\beta_t\sigma\sigma_t^P\right\}dt.$$

Because $e^x \geq 1 + x$ and $W_t < 0$ (since $v_t < 0$), the drift of $e^{rt}G_t$ is bounded above by

$$rW_t\left\{\begin{array}{l} -\frac{a}{\phi}\Delta_t\left[(\phi+r)p_t + \beta_t\left(ar\sigma\sigma_t^P - \phi\right) + \frac{\sigma_t^P}{\sigma}(\delta_t + \Delta_t)\right] - a\left(\widehat{c}_t - \frac{1}{2}\widehat{\mu}_t^2\right) + a\left(c_t - \frac{1}{2}\mu_t^2\right) \\ + ar\left(S_t + \frac{1}{\phi}p_t\Delta_t + 0.5k\Delta_t^2\right) - a(rS_t + c_t - \widehat{c}_t) - a(\beta_t - p_t - \phi k\Delta_t)(\delta_t + \Delta_t) + \Delta_t\frac{a^2r}{\phi}\beta_t\sigma\sigma_t^P \end{array}\right\}$$

$$= arW_t\left[\frac{1}{2}\delta_t^2 + (\phi k - \frac{\sigma_t^P}{\phi\sigma})\delta_t\Delta_t + (0.5rk + \phi k - \frac{\sigma_t^P}{\phi\sigma})\Delta_t^2\right]$$

$$= arW_t\left[\frac{1}{2}\left(\delta_t + \left(\phi k - \frac{\sigma_t^P}{\phi\sigma}\right)\Delta_t\right)^2 + \left(0.5rk + \phi k - \frac{\sigma_t^P}{\phi\sigma} - \frac{1}{2}\left(\phi k - \frac{\sigma_t^P}{\phi\sigma}\right)^2\right)\Delta_t^2\right]. \tag{A20}$$

Notice that all the terms that are linear in the instantaneous deviation δ_t and cumulative past deviations Δ_t all get cancelled, thanks to the first-order *Incentive Compatibility* condition for the agent.

Our goal is to show that the sum of all the quadratic terms in (A20) is negative always. Then, because the drift of $e^{rt}G_t$ is bounded above by (A20), it immediately implies the drift for dG_t is always negative. For all possible instantaneous deviations δ_t and cumulative past deviations Δ_t, because $W_t < 0$, the necessary and sufficient condition for (A20) to be negative always is that the term in the bracket is positive for the all equilibrium values σ_t^P in the optimal contract. In other words, we require that for we choose some k so that for in the optimal contract, σ_t^P satisfies

$$\frac{1}{2}\left(\phi k - \frac{\sigma_t^P}{\phi\sigma}\right)^2 - \left(0.5rk + \phi k - \frac{\sigma_t^P}{\phi\sigma}\right) \leq 0$$

$$\Leftrightarrow \phi^2 k^2 - (r + 2\phi + 2\frac{\sigma_t^P}{\sigma})k + (\frac{\sigma_t^P}{\sigma\phi} + 1)^2 - 1 \leq 0 \tag{A21}$$

Obviously, this imposes certain sufficient condition on the range of $\{\sigma_t^P\}$ in the optimal contract. We give two particular examples of the sufficient conditions.

1. Proposition 2 shows that $\{\sigma_t^P\}$ is bounded in the optimal contract. What is more, although we cannot prove it rigorously, the optimal contract exhibits the "option-like" feature so that $\sigma_t^P > 0$ which holds in all of our numerical examples. The range of σ_t^P is represented by $[0, L_{\sigma p}]$ with $L_{\sigma p} \equiv \sup \sigma_t^P > 0$. Then the two real roots for the left-hand side of the

quadratic Equation (A21), represented by $k_-\left(\sigma_t^P\right) < k_+\left(\sigma_t^P\right)$, are

$$k_+\left(\sigma_t^P\right) = \frac{r+2\phi}{2\phi^2} + \frac{\sigma_t^P}{\sigma\phi^2} + \frac{\sqrt{4\phi^2+r^2+4r(\phi+\frac{\sigma_t^P}{\sigma})}}{2\phi^2},$$

$$k_-\left(\sigma_t^P\right) = \frac{r+2\phi}{2\phi^2} + \frac{\sigma_t^P}{\sigma\phi^2} - \frac{\sqrt{4\phi^2+r^2+4r(\phi+\frac{\sigma_t^P}{\sigma})}}{2\phi^2}.$$

Since $\sigma_t^P > 0$, both roots are increasing in σ_t^P because $\frac{\partial k_+}{\partial \sigma_t^P} = \frac{1}{\sigma\phi^2}(1 + \frac{r}{\sqrt{(2\phi+r)^2+4r\sigma_t^P/\sigma}}) > 0$,

and $\frac{\partial k_-}{\partial \sigma_t^P} = \frac{1}{\sigma\phi^2}(1 - \frac{r}{\sqrt{(2\phi+r)^2+4r\sigma_t^P/\sigma}}) > 0$. As a result, if $k_-\left(\sigma_t^P = L_{\sigma_p}\right) < k_+\left(\sigma_t^P = 0\right)$, then there exists a positive k lying inside the interval

$$k \in \left[k_-\left(\sigma_t^P = L_{\sigma_p}\right), k_+\left(\sigma_t^P = L_{\sigma_p}\right)\right],$$

so that the drift for dG_t is always negative for all $\sigma_t^P \in [0, L_{\sigma_p}]$. To summarize, we need the following sufficient condition

$$\frac{L_{\sigma_p}}{\sigma\phi^2} - \frac{\sqrt{4\phi^2+r^2+4r(\phi+\frac{L_{\sigma_p}}{\sigma})}}{2\phi^2} \leq \frac{r+2\phi}{2\phi^2}, \text{ where } L_{\sigma_p} \equiv \sup\sigma_t^P \text{ and } \sigma_t^P > 0. \quad \text{(A22)}$$

Because the left-hand side is increasing in $L_{\sigma_p} > 0$, this condition requires the volatility of information rent σ_t^P cannot be too high. A slightly more relaxed sufficient condition for (A22) is to set the L_{σ_p} inside the square root to be zero, and the condition simplifies to $L_{\sigma_p} \leq \sigma(r+2\phi)$.

2. If we do not impose that $\sigma_t^P > 0$, one can still follow the above logic to give a similar but more complicated sufficient condition for the range of $\{\sigma_t^P\}$. But there is a natural choice of $k = 1$ that makes the condition (A21) transparent without $\sigma_t^P > 0$. In this case, (A21) can be simplified to

$$(\sigma_t^P)^2 \leq \sigma^2\phi^2\left(r+2\phi-\phi^2\right), \ k=1 \quad \text{(A23)}$$

which gives an upper bound on the absolute magnitude of σ_t^P. Note this condition does not require the preassumption that $\sigma_t^P > 0$.

To summarize, we have shown that under the measure induced by $\widehat{\mu}$,

$$dG_t = \text{negative drift} + [-ar\beta_t W_t \sigma - ar W_t \frac{1}{\phi}\Delta_t\sigma_t^P]dB_t^{\widehat{\mu}},$$

if the volatility of information rent, σ_t^P, is not excessively high (in the sense of either (A22) or (A23), which easily holds in our numerical example). Intuitively, all else equal, the agent's global deviation value tends to be increasing in the volatility σ_t^P of his deviation state variable, because the agent's has the "option" to adjust his optimal strategy swiftly following a sequence of deviations and performance shocks.

The last routine step to ensure G_t being a supermartingale is to check the following condition:

$$\mathbb{E}_0^{\widehat{\mu}}\left[\int_0^T (ar\beta_t W_t \sigma + ar W_t \frac{1}{\phi}\Delta_t\sigma_t^P)dB_t^{\widehat{\mu}}\right] = 0 \text{ for all } T.$$

Since $|S_t|$, $|\Delta_t|$, $|p_t|$, $|\beta_t|$ and $|\sigma_t^P|$ are bounded, we only need to ensure the square integrability condition (Revuz and Yor 1999, 139):

$$\mathbb{E}_0^{\widehat{\mu}}\left[\int_0^T \left(e^{-rt}v_t\right)^2 dt\right] < \infty \text{ for all } T.$$

Under $\widehat{\mu}$, using (A3) we have $\frac{dv_t}{v_t} = -ar\beta_t\sigma\,dB_t^{\widehat{\mu}} - ar\beta_t(\delta_t - \Delta_t)dt$, which implies that

$$v_t = v_0\exp\left[\int_0^t -ar\beta_s\sigma\,dB_s^{\widehat{\mu}} - \int_0^t 0.5a^2r^2\beta_s^2\sigma^2\,ds - \int_0^t ar\beta_s[\delta_s - \Delta_s]ds\right].$$

L_μ and M represent the upper bounds of $|\mu_t|$ and $|\beta_t|$, respectively; that is, $|\mu_t| < L_\mu$ and $|\beta_t| < M$. Then we have

$$\left|\int_0^t ar\beta_s[\delta_s - \Delta_s]ds\right| = ar\left|\int_0^t \delta_s\left[\beta_s - \phi\int_s^t e^{\phi(s-u)}\beta_u\,du\right]ds\right|$$

$$< ar\int_0^t |\delta_s|\left|\beta_s - \phi\int_s^t e^{\phi(s-u)}\beta_u\,du\right|ds$$

$$< ar\int_0^t 2L_\mu\max\left(\beta_s, \phi\int_s^t e^{\phi(s-u)}\beta_u\,du\right)ds < 2arL_\mu Mt,$$

$$\left|\int_0^t 0.5a^2r^2\beta_s^2\sigma^2\,ds\right| < 0.5a^2r^2\sigma^2M^2t.$$

Hence,

$$\int_0^T \left(e^{-rt}v_t\right)^2dt = \int_0^T \left(e^{-rt}v_0\right)^2\exp\left[\int_0^t -2ar\beta_s\sigma\,dB_s^{\widehat{\mu}} - \int_0^t a^2r^2\beta_s^2\sigma^2\,ds\right.$$

$$\left. - \int_0^t 2ar\beta_s[\delta_s - \Delta_s]ds\right]dt$$

$$< \int_0^T (v_0)^2 e^{4arL_\mu Mt - 2rt}\exp\left[\int_0^t -2ar\beta_s\sigma\,dB_s^{\widehat{\mu}}\right.$$

$$\left. - \int_0^t 2a^2r^2\beta_s^2\sigma^2\,ds\right]\exp\left[\int_0^t a^2r^2\beta_s^2\sigma^2\,ds\right]dt$$

$$< \int_0^T (v_0)^2 e^{4arL_\mu Mt + a^2r^2M^2\sigma^2t - 2rt}\exp\left[\int_0^t -2ar\beta_s\sigma\,dB_s^{\widehat{\mu}}\right.$$

$$\left. - \int_0^t 2a^2r^2\beta_s^2\sigma^2\,ds\right]dt.$$

Because $\int_0^t (2ar\beta_s\sigma)^2\,ds < (2arM\sigma)^2t$ for all t, $\exp\left[\int_0^t -2ar\beta_s\sigma\,dB_s^{\widehat{\mu}} - \int_0^t 2a^2r^2\beta_s^2\sigma^2\,ds\right]$ is an exponential martingale under the measure induced by $\widehat{\mu}$. Therefore, for all T, we have

$$\mathbb{E}_0^{\widehat{\mu}}\left[\int_0^T \left(e^{-rt}v_t\right)^2dt\right] < \int_0^T e^{4arL_\mu Mt + a^2r^2\sigma^2M^2t - 2rt}(v_0)^2$$

$$\mathbb{E}_0^{\widehat{\mu}}\left[\exp\left(\int_0^t -2ar\beta_s\sigma\,dB_s^{\widehat{\mu}} - \int_0^t 2a^2r^2\beta_s^2\sigma^2\,ds\right)\right]dt$$

$$= (v_0)^2\int_0^T e^{4arL_\mu Mt + a^2r^2\sigma^2M^2t - 2rt}\,dt < \infty.$$

Now, given the fact that G_t is a supermartingale, we have

$$W(S_0, \Delta_0) = G_0 \geq \mathbb{E}_0^{\widehat{\mu}}\lim_{t\to\infty}[G_t] = \mathbb{E}_0^{\widehat{\mu}}\left[\int_0^\infty e^{-rs}u(\widehat{c}_s, \widehat{\mu}_s)ds + \lim_{t\to\infty}\left[e^{-rt}W_t\right]\right]$$

$$= \mathbb{E}_0^{\widehat{\mu}}\left[\int_0^\infty e^{-rs}u(\widehat{c}_s, \widehat{\mu}_s)ds\right],$$

which is the agent's deviation payoff. Here, the last equality requires the transversalilty condition which is ensured by the assumption of bounded $|S_t|$ and $|\Delta_t|$. This implies that (A19) is indeed the upper bound for the agent's deviation value. ∎

We have shown that $W(v_t, p_t; S_t, \Delta_t)$ is an upper bound for the agent's potential deviation value given the deviated states (S_t, Δ_t). Then, for an agent who has not deviated yet with $S_t = \Delta_t = 0$, the upper bound of his deviation value is just v_t. Because the equilibrium strategy achieves this upper bound v_t, the equilibrium strategy is indeed globally optimal. As a result, we have shown that the equilibrium strategy that achieves v_t is indeed optimal. Q.E.D.

A.9 What If the Agent Cannot Privately Save?
This appendix analyzes the case in which the agent cannot private save. Recall the following definitions

$$v_t \equiv \mathbb{E}_t\left[\int_t^\infty e^{-r(s-t)} u(c_s, \mu_s) ds\right], \widetilde{p}_t \equiv \mathbb{E}_t\left[\int_t^\infty \phi e^{-(r+\phi)(s-t)} \widetilde{\beta}_s ds\right], \widetilde{\beta}_t = (-arv_t)\beta_t,$$

and the associated volatilities $\widetilde{\beta}_t$ and $\widetilde{\gamma}_t$ so that

$$dv_t = (rv_t - u_t) dt + \widetilde{\beta}_t \sigma d B_t,$$

$$d\widetilde{p}_t = \left((r+\phi)\widetilde{p}_t - \phi\widetilde{\beta}_t\right) dt + \widetilde{\gamma}_t \sigma d B_t.$$

$\widetilde{J}(v, \widetilde{p})$ denotes the principal's value function (ignoring the posteior mean of the project m), which satisfies the HJB

$$r\widetilde{J} = \max_{c, \widetilde{\beta}, \widetilde{\gamma}} \mu\left(c, \widetilde{\beta}; \widetilde{p}\right) - c + \widetilde{J}_v\left(rv - u\left(c, \mu\left(c, \widetilde{\beta}; \widetilde{p}\right)\right)\right) + \widetilde{J}_{\widetilde{p}}\left((r+\phi)\widetilde{p} - \phi\widetilde{\beta}\right)$$

$$+ \frac{\sigma^2}{2}\left[\widetilde{J}_{vv}\widetilde{\beta}^2 + \widetilde{J}_{\widetilde{p}\widetilde{p}}\widetilde{\gamma}^2 + 2\widetilde{J}_{v\widetilde{p}}\widetilde{\beta}\widetilde{\gamma}\right], \tag{A24}$$

with the agent's optimal effort $\mu\left(c, \widetilde{\beta}; \widetilde{p}\right)$ satisfying the first-order condition $-u_\mu\left(c_t, \mu\left(c_t, \widetilde{\beta}_t; \widetilde{p}lt\right)\right) = \widetilde{\beta}_t - \widetilde{p}_t$.

Define

$$p_t \equiv \frac{\widetilde{p}_t}{(-arv_t)}; \beta_t \equiv \frac{\widetilde{\beta}_t}{(-arv_t)}; \gamma_t \equiv \frac{\widetilde{\gamma}_t}{(-arv_t)}, \text{ and } \sigma_t^P \equiv \sigma\left(\gamma_t + arp_t\beta_t\right).$$

A derivation similar to that in Appendix A.6 leads to the evolution of p_t in the no private savings case (with $rv_t \neq u_t$) as

$$dp_t = \left[(r+\phi)p_t - \frac{p_t}{v_t}(rv_t - u_t) + \beta_t\left(ar\sigma\sigma_t^P - \phi\right)\right] dt + \sigma_t^P d B_t.$$

Furthermore, the agent's optimal effort policy can be redefined as $\mu(c, \beta; v, p)$. For simplicity, we omit (v, p) in the expression. As explained in the main text, the agent's incentive-compatibility condition is $u_\mu(c_t, \mu_t) = arv_t(\beta_t - p_t)$, and since under CARA we have $u_\mu = a\mu u$, it becomes

$$u(c_t, \mu_t) = \frac{rv_t}{\mu_t}(\beta_t - p_t).$$

This implies that

$$c_t = \frac{1}{2}\mu_t^2 - \frac{1}{a}\ln(-arv_t) - \frac{1}{a}\ln\frac{\beta_t - p_t}{\mu_t}. \tag{A25}$$

The principal's value function can be redefined as $J(v,p) \equiv \tilde{J}\left(v, \frac{\tilde{p}_t}{-arv_t}\right)$; and (A24) implies that

$$rJ = \max_{c,\beta,\sigma^P} \mu(c,\beta) - c + J_v(rv - u(c,\mu(c,\beta)))$$

$$+ J_p\left((r+\phi)p - \frac{p}{v}(rv - u(c,\mu(c,\beta))) + \beta(ar\sigma\sigma^P - \phi)\right)$$

$$+ \frac{1}{2}\left[\sigma^2 J_{vv}(-arv)^2\beta^2 + J_{pp}(\sigma^P)^2 + 2J_{vp}(-arv)\beta\sigma\sigma^P\right].$$

Given the CARA preferences, we conjecture that (*ns* stands for "no savings")

$$J(v,p) = \frac{\ln(-arv)}{ar} + V^{ns}(p),$$

with $J_v = \frac{1}{arv}$, $J_p = V_p^{ns}$, $J_{vv} = -\frac{1}{arv^2}$ and $J_{pp} = V_{pp}^{ns}$. Using the expression of c in (A25), and observing that v cancels, one can simplify the HJB equation to the following ODE for $V(p)$:

$$rV^{ns} = \max_{\mu,\beta,\sigma^P} \mu - \frac{1}{2}\mu^2 + \frac{1}{a}\ln\frac{\beta-p}{\mu} + \frac{1}{a}\left(1 - \frac{\beta-p}{\mu}\right) - \frac{1}{2}ar\sigma^2\beta^2$$

$$+ V_p^{ns}\left(-\phi(\beta-p) + rp\frac{\beta-p}{\mu} + ar\sigma\sigma^P\beta\right) + \frac{1}{2}(\sigma^P)^2 V_{pp}^{ns}, \tag{A26}$$

with the following first-order-conditions:

$$\mu\left(1 + a\mu^2 - a\mu\right) = \left(1 - arpV_p^{ns}\right)(\beta-p), \tag{A27}$$

$$\sigma^P = -ar\sigma \frac{V_p^{ns}}{V_{pp}^{ns}}\beta, \tag{A28}$$

$$\beta = \frac{1}{ar\sigma^2}\frac{(1-\mu)\left(1 - arpV_p^{ns}\right)}{1 + a\mu^2 - a\mu} + V_p^{ns}\cdot\left(-\frac{\phi}{ar\sigma^2} + \frac{\sigma^P}{\sigma}\right). \tag{A29}$$

Remark 3. Comparing to private savings case in which $\mu = \beta - p$, now in (A27) we have μ to solve a cubic equation

$$\mu\left(1 + a\mu^2 - a\mu\right) = \left(1 - arpV_p^{ns}\right)(\beta-p),$$

with its right hand side being dependent on the value function itself via V_p^{ns}. The reason why V_p^{ns} matters is that when changing μ, the principal controls the agent's utility (recall c can be fixed when there are no private savings). This affects the drift of v_t which is $rv_t - u(c_t, \mu_t)$ (recall (30)), and in turn the evolution of the information rent $\frac{\tilde{p}_t}{(-arv_t)}$. In contrast, with hidden savings, the agent's consumption adjusts with implemented μ_t so that $u_t = rv_t$ always.

We also have a much more complicated flow term in (A26), which involves two new highly nonlinear terms $\frac{1}{a}\ln\frac{\beta-p}{\mu} + \frac{1}{a}\left(1 - \frac{\beta-p}{\mu}\right)$. In contrast, these two terms vanish when $\mu = \beta - p$ in the case of private savings.

The final ODE about $V(p)$ becomes extremely complicated. We have tried the following. From Eq. (A28) and (A29), we have

$$\beta = \frac{\frac{1}{ar\sigma^2}\frac{(1-\mu)\left(1-arpV_p^{ns}\right)}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns}}{1 + ar\frac{\left(V_p^{ns}\right)^2}{V_{pp}^{ns}}} \quad \text{and} \quad \sigma^P = -ar\sigma\frac{V_p^{ns}}{V_{pp}}\frac{\frac{1}{ar\sigma^2}\frac{(1-\mu)\left(1-arpV_p^{ns}\right)}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns}}{1 + ar\frac{\left(V_p^{ns}\right)^2}{V_{pp}^{ns}}}.$$

Substituting these two equations into the HJB, we have

$$
rV^{ns} = \max_{\mu} \mu - \frac{1}{2}\mu^2 + \frac{1}{a}\ln \frac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns} - p\left(1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\right)}{\mu\left(1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\right)} \tag{A30}
$$

$$
+ \frac{1}{a}\left(1 - \frac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns} - p\left(1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\right)}{\mu\left(1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\right)}\right)
$$

$$
- \frac{1}{2}ar\sigma^2\left(\frac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns}}{1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}}\right)^2
$$

$$
+ V_p \quad + \quad \left(
\begin{array}{l}
-\phi\left(\dfrac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns}}{1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}} - p\right) \\[3em]
+\; rp\,\dfrac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns} - p\left(1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\right)}{\mu\left(1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\right)} \\[3em]
-\, a^2 r^2\sigma^2 \dfrac{V_p^{ns}}{V_{pp}^{ns}}\left(\dfrac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns}}{1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}}\right)^2
\end{array}
\right)
$$

$$
+ \frac{a^2 r^2\sigma^2}{2}\frac{(V_p^{ns})^2}{V_{pp}^{ns}}\left(\frac{\frac{1}{ar\sigma^2}\frac{(1-\mu)(1-arpV_p^{ns})}{1+a\mu^2-a\mu} - \frac{\phi}{ar\sigma^2}V_p^{ns}}{1+ar\frac{(V_p^{ns})^2}{V_{pp}^{ns}}}\right)^2 .
$$

This is way more complicated than the final ODE (39) in the case with private savings; note that in the above ODE we have not even expressed the optimal μ explicitly yet.

We conclude by stating that it is quite challenging to even numerically solve the case without privates savings. In unreported results that are available upon request, we analyze an important benchmark case with deterministic contracts (i.e., implies $\sigma^p = 0$). Recall that in the private saving case, the resultant value function under deterministic contracts is a quadratic function, and the solution is derived in closed form (see Proposition 4). This solution to this benchmark case helps quite a bit in guessing the structure of value function in the general case ($\underline{p} = 0$ and $\overline{p} = p^d$). Unfortunately, in the case of no private savings, the ODE for deterministic contracts becomes highly nonlinear in p, and no longer tractable like in the case of private savings.

A.10 Proof for Proposition 4

We first conjecture that the value function for the deterministic policy, $V^d(p)$, has the following quadratic form

$$V^d(p) = -\frac{1}{2}A^d p^2 + B^d p + C^d.$$

Plugging the above conjecture into the following ODE for the deterministic value function

$$rV^d(p) = \frac{1}{2}\frac{\left(1+p-\phi V^d(p)\right)^2}{1+ar\sigma^2} - p - \frac{1}{2}p^2 + V_p^d(p)(\phi+r)p,$$

we can easily show that $B^d = \frac{1}{\phi}$, $C^d = 0$, and A^d satisfies

$$-\frac{1}{2}rA^d p^2 = \frac{1}{2}\frac{\left(1+\phi A^d\right)^2 p^2}{1+ar\sigma^2} - \frac{1}{2}p^2 - A^d(\phi+r)p^2.$$

Rearranging the above equation, we have

$$\phi^2\left(A^d\right)^2 - A^d\phi\left[\frac{r}{\phi}\left(1+ar\sigma^2\right)+2ar\sigma^2\right] - ar\sigma^2 = 0,$$

which gives the solution for A^d:

$$A^d = \frac{1}{2\phi}\left[\frac{r}{\phi}\left(1+ar\sigma^2\right)+2ar\sigma^2 + \sqrt{\left[\frac{r}{\phi}\left(1+ar\sigma^2\right)+2ar\sigma^2\right]^2 + 4ar\sigma^2}\right].$$

The optimal initial $p_0^d = \frac{B^d}{A^d}$ follows easily from the first-order equation.

The incentive, as a function of information rent p_t, is

$$\beta_t^d = \frac{1+p_t-V_p^d\phi}{1+ar\sigma^2} = \frac{1+A^d\phi}{1+ar\sigma^2}p_t.$$

Using Equation (35), we can derive the evolution of information rent p_t to be

$$\frac{dp_t^d}{p_t^d} = (\phi+r)dt - \frac{\beta_t^d}{p_t}\phi dt = \left(\phi+r - \frac{1+A^d\phi}{1+ar\sigma^2}\phi\right)dt \equiv -\lambda dt.$$

To show that $\lambda = \frac{1+A^d\phi}{1+ar\sigma^2}\phi - (\phi+r) > 0$, it is equivalent to show that $A^d > \left(1+\frac{r}{\phi}\right)\frac{ar\sigma^2}{\phi} + \frac{r}{\phi^2}$, which always holds by Lemma A.2 in the companion Internet Appendix. Finally, the optimal effort can be calculated as $\mu_t^d = \beta_t^d - p_t^d$.

A.11 Proof for Proposition 5

Suppose along the equilibrium path the agent's continuation payoff is v_t. Similar to Equation (12), we want to show that $dv_t = -arv_t\beta_t(dY_t - \mu_t - m_t dt)$ where β_t is the short-term incentive slope offered along the equilibrium path. Because the agent's future rents are always zero (principals have all the bargaining power), it is easy to show that under the optimal saving policy the private saving balance follows $S_t = -\frac{1}{ar}\ln(-arv_t)$ with consumption policy $c_t = g(\mu_t) - \frac{1}{a}\ln(-arv_t)$. Because the principal has all the bargaining power, the fixed wage α_t satisfies

$$\alpha_t = g(\widehat{\mu}_t) + \frac{1}{2}ar\sigma^2\beta_t^2. \tag{A31}$$

Intuitively, the principal reimburses the agent's effort cost $\frac{\widehat{\mu}_t^2}{2}$, and compensates the risk premium $\frac{1}{2}ar\sigma^2\beta_t^2$ borne by the agent; they are just enough to convince the agent to take the offer. Then, in

equilibrium, $\mu_t = \widehat{\mu}_t$ and the agent's budget constraint reads

$$dS_t = r S_t dt - c_t dt + \alpha_t dt + \beta_t (dY_t - \mu_t - m_t dt)$$

$$= \frac{1}{2} ar\sigma^2 \beta_t^2 dt + \beta_t (dY_t - \mu_t - m_t dt).$$

Since $v_t = -\frac{1}{ar} \exp(-ar S_t)$, using Ito's lemma we have

$$dv_t = \exp(-ar S_t) dS_t + \frac{ar}{2} \exp(-ar S_t)(dS_t)^2 = -ar v_t \beta_t (dY_t - \mu_t - m_t dt).$$

Thus, the agent's continuation value process is identical to Equation (12). This also verifies that α_t in Equation (A31) is the minimum fixed wage needed to attract the agent.

Proposition 1 implies that the agent's incentive compatibility constraint satisfies

$$\mu_t = \beta_t - \mathbb{E}_t \left[\int_t^\infty \phi \beta_s e^{-(\phi+r)(s-t)} \exp\left(-\int_t^s ar\beta_u \sigma d B_u - \frac{1}{2} \int_t^s a^2 r^2 \beta_u^2 \sigma^2 du \right) ds \right] = \beta_t - p_t.$$

Importantly, because the principal t takes future β_{t+s} as given, the principal t is taking p_t as given and choosing β_t to maximize

$$\mathbb{E}_t[dY_t]/dt - \mathbb{E}_t[\alpha_t dt - \beta_t(dY_t - \mu_t dt - m_t dt)]/dt$$

$$= \mu_t + m_t - \alpha_t = (\beta_t - p_t) + m_t - g(\beta_t - p_t) - \frac{1}{2} ar\sigma^2 \beta_t^2.$$

Hence, the first-order condition for the optimal incentive β_t is $1 - (\beta_t - p_t) - ar\sigma^2 \beta_t = 0$, which implies that

$$\beta_t = \frac{1+p_t}{1+ar\sigma^2}.$$

We will see that this optimality condition does not hold in the long-term contracting case.

Conjecture that $\beta_t = \beta^{ST}$ and $p_t = p^{ST}$ are constants (we will verify this property shortly.) Then, since $p^{ST} = \frac{\phi}{\phi+r} \beta^{ST}$,

$$\beta^{ST} = \frac{\phi+r}{r+ar\sigma^2(\phi+r)}, p^{ST} = \frac{\phi}{r+ar\sigma^2(\phi+r)},$$

and the equilibrium effort is

$$\mu^{ST} = \beta^{ST} - p^{ST} = \frac{r}{r+ar\sigma^2(\phi+r)}.$$

Now let us rule out the case of timing-varying β. Recall that the feasible set of β_t is bounded by $[-M, M]$. Define $\overline{\beta} \equiv \sup\{\beta_t\} \in [-M, M]$. The optimality of short-term incentive implies that there exists t, so that $p_t = (1+ar\sigma^2)\overline{\beta} - 1 - \varepsilon$ for some sufficiently small ε. However, similar to the argument in Lemma 3, $p_t \leq \frac{\phi}{\phi+r}\overline{\beta}$, which implies

$$\left(1+ar\sigma^2\right)\overline{\beta} - 1 - \varepsilon \leq \frac{\phi}{\phi+r}\overline{\beta} \Rightarrow \left(\frac{r}{\phi+r} + ar\sigma^2\right)\overline{\beta} \leq 1+\varepsilon. \tag{A32}$$

Similarly, define $\underline{\beta} \equiv \inf\{\beta_t\} \in [-M, M]$, and we will have

$$\left(1+ar\sigma^2\right)\underline{\beta} - 1 + \varepsilon \geq \frac{\phi}{\phi+r}\underline{\beta} \Rightarrow -\left(\frac{r}{\phi+r} + ar\sigma^2\right)\underline{\beta} \leq -1+\varepsilon. \tag{A33}$$

Summing Equations (A32) and (A33), we have

$$\overline{\beta} - \underline{\beta} \leq \frac{2\varepsilon}{\frac{r}{\phi+r} + ar\sigma^2}.$$

Since ε is arbitrarily small, it must be $\overline{\beta} = \underline{\beta}$, and β_t is constant. Q.E.D.

References

Adrian, T., and M. Westerfield. 2009. Disagreement and learning in a dynamic contracting model. *Review of Financial Studies* 22:3873–906.

Baron, D. P., and D. Besanko. 1984. Regulation and information in a continuing relationship. *Information Economics and Policy* 1:267–302.

Bergemann, D., and U. Hege. 1998. Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance* 22:703–35.

Bergman, N., and D. Jenter. 2007. Employee sentiment and stock option compensation. *Journal of Financial Economics* 84:667–712.

Berk, J., and R. Green. 2004. Mutual fund flows and performance in rational markets. *Journal of Political Economy* 112:1269-95.

Biais, B., T. Mariotti, G. Plantin, and J. Rochet. 2007. Dynamic security design: Convergence to continuous time and asset pricing implications. *Review of Economic Studies* 74:345–90.

Core, J., and W. Guay. 1999. The use of equity grants to manage optimal equity incentive levels. *Journal of Accounting and Economics* 28:151–84.

Cosimano, T., A. Speight, and H. Yun. 2011. Executive compensation in a changing environment: Cautious investors and sticky contracts. Working Paper, University of Notre Dame.

Cvitanic, J., X. Wan, and H. Yang. 2013. Dynamics of contract design with screening. *Management Science* 59:1229–44.

DeMarzo, P., and M. Fishman. 2007. Optimal long-term financial contracting. *Review of Financial Studies* 20:2079–128.

DeMarzo, P., M. Fishman, Z. He, and N. Wang. 2012. Dynamic agency and the q-theory of investment. *Journal of Finance* 67:2295–340.

DeMarzo, P., and Y. Sannikov. 2006. Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance* 61:2681–724.

DeMarzo, P. M., and Y. Sannikov. 2017. Learning, termination and payout policy in dynamic incentive contracts. *Review of Economic Studies* 84:182–236.

Di Tella, S., and Y. Sannikov. 2016. Optimal asset management contracts with hidden savings. Working Paper, Stanford GSB.

Dittmann, I., and E. Maug. 2007. Lower salaries and no options? On the optimal structure of executive pay. *Journal of Finance* 62:303–43.

Edmans, A., and X. Gabaix. 2011. Tractability in incentive contracting. *Review of Financial Studies* 24:2865–94.

Eisfeldt, A., and A. Rampini. 2008. Managerial incentives, capital reallocation, and the business cycle. *Journal of Financial Economics* 87:177–99.

Fleming, W., and R. Rishel. 1975. *Deterministic and stochastic optimal control*. New York: Springer.

Fudenberg, D., Holmstrom, B., and P. Milgrom. 1990. Short-term contracts and long-term agency relationships. *Journal of Economic Theory* 51:1–31.

Garrett, D., and A. Pavan. 2012. Managerial turnover in a changing world. *Journal of Political Economy* 120: 879–925.

Gershkov, A., and M. Perry. 2012. Dynamic contracts with moral hazard and adverse selection. *Review of Economic Studies* 79:268–306.

Gibbons, R., and K. Murphy. 1992. Optimal incentive contracts in the presence of career concerns: Theory and evidence. *Journal of Political Economy* 100:468–505.

Golosov, M., M. Troshkin, and A. Tsyvinski. 2012. Optimal dynamic taxes. Working Paper, Princeton, Yale, and Minnesota.

Halac, M., N. Kartik, and Q. Liu. 2016. Optimal contracts for experimentation. forcoming in *Review of Economic Studies* 83:1040–91.

Hall, B., and J. Liebman. 1998. Are CEOs really paid like bureaucrats. *Quarterly Journal of Economics* 113: 653–91.

He, Z. 2009. Optimal executive compensation when firm size follows geometric brownian motion. *Review of Financial Studies* 22:859–92.

———. 2011. A model of dynamic compensation and capital structure. *Journal of Financial Economics* 100: 351–66.

———. 2012. Dynamic compensation contracts with private savings. *Review of Financial Studies* 25:1494–549.

He, Z., S. Li, B. Wei, and J. Yu. 2014. Uncertainty, risk, and incentives: Theory and evidence. *Management Science* 60:206–26.

Holmstrom, B. 1999. Managerial incentive problems: a dynamic perspective. *Review of Economic Studies* 66:169–82.

Holmstrom, B., and P. Milgrom. 1987. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica* 55:303–28.

Horner, J., and L. Samuelson. 2013. Incentives for experimenting agents. *Rand Journal of Economics* 44:632–63.

Ittner, C., R. Lambert, and D. Larcker. 2003. The structure and performance consequences of equity grants to employees of new economy firms. *Journal of Accounting and Economics* 34:89–127.

Ju, N., and X. Wan. 2012. Optimal compensation and pay-performance sensitivity in a continuous-time principal-agent model. *Management Science* 58:641–57.

Kocherlakota, N. 2004. Figuring out the impact of hiden savings on optimal unemployment insurance. *Review of Economic Dynamics* 7:541–54.

Liptser, R., and A. Shiryaev. 1977. *Statistics of random processes II: Applications*. In *Applications of mathematics*, ed. A. N. Shiryaev, Volume 6. New York: Springer.

Malenko, A. 2013. Optimal dynamic capital budgeting. Working Paper, MIT.

Marinovic, I., and F. Varas. 2016. CEO horizon, optimal duration and the escalation of short-termism. Working Paper, Stanford and Duke.

Medoff, J., and K. Abraham. 1980. Experience, performance and earnings. *Quarterly Journal of Economics* 95:703–36.

Murphy, K. 2003. Stock-based pay in new economy firms. *Journal of Accounting and Economics* 34:129–47.

Pavan, A., I. Segal, and J. Toikka. 2014. Dynamic mechanism design. *Econometrica* 82:601–53.

Prat, J., and B. Jovanovic. 2014. Dynamic contracts when agent's quality is unknown. *Theoretical Economics* 9:865–914.

Piskorski, T., and A. Tchistyi. 2010. Optimal mortgage design. *Review of Financial Studies* 23:3098–140.

Revuz, D., and M. Yor. 1999. *Continuous martingale and brownian motion*, 3rd ed. Springer.

Sannikov, Y. 2007. Agency problems, screening and increasing credit lines. Working Paper, Princeton University.

———. 2008. A continuous-time version of the principal-agent problem. *Review of Economic Studies* 75: 957–984.

———. 2014. Moral hazard and long-run incentives. Working paper, Princeton University.

Stulz, R., and C. Smith. 1985. The determinants of firm's dedging policies. *Journal of Financial and Quantitative Analysis* 20:391–406.

Sung, J. 2005. Optimal contracts under adverse selection and moral hazard: A continuous-time approach. *Review of Financial Studies* 18:1021–73.

Williams, N. 2009. On dynamic principal-agent problems in continuous time. Working Paper, University of Wisconsin-Madison.

———. 2011. Persistent private information. *Econometrica* 79:1233–74.

Zhang, Y. 2009. Dynamic contracting with persistent shocks. *Journal of Economic Theory* 144:635–75.