Incentives and Performance with Optimal Money Management Contracts

Stefano Pegoraro*

October 29, 2021

Abstract

I characterize the dynamics of incentives in an optimal contract with investment delegation, moral hazard, and uncertainty about the agent's productivity. The principal increases the agent's incentives after good performance in order to delegate more capital to an agent with higher perceived productivity, thus implementing a convex pay-for-performance scheme. Moreover, the principal commits to reduce the agent's future incentives in order to mitigate ex-ante investment distortions. Methodologically, I provide a duality-based strategy to overcome technical challenges common to continuous-time contracting models with state variables. I also derive a sufficient condition to verify the validity of the first-order approach.

Keywords: Dynamic contracts, managerial compensation, duality, dynamic programming, delegated investment.

1 Introduction

In many situations, principals delegate the management of capital to agents with uncertain productivity. For example, investors rely on fund managers with unknown skill to invest their savings; venture capitalists finance entrepreneurs developing innovative

*University of Notre Dame, Mendoza College of Business – s.pegoraro@nd.edu. I am immensely grateful to Lars Peter Hansen, Zhiguo He, Douglas Diamond, and Pietro Veronesi. I would also like to thank Ehsan Azarmsa, Markus Broman, Benjamin Brooks, Kirill Chernomaz, Sebastian Di Tella, Jacques Cremer, David Feldman, Jian Li, Nicola Limodio, Johnathan Payne, Lubos Pastor, Martin Schmalz, Luis Simon, Panagiotis Souganidis, Harald Uhlig, Yuyao Wang, three anonymous referees, and seminar participants at the University of Chicago, SITE (Banks and Financial Frictions), Chicago Booth Asset Pricing Conference, Illinois Economic Association, Fed Board, Texas A&M, UCL, Bocconi, Stockholm School of Economics, Wharton, Notre Dame, CESIfo, Northern Finance Association, Midwestern Finance Association, Western Economics Association, Midwest Economic Association, Royal Economic Society, French Finance Association, Canadian Economic Association, French Economic Association, Econometric Society (Asia and China meetings), and Midwest Trade and Theory Conference. For the latest version of this paper, please visit https://sites.nd.edu/stefano-pegoraro/research/.

products; and CEOs allocate resources to division managers with undetermined ability. If the agent is subject to moral hazard, the principal designs an incentive contract for the agent. As the principal learns about the agent's productivity over time, she may dynamically change the agent's incentives and capital under management. However, models of dynamic incentive provision with learning are difficult to solve. Therefore, although delegated money management is widespread in the world, the theoretical literature presents a noticeable gap. In this paper, I fill this gap.

I make three main contributions. First, I develop a continuous-time contracting model of money management with uncertainty about the agent's productivity,[1] and I derive a sufficient condition for the validity of the first-order approach. Second, I uncover a novel mechanism that generates a convex compensation scheme for the agent. This mechanism highlights how the agent's incentives change in response to past performance. Third, I provide a duality-based methodology to overcome the technical challenges posed by contracting models with state variables. This methodology can be directly applied to a wide class of such models.

In the model, a risk-neutral principal writes a long-term contract for a risk-averse agent with uncertain productivity. The agent operates a constant-returns-to-scale technology, but suffers a private cost of effort proportional to the capital under management. To prevent shirking, the principal provides incentives to the agent by exposing him to performance risk. Along the equilibrium path, the principal and the agent symmetrically learn about the agent's productivity by observing performance, and they share common beliefs. However, an off-equilibrium deviation by the agent causes a persistent distortion in the principal's beliefs, giving the agent an information rent. Because of the agent's information rent, agency frictions are exacerbated: Everything else being equal, the principal delegates a smaller amount of capital to the agent when the agent's information rent is larger.

The agent enjoys larger information rents when he expects larger exposures to risk in the future. Hence, to lower the agent's information rent, the principal commits to reduce the agent's future incentives, along with delegated capital.[2] Although optimal ex

[1] The dynamic contracting literature has been predominately applied outside the context of investment delegation. Typical applications of the dynamic contracting theory include executive compensation (Edmans et al., 2012; Garrett and Pavan, 2012; He, 2012), delegation of experimentation (Guo, 2016; Halac et al., 2016; Hörner and Samuelson, 2013), optimal taxation (Farhi and Werning, 2013; Golosov et al., 2003; Kocherlakota, 2005), dynamic price discrimination (Battaglini, 2005) and security design (Biais et al., 2007; DeMarzo and Sannikov, 2006).

[2] Such commitment is common in models where future risk exposure exacerbates current frictions either because of information rents (He et al., 2017; Prat and Jovanovic, 2014), or because of the agent's precautionary saving motive (Di Tella and Sannikov, 2021). However, the path of incentives differs substantially from models with career concerns (Fama, 1980; Holmstrom, 1999) where implicit incentives from career concerns

ante, such commitment is costly ex post for the principal. To diminish the expected cost of commitment, the principal optimally allocates her commitment as a function of past performance. After good performance, the principal rewards the agent with higher future compensation, and therefore suffers higher costs if committed to reduce the agent's incentives and capital. As a result, the principal relaxes her commitment after good performance and reinforces it after bad performance.

The model predicts the agent's incentives increase after a history of good performance and they decline after a history of bad performance. The result relies on the interaction between moral hazard and learning. Because of learning, the agent's estimated productivity increases after good performance. The principal then wants to allocate more capital to the agent.[3] However, because of moral hazard, the principal must increase the agent's incentives along with delegated capital in order to prevent shirking. Furthermore, after good performance, the principal also relaxes her commitment and allows herself to provide steeper incentives and more capital to the agent.

Based on this mechanism, the model generates a convex relation between the agent's compensation and his cumulative performance. With a convex compensation scheme, the agent's compensation becomes more (less) sensitive to current performance after a history of good (bad) performance, thus reflecting the optimal allocation of incentives described above. The model's mechanism, and the resulting convex incentive scheme, has been observed empirically in various settings. Corporate executives frequently receive performance bonuses in the form of company shares or stock options (Edmans et al., 2017; Hall and Liebman, 1998); hedge fund managers earn fees on the fund's profits and typically reinvest the fees back in the fund (Agarwal et al., 2009); entrepreneurs are granted larger equity stakes by venture capitalists after good performance, especially when the venture's quality is uncertain (Kaplan and Strömberg, 2003). In all these cases, the agent's risk exposure increases after good performance, resulting in a convex incentive scheme.

Methodologically, I provide a general technique to feasibly solve continuous-time contracting models with state variables. These models often require a solution to a differential equation with an endogenous boundary in the state space. Previous contributions developed specialized methods that could not be directly extended outside their model.[4]

decline over time, but explicit incentives from short-term contracts increase (Gibbons and Murphy, 1992).

[3] Empirically, Graham et al. (2015) document that CEOs and CFOs allocate capital to division managers based on the managers' reputation and past performance.

[4] DeMarzo and Sannikov (2017) exploit the structure of their model to reformulate the problem in terms of dynamically changing termination thresholds. He et al. (2017) use a guess-and-verify numerical approach that cannot be feasibly applied to differential equations in more than one dimension. Prat and Jovanovic (2014) and Williams (2011) make parametric assumptions on the utility function and on the law of motion of the state variable to obtain analytical solutions.

In this paper, I use duality methods to provide a general solution strategy that applies to a wide class of models, including those with differential equations in more than one dimension.

In the dual problem, a new state variable, a multiplier, directly captures the commitment of the principal to reduce future incentives. Besides allowing for a more transparent representation of the dynamic allocation of commitment, the dual problem offers a key simplification relative to the standard approach. In the dual problem, the endogenous boundary on the agent's information rent is replaced by an exogenous non-negativity constraint on the multiplier. The optimal contract can then be conveniently characterized by the solution to the dual problem after simply initializing the multiplier to zero.

RELATION TO THE LITERATURE. This paper belongs with the growing literature that uses recursive methods to solve for optimal contracts (Phelan and Stacchetti, 2001; Phelan and Townsend, 1991; Sannikov, 2008; Spear and Srivastava, 1987; Williams, 2009). The model features persistent private information, similar to Battaglini (2005), Di Tella and Sannikov (2021), Kapička (2013), Farhi and Werning (2013), Fernandes and Phelan (2000), Pavan et al. (2014), and Williams (2011). Therefore, the first-order approach is not guaranteed to be valid (Battaglini and Lamba, 2019; Kocherlakota, 2004). I contribute to this literature by providing a sufficient incentive-compatibility condition that I use to verify the validity of the first-order approach. This condition states that, if a contract prevents one-shot deviations and reduces the marginal value of the agent's private information after bad performance, the contract is incentive compatible.

I also extend the literature on dynamic contracting with learning (Bergemann and Hege, 1998, 2005; DeMarzo and Sannikov, 2017; Halac et al., 2016; He et al., 2017; Hörner and Samuelson, 2013; Prat and Jovanovic, 2014) by allowing for endogenous capital under management. Because of the interaction of learning, moral hazard, and investment delegation, the model generates a convex incentive scheme for the manager. As the manager's assessed productivity increases with past performance, the principal desires to allocate more capital to him. She thus offers steeper incentives because more money is at stake.

This mechanism for a convex incentive scheme is new. In Dittmann et al. (2010), convex compensation contracts encourage the agent to take risks, whereas in Li and Tiwari (2009), they motivate the agent to gather and use private information when forming a portfolio of assets. In dynamic settings, Edmans and Gabaix (2011) obtain convex compensation contracts when marginal costs of effort are high, whereas He et al. (2017) obtain option-like incentives if the principal can commit to establish a negative correlation between the agent's incentives and marginal utility of consumption.

Finally, I develop a duality-based approach to overcome the mathematical challenges that are common to continuous-time contracting models with state variables. This approach could be applied to models with learning, like in this paper, models with hidden savings like the model in Di Tella and Sannikov (2021), or models with arbitrary state variables, like the model in Williams (2011). Within the literature that uses duality methods in contracting models, my paper is therefore close in spirit to Marcet and Marimon (2019) and Pavoni et al. (2018), who focus on discrete-time models. Other papers applying duality theory in contracting models, although for different purposes, include Miao and Zhang (2015), Sannikov (2014), Arie (2016), Carrasco et al. (2019), Myerson (1984), and Myerson (1985).

2 MODEL

A principal (she) writes a long-term contract for an agent (he). The principal delegates the management of capital to the agent. However, the agent's ability to generate returns is uncertain, and both the principal and the agent learn about the agent's productivity by observing realized returns. Moreover, the agent is subject to moral hazard, because he might secretly shirk and obtain a private benefit at the expense of the principal. Time is continuous and starts at 0. A complete probability space $(\Omega, \bar{\mathcal{F}}, P)$ with filtration $(\bar{\mathcal{F}}_t)_{t\geq 0}$ is given.

2.1 PLAYERS AND TECHNOLOGY

The principal hires the agent to invest capital in a constant-returns-to-scale technology. However, the agent's productivity is unknown to all players, including the agent himself. The agent's hidden type, $h \in \{0, 1\}$, is a random variable on the probability space $(\Omega, \bar{\mathcal{F}}, P)$ and it indicates whether the agent is unskilled ($h = 0$) or skilled ($h = 1$.) Players have a common prior $\phi_0 := \mathrm{E}[h|\bar{\mathcal{F}}_0] \in [0, 1]$. The agent's unknown skill constitutes the source of uncertain productivity.[5]

The principal obtains capital from a competitive capital market with perfectly elastic supply. When managed by the agent, capital produces excess returns given by

$$dR = (\sigma\eta h - m_t)dt + \sigma dW_t,$$

[5]Uncertain productivity may be due to factors other than the agent's skill. For example, the quality of the underlying technology itself may be uncertain because the agent is making unconventional investments.

where $\eta \geq 0$ and $\sigma \geq 0$ are known parameters, and where $(W_t)_{t\geq 0}$ is a $\bar{\mathcal{F}}_t$-adapted standard Brownian motion on the complete probability space $(\Omega, \bar{\mathcal{F}}, P)$, and it is independent of h.

Returns depend on the uncertain skill of the agent, h, and on the agent's hidden action, $m_t \geq 0$, which is positive if the agent shirks. The agent generates superior returns only if he is skilled ($h = 1$). The parameter $\eta > 0$ is the signal-to-noise ratio of returns, which measures how informative returns are regarding the agent's productivity. If the agent shirks to obtain a private benefit, he reduces returns for the principal at rate m_t for each unit of capital invested.

At time t, the principal invests capital $K_t \geq 0$, collects the excess returns the agent produces, dR_t, and offers a compensation $C_t \geq 0$ to the agent. The principal is risk neutral, and her cost of employing the agent is

$$\mathrm{E}\left[\int_0^\infty e^{-rt}(C_t dt - K_t dR_t)\Big|\bar{\mathcal{F}}_0\right], \tag{1}$$

where $r > 0$.

At time t, the agent receives compensation C_t from the principal, actively invests capital K_t, and shirks at rate m_t. His lifetime utility is

$$\mathrm{E}\left[\int_0^\infty e^{-\delta t}u(C_t + m_t\lambda K_t)\,dt\Big|\bar{\mathcal{F}}_0\right], \tag{2}$$

where $\lambda \in (0,1)$, $\delta \geq r(1-\rho)$, and $u(x) = \frac{x^{1-\rho}}{1-\rho}$ with $\rho \in (0, 1/2)$.[6]

If the agent shirks at rate m_t, he reduces the principal's cash-flow rate by $m_t K_t$. However, the agent obtains a private consumption benefit of only $\lambda m_t K_t$. Shirking is therefore inefficient because the agent destroys more value that what he obtains, and $1-\lambda$ represents the inefficiency of shirking. Full effort coincides with $m_t = 0$, and the consumption value of shirking can be equivalently interpreted as the cost of effort.

2.2 Contracting Environment and Learning

Returns constitute the only source of information for the players. Players form estimates of the agent's productivity by observing his performance. Moreover, because return realizations depend on the agent's hidden action, the principal may reward or punish the agent based on his performance in order to deter shirking. In this framework, players'

[6]The assumption that $\rho < 1/2$ is needed to obtain a finite solution to the model. If $\rho > 1/2$, the agent's marginal utility of consumption would decline quickly enough that the principal would find it profitable to give infinite capital and infinite consumption to the agent and overcome the incentive problem. See also Di Tella (2019).

information is generated by the history of returns. Let $(R_s)_{0\leq s\leq t}$ denote the history of returns up to time t. I define $\mathcal{F}_t := \{(R_s)_{0\leq s\leq t}\}$ as the smallest σ-algebra for which $(R_s)_{0\leq s\leq t}$ is measurable, possibly augmented by the P-null sets. Thus, $\mathcal{F}_0 = \bar{\mathcal{F}}_0$ and $(\mathcal{F}_t)_{t\geq 0}$ is the augmented filtration generated by the history of returns.

A contract between the principal and the agent specifies the agent's compensation and his capital under management as a function of past performance.

DEFINITION 1 (CONTRACT). *A contract $\mathcal{C}$ is a $\mathcal{F}_t$-adapted process $(C_t, K_t)_{t\geq 0}$.*

Although the principal cannot directly control the agent's hidden action, she understands the implications of a contract on the agent's incentives to shirk. In general, we say that a contract $\mathcal{C} = (C_t, K_t)_{t\geq 0}$ is *incentive compatible with the shirking process* $(m_t)_{t\geq 0}$ if the latter is the agent's best response to contract $\mathcal{C}$.

Given a contract $\mathcal{C} = (C_t, K_t)_{t\geq 0}$ that is incentive compatible with $(m_t)_{t\geq 0}$, players symmetrically learn about the agent's productivity by observing realized returns. They therefore possess common beliefs about the agent's skill,

$$\phi_t := \mathrm{E}[h|\mathcal{F}_t],$$

starting from the common prior ϕ_0. As players observe returns, beliefs ϕ_t evolve as

$$d\phi_t = \eta\phi_t(1-\phi_t)dW_t^{\mathcal{C}}, \tag{3}$$

where

$$dW_t^{\mathcal{C}} := \frac{1}{\sigma}\left[dR_t - (\sigma\eta\phi_t - m_t)\right], \tag{4}$$

is the increment of a standard Brownian motion under the measure of returns induced by $\mathcal{C}$. Lemma O.1 in the online appendix provides a formal proof of (3).

Equation (3) highlights that changes in beliefs, $d\phi_t$, are positively correlated with the agent's performance dR_t, reflecting the intuition that an agent that performs well (poorly) is more (less) likely to be skilled. The sensitivity of beliefs to performance, $\eta\phi_t(1-\phi_t)$, quantifies the amount of information that players obtain from returns and hence the speed of learning.

When optimally designing a contract, the principal minimizes the costs of delegation, after taking into account the incentives that the agent receives from the contract. Formally, an optimal contract can be defined as follows.

DEFINITION 2 (OPTIMAL CONTRACT). *Given a prior ϕ_0 and an outside option V_0 for the agent, a contract $\mathcal{C} = (C_t, K_t)_{t\geq 0}$ is optimal if it minimizes the principal's cost* (1)*, while offering the agent an expected lifetime utility* (2) *at least as large as V_0.*

The principal is aware the agent might shirk and gain private benefits. However, the following lemma shows that any optimal contract is designed so that the agent is not tempted to shirk. Therefore, when looking for an optimal contract, I can restrict the search over the class of contracts that are incentive compatible with full effort.

LEMMA 1. *Any optimal contract is incentive compatible with no shirking, that is,* $m_t = 0 \, \forall t \geq 0$.

All proofs are in the appendix.

The intuition for this lemma is the following. First, by the revelation principle, we lose no generality if we restrict attention to incentive-compatible contracts. Second, optimal contracts must align the agent's incentives with the principal's objectives. Because shirking is inefficient, the principal designs an optimal contract that motivates the agent to exert full effort.

2.3 INCENTIVE COMPATIBILITY AND INFORMATION RENT

Because of Lemma 1, I now focus on contracts that are incentive compatible with no shirking and, for brevity, I simply refer to them as *incentive-compatible contracts*. I provide a necessary condition that incentive-compatible contracts must satisfy. As in static principal-agent problems, this condition requires the agent to be exposed to some risk. Whereas in static models, the principal exposes the agent to risk by giving him performance-contingent pay, in a dynamic model, the principal exposes the agent to risk by adjusting his future continuation value based on performance.

The agent's continuation value at time t, denoted by V_t, represents the future expected utility for an agent that does not shirk. The continuation value V_t is a function of the continuation contract, $\mathcal{C}_t$,[7] and beliefs at time t, and it can be expressed as[8]

$$V_t = \tilde{V}(\mathcal{C}_t, \phi_t) \coloneqq \mathrm{E}\left[\int_t^\infty e^{-\delta(s-t)} u(C_s)\, ds \middle| \mathcal{F}_t\right]. \tag{5}$$

Using the martingale-representation approach developed in previous literature (Sannikov, 2008; Williams, 2009), I obtain the law of motion for the agent's continuation value

[7] A continuation contract at time t, $\mathcal{C}_t$, is a $\mathcal{F}_s$-adapted process $(C_s, K_s)_{s \geq t}$.

[8] Given a continuation contract $\mathcal{C}_t$, beliefs are a sufficient statistic for the probability measure of future returns. To see why, write

$$V_t = (1-\phi_t)\mathrm{E}\left[\int_t^\infty e^{-\delta(s-t)} u(C_s)\, ds \middle| h = 0\right] + \phi_t \mathrm{E}\left[\int_t^\infty e^{-\delta(s-t)} u(C_s)\, ds \middle| h = 1\right].$$

The conditional expectations on the right-hand side of this equation are functions of the continuation contract $\mathcal{C}_t$ only.

V_t as

$$dV_t = (\delta V_t - u(C_t))dt + \beta_t dW_t^{\mathcal{C}} \quad \text{with} \quad \lim_{t\to\infty} \mathrm{E}\left[e^{-\delta t} V_t | \mathcal{F}_0\right] = 0, \tag{6}$$

for some $\mathcal{F}_t$-adapted process $(\beta_t)_{t\geq 0}$.

If β_t is positive, the future utility of the agent increases with performance. Hence, β_t represents the risk exposure of the agent. Although the risk-neutral principal should fully insure the risk-averse agent in an efficient allocation, the agent will be exposed to risk in the optimal contract in order to mitigate the moral hazard friction. If $\beta_t > 0$, the agent suffers a loss of future utility when returns decline because of shirking. Therefore, the agent might be deterred from shirking if his risk exposure, β_t, is large enough.

As a benchmark, assume for the moment that the agent's type, h, is known. In this case, the principal prevents shirking by offering a contract in which the agent's risk exposure, β_t, offsets the marginal consumption value of shirking. A necessary and sufficient condition for incentive compatibility with no shirking is

$$u'(C_t)\lambda\sigma K_t \leq \beta_t. \tag{7}$$

The intuition for this condition is the following.[9] If the agent shirks, he reduces returns by m_t, and hence $dW_t^{\mathcal{C}}$ acquires a negative drift $-\frac{m_t}{\sigma}$. The agent therefore suffers a loss of continuation value equal to $\beta_t \frac{m_t}{\sigma}$. However, his current utility increases from $u(C_t)$ to $u(C_t + m_t\lambda K_t)$. Condition (7) ensures $m_t = 0$ is the agent's best response to this trade off; that is, $0 = \arg\max_{m\geq 0}\{[u(C_t + m\lambda K_t) - u(C_t)] - \beta_t \frac{m}{\sigma}\}$.

However, if the agent's type is uncertain, the principal faces some additional challenges in designing an incentive-compatible contract. Suppose the agent deviates to $m_t = m > 0$ for a small amount of time between s and $s + \Delta s$. With learning, the agent not only gains consumption value from the deviation, but also earns an informational advantage over the principal. Unaware of the agent's deviation, the principal updates beliefs according to equations (3) and (4) with $m_t = 0$ for all t. Because the correct expected return is $\sigma\eta\phi_t - m$ for $t \in [s, s + \Delta s]$, the principal's beliefs, ϕ_t, acquire a negative drift relative to the agent's beliefs, ϕ_t^A. Immediately after the deviation, the difference between the agent's and the principal's beliefs, $\phi_t^A - \phi_t$, is given by $\phi_{s+\Delta s}^A - \phi_{s+\Delta s} \approx \eta\phi_s(1-\phi_s)\frac{m}{\sigma}\Delta s$.

This difference in beliefs is persistent and causes persistent distortions in the provision of incentives. Following a deviation, the agent will always be more optimistic than the principal; that is, $\phi_t^A > \phi_t$ for all $t > s + \Delta s$. The agent is not only more optimistic, but also aware of possessing correct beliefs. By having more accurate and optimistic beliefs,

[9] A proof can be found in Di Tella (2019) and Di Tella and Sannikov (2021).

the agent earns an information rent over the principal.[10]

To see how the agent earns an information rent, consider equation (6). If $\phi_t^A > \phi_t$, shocks $dW^{\mathcal{C}} := \frac{1}{\sigma}\left[dR_t - \sigma\eta\phi_t dt\right]$ have a positive drift equal to $\eta(\phi_t^A - \phi_t)dt > 0$ based on the agent's information. Therefore, the agent's continuation value acquires an additional drift equal to $\beta_t\eta(\phi_t^A - \phi_t)dt$. This additional drift captures the surprise the agent expects the principal to receive. Suppose that after a (hidden) deviation, the principal promises the agent a continuation value $V_{s+\Delta s}$. The future expected utility of the agent, conditional on not shirking is the future, is larger than $V_{s+\Delta s}$ by an amount equal to

$$\mathrm{E}\left[\int_{s+\Delta s}^{\infty} e^{-\delta(t-s-\Delta s)}\beta_t\eta(\phi_t^A - \phi_t)dt\Big|\mathcal{F}_{s+\Delta s}\right].$$

This additional future utility constitutes the information rent the agent acquired after the deviation.

This reasoning suggests that when players are learning, an incentive-compatible contract should provide risk exposure β_t large enough to offset the agent's marginal utility of shirking, as well as the marginal information rent that he could earn.

THEOREM 1 (NECESSARY CONDITION)**.** *If $\mathcal{C} = (C_t, K_t)_{t\geq 0}$ is incentive compatible with no shirking, then*

$$u'(C_t)K_t\lambda\sigma \leq \beta_t - \eta\xi_t, \tag{8}$$

where ξ_t follows

$$d\xi_t = (\delta\xi_t - \beta_t\eta\phi_t(1-\phi_t))dt + \omega_t dW_t^{\mathcal{C}} \quad \textit{with} \quad \lim_{t\to\infty}\mathrm{E}\left[e^{-\delta t}\xi_t|\mathcal{F}_0\right] = 0 \tag{9}$$

for some $\mathcal{F}_t$-adapted process $(\omega_t)_{t\geq 0}$.

Similar to (7), condition (8) highlights that, everything else being equal, the principal has to increase the agent's risk exposure β_t to increase the amount of delegated capital K_t in an incentive-compatible contract. Thus, because the agent is risk averse, delegation is costly. In the optimal contract, the principal will therefore trade off the costs and benefits of delegation when deciding the level of investment and the agent's risk exposure.

The term $\eta\xi_t$ in (8), which is missing in (7), accounts for the marginal information rent that the agent acquires after an instantaneous deviation. Solving (9), we can express the information rent ξ_t as the present value of the agent's future risk exposure and speed of

[10]Intuitively, a skilled agent who is believed to be unskilled is better off than an agent who is actually unskilled. The skilled agent can expect to surprise the principal in the future thanks to his superior skill. The truly unskilled agent cannot expect to surprise anyone.

learning; specifically,

$$\xi_t = \mathrm{E}\left[\int_t^\infty e^{-\delta(s-t)}\beta_s\eta\phi_s(1-\phi_s)\,ds\middle|\mathcal{F}_t\right], \tag{10}$$

which is a positive quantity in any incentive-compatible contract, by Lemma O.2 in the online appendix.

Theorem 1 thus also implies that because of the agent's information rent, investment distortions are exacerbated relative to a model with perfect information. The larger the information rent, the lower the amount of capital that can be delegated to the agent for a given compensation C_t and risk exposure β_t. By reducing the amount of capital under the agent's management, the principal foregoes profitable investment opportunities.

Therefore, from an ex-ante perspective, the principal prefers to lower the agent's information rent. To do so, she has to commit to reduce the agent's future risk exposures, $(\beta_t)_{t\geq 0}$, because of (10). However, the principal may be tempted to re-negotiate the contract and increase risk exposure ex post. Consequently, to achieve ex-ante optimality, the principal must fully commit to the terms of an initial contract.

To offer an interpretation of (8), consider expression (5). Let $\partial_\phi \tilde{V}(\mathcal{C}_t, \phi_t)$ measure the marginal change in the continuation value due to a marginal change in beliefs while keeping the continuation contract fixed. Lemma O.2 in the online appendix shows that

$$\xi_t = \phi_t(1-\phi_t)\partial_\phi \tilde{V}(\mathcal{C}_t, \phi_t). \tag{11}$$

From equation (6), we know the risk exposure β_t corresponds to the total volatility of the continuation value V_t. Combining (3) and (11), we see $\eta\xi_t = \eta\phi_t(1-\phi_t)\partial_\phi V(\mathcal{C}_t, \phi_t)$ represents the volatility of the continuation value driven by changes in beliefs. Therefore, we can interpret the quantity

$$\beta_t - \eta\xi_t = \beta_t - \eta\phi_t(1-\phi_t)\partial_\phi V(\mathcal{C}_t, \phi_t)$$

as the volatility of the agent's continuation value driven by changes in the continuation contract while keeping beliefs fixed.

Therefore, equation (8) states that to provide incentives to the agent, the principal cannot rely on changes in beliefs to punish him for bad performance. Although changes in beliefs do affect the volatility of the continuation value along the equilibrium path, they cannot be exploited to prevent off-equilibrium deviations. In fact, following a deviation m_t, the agent's true expected future utility declines by $(\beta_t - \eta\xi_t)\frac{m_t}{\sigma}$, and not by $\beta_t\frac{m_t}{\sigma}$, as the principal incorrectly thinks. Therefore, in an incentive-compatible contract, the quantity

$\beta_t - \eta\xi_t$ is what matters for incentive provision, and must be such that

$$0 = \arg\max_{m \geq 0} \quad \underbrace{u(C_t + m\lambda K_t) - u(C_t)}_{\text{benefit of shirking}} - \underbrace{(\beta_t - \eta\xi_t)\frac{m}{\sigma}}_{\text{cost of shirking}}.$$

In other words, the incentive-compatibility condition asserts that to prevent shirking, the principal must adjust the continuation contract so that, keeping beliefs constant, the risk exposure of the agent exceeds the private benefit of shirking.

2.4 Sufficient Incentive-Compatibility Condition

Theorem 1 offers a condition that prevents the agent from engaging in an instantaneous, "one-shot", deviation. Although this condition is necessary for incentive compatibility, alone it does not guarantee incentive compatibility. With learning, any hidden shirking will cause a persistent wedge between the agent's and principal's beliefs. Given this wedge, the agent's best response to the contract may still involve a dynamic shirking strategy even if (8) holds. Next, I provide a sufficient condition that ensures that no shirking is the agent's best response to a contract.

Previous literature has long recognized the challenge that state variables, such as beliefs, pose in the design of an optimal contract. The common approach is to solve for an optimal contract by imposing the necessary condition (8) only. This approach is often referred to as the relaxed-problem approach. Then, one should verify whether the contract so obtained satisfies a sufficient incentive-compatibility condition. This strategy is the one undertaken by He (2012) and Di Tella and Sannikov (2021) for private savings, by Prat and Jovanovic (2014), DeMarzo and Sannikov (2017), He et al. (2017), and Cisternas (2018) for learning, and Williams (2011) for generic state variables. I take the same approach and use the following theorem to verify whether the solution to the relaxed problem is actually incentive compatible.

Theorem 2 (Sufficient Condition)**.** *If a contract* $\mathcal{C} = (C_t, K_t)_{t \geq 0}$ *is such that* (8) *and*

$$\omega_t - \eta(1 - 2\phi_t)\xi_t \geq 0 \tag{12}$$

hold for all $t \geq 0$*, then the contract is incentive compatible with no shirking.*

To interpret the sufficient condition for incentive compatibility, it is useful to refer to the proof of this theorem in the appendix, where I show $\omega_t - \eta(1 - 2\phi_t)\xi_t$ is proportional to the volatility of $\partial_\phi \tilde{V}(\mathcal{C}_t, \phi_t)$. Theorem 2 therefore states that if a contract prevents "one-shot" deviations thanks to (8), and it reduces the marginal value of beliefs after a negative shock thanks to (12), then the contract is incentive compatible with no shirking.

This result has an intuitive appeal. If the contract lowers the marginal value of the agent's beliefs $\partial_\phi \tilde{V}(\mathcal{C}_t, \phi_t)$ after a bad shock, then, following a deviation, the agent would experience a decline not only in his continuation value, but also in his information rent. The agent loses part of the option to "impress" the principal in the future, thus lowering the value of his informational advantage.

Equation (12) is likely to hold in an optimal contract. Looking at equation (10), we see ξ_t depends on future risk exposures $(\beta_s)_{s\geq t}$. After a positive shock, the agent's estimated productivity increases by (3). Hence, the principal may want to increase future capital under management K_t, to take advantage of increased expected returns. By the incentive-compatibility condition (8), the principal will then have to increase the agent's future risk exposure as well. We therefore have reasonable economic motivations to expect that in the optimal contract, ξ_t increases after good performance (or at least it does not decline too much), in such a way that condition (12) is satisfied.

3 OPTIMAL CONTRACT AND DUALITY

To solve for the optimal contract, I proceed as follows. First, I define the relaxed optimal contract, in which the principal minimizes costs subject to the necessary incentive-compatibility condition (8). Then, I introduce the dual problem and establish its relation to the relaxed optimal contract. Finally, I provide a verification result to characterize the relaxed optimal contract using the Hamilton-Jacobi-Bellman (HJB) equation associated with the dual problem, and to verify whether the contract is an optimal (incentive-compatible) contract.

3.1 THE RELAXED PROBLEM

I adopt a first-order approach and solve for a contract that minimizes the cost for the principal within the class of contracts satisfying (8). This is the so-called relaxed optimal contract.

DEFINITION 3 (RELAXED OPTIMAL CONTRACT)**.** *Given the prior ϕ_0 and the agent's outside option V_0, a relaxed optimal contract is a contract that minimizes the cost for the principal* (1) *subject to the necessary incentive-compatibility condition* (8) *and subject to delivering expected lifetime utility V_0 to the agent if he does not shirk.*

Because (8) is a necessary but not sufficient condition for the agent not to shirk, the cost of a relaxed optimal contract provides a lower bound for the cost of an optimal contract. Although I only impose the necessary condition (8) in a relaxed optimal contract, I use the result in Theorem 2 to numerically verify whether the relaxed optimal contract is fully incentive compatible with no shirking. In this case, the relaxed optimal contract is an optimal contract.

The principal fully commits to the contract she offers.[11] Because of this commitment, the agent's continuation value V_t and information rent ξ_t are recursive state variables of the contract-design problem. At every point in time, the (forward-looking) continuation value and information rent of the agent are determined by past promises of the principal, and the continuation contract should be optimal given these promises.

For a given information rent ξ_0 at time 0, consider the following cost-minimization problem for the principal:

$$
\begin{aligned}
J^*(V_0, \xi_0, \phi_0) = \inf_{\mathcal{C}} \mathrm{E}\left[\int_0^\infty e^{-rt}\left(C_t - \eta\phi_t \frac{C_t^\rho}{\lambda}(\beta_t - \eta\xi_t)\right) dt \middle| \mathcal{F}_0\right] \\
\text{s.t. } C_t \geq 0, \beta_t \geq \eta\xi_t \\
(6), (9), \text{ and } (3) \text{ with } m_t = 0\ \forall t \geq 0.
\end{aligned}
\tag{13}
$$

I restrict my attention to admissible contracts, as in Di Tella and Sannikov (2021). These are contracts in the form $\mathcal{C} = (C_t, K_t)_{t\geq 0}$ such that

$$
\mathrm{E}\left[\int_0^\infty e^{-rt}(|C_t| + |K_t| + |\beta_t|)\, dt \middle| \mathcal{F}_0\right] < \infty, \tag{14}
$$

thus ensuring the principal's objective function (1) is well defined and that a strong solution to the stochastic differential equation in (6) exists starting from a given V_0.

In this formulation, the principal's objective function already incorporates the necessary incentive-compatibility constraint (8) as an equality. Because expected returns are always positive, the principal optimally increases the level of assets under management as much as the incentive-compatibility constraint permits. Therefore, (8) is always binding. Hence, going forward, a contract $\mathcal{C}$ could equivalently specify compensation and capital, $(C_t, K_t)_{t\geq 0}$, or compensation and risk exposure, $(C_t, \beta_t)_{t\geq 0}$.

The constraints $C_t \geq 0$ and $\beta_t \geq \eta\xi_t$ impose the non-negativity of consumption and capital, whereas the constraints (6) and (9) represent promise-keeping constraints on the

[11] With full commitment, the principal implements an allocation that yields the best outcome given the frictions of the model. In the supplemental material, I focus on contracts with limited commitment, and I highlight which results hold independently of the commitment assumption.

agent' continuation value and information rent. Finally, (3) with $m_t = 0$ is the law of motion of beliefs under Bayesian learning and no shirking.

Therefore, a solution to problem (13) is a relaxed optimal contract if ξ_0 is an optimal choice of the principal. In fact, according to Definition 3, the initial outside option of the agent, V_0, and the initial beliefs, ϕ_0, are exogenously given. However, the definition imposes no constraints on the initial information rent of the agent ξ_0. Therefore, the principal chooses the initial information rent to minimize costs; that is,

$$\xi_0 = \arg\min_{\xi \geq 0} J^*(V_0, \xi, \phi_0). \tag{15}$$

In choosing the initial information rent for the agent, the principal faces a trade-off. By choosing a small information rent ξ_0, the principal reduces the ex-ante distortions in the allocation of capital captured by the incentive-compatibility condition (8). However, to deliver a small information rent, the principal must commit to limit the agent's risk exposure in the future, as expression (10) indicates. In the extreme case in which the information rent ξ_0 is set to zero, the principal would have to immediately "retire" the agent, that is, offer a safe contract with no risk exposure for the agent and no capital under management.

Normally, to solve for a relaxed optimal contract, one would then proceed in two steps. First, one would solve problem (13) using dynamic programming methods. Second, one would solve for ξ_0 using (15). Unfortunately in the current situation, this approach is not feasible, either analytically or numerically. In the optimal contract, the information rent at any time $t \geq 0$, ξ_t, is bounded from above by a function of V_t and ϕ_t.[12] For values of information rent above the upper bound, admissible optimal contracts might not exist. To solve for the optimal contract, we need to know the bound of the agent's information rent. However, this bound is endogenous to the contract; therefore, to derive this bound, we need to know the optimal contract.[13]

To overcome this challenge, I introduce the dual problem of (13), derive its properties, and show the dual problem offers an efficient and intuitive way to characterize the relaxed optimal contract.

[12] This function is implicitly defined as $\arg\min_\xi J^*(V_t, \xi, \phi_t)$, and equation (20) provides a formal statement.

[13] Information rents in contracting models tend to be bounded, and similar bounds are obtained in DeMarzo and Sannikov (2017), Di Tella and Sannikov (2021), and He et al. (2017).

3.2 THE DUAL PROBLEM

To define the dual formulation of (13), I introduce a new state variable that directly captures the principal's commitment to reduce the agent's risk exposure. As discussed in section 2.3, the principal has to commit to reduce the agent's future risk exposure in order to lower the agent's ex-ante information rent. In the dual problem, a positive multiplier, Y_t, replaces the information rent, ξ_t, as a state variable. The multiplier Y_t measures the principal's commitment at time t: The larger the multiplier Y_t, the more binding the commitment to reduce the agent's risk exposure, β_t.

For a given initial value of the multiplier at time 0, $Y_0 \geq 0$, let the multiplier evolve according to

$$dY_t = (r-\delta)Y_t dt + \phi_t \frac{\eta^2}{\lambda} C_t^\rho dt + \eta dU_t, \tag{16}$$

where $(U_t)_{t\geq 0}$ is a right-continuous, non-decreasing process with left limits and with $U_0 = 0$.

Then, consider the following problem:

$$\begin{aligned} G^*(V_0, Y_0, \phi_0) = \inf_{(C_t,\beta_t)_{t\geq 0}} \sup_{(U_t)_{t\geq 0}\in \mathcal{I}} \mathrm{E}\left[\int_0^\infty e^{-rt}\left\{\left(C_t - \eta\phi_t \frac{\beta_t}{\lambda}C_t^\rho + Y_t\beta_t\eta\phi_t(1-\phi_t)\right)dt - \beta_t dU_t\right\}\middle|\mathcal{F}_0\right] \\ \text{s.t.} \quad C_t \geq 0, \beta_t \geq 0 \\ (6), (16), \text{ and (3) with } m_t = 0\ \forall t \geq 0, \end{aligned} \tag{17}$$

where $\mathcal{I}$ is the set of right-continuous, non-decreasing processes with left limits and starting at 0. I call (17) the dual problem, as opposed to problem (13), which I call the primal problem, hereafter.

When looking for a solution to (17), I focus on admissible contracts for the dual problem, which are characterized by (14) and

$$\mathrm{E}\left[\int_0^\infty e^{-rt}|Y_t\beta_t\eta\phi_t(1-\phi_t)|\,dt\middle|\mathcal{F}_0\right] < \infty.$$

The dual problem in (17) has an intuitive appeal. First, consider the term $C_t - \eta\phi_t\frac{\beta_t}{\lambda}C_t^\rho$, which represents the flow cost of a principal who writes a contract with no information-rent problem. We can think of this hypothetical situation as one in which the hidden action m_t is observable but not contractible. In this case, the incentive-compatibility condition is the same as in (7) because shirking does not induce any belief distortion or any information rent for the agent.

Second, consider the term $Y_t\beta_t\eta\phi_t(1-\phi_t)$, which represents a flow cost of risk exposure β_t. In the dual formulation, this flow cost captures the principal's commitment to reduce the agent's risk exposure in order to lower ex-ante information rents. The larger $Y_t\eta\phi_t(1-\phi_t)$, the more binding the commitment of the principal at time t. In this sense, the multiplier Y_t measures the commitment of the principal to reduce risk exposures. Also note commitment becomes more binding when the speed of learning, $\eta\phi_t(1-\phi_t)$, increases. When returns are more informative, the information-rent distortions are more severe as (10) indicates; hence, the principal needs to further reduce risk exposures to honor her commitment to the information rent she initially promised.

Finally, the process $(U_t)_{t\geq 0}$ introduces possibly non-continuous changes in the multiplier and in the flow cost. This process serves the same function as the constraint on the non-negativity of capital, $\beta_t \geq \eta\xi_t$, in the primal problem. As I show in equation (A.3) in the proof of Theorem 3, a positive penalty $dU_t > 0$ materializes whenever $\beta_t < \eta\xi_t$, where ξ_t is the information rent implied by the dual contract at time t. Therefore, any contract $(C_t, \beta_t)_{t\geq 0}$ that is optimal for the dual problem will be such that $\beta_t \geq \eta\xi_t$ for all t, in order to avoid an unbounded penalty when this inequality is violated.

3.3 The Relation between the Dual and Primal Problem

The dual problem (17) offers a major simplification relative to the primal problem (13). To begin with, in the dual problem (17), a backward-looking state variable, the multiplier Y_t, replaces the forward-looking information rent ξ_t. Moreover, in the dual problem, the multiplier is subject to an exogenous lower bound, $Y_t \geq 0$, whereas the information rent in the primal problem is subject to an endogenous upper bound.

Importantly, as the following theorem shows, the primal and the dual problems are closely connected, and a relaxed optimal contract can be derived as a solution to the dual problem when Y_0 is initialized to zero.

Theorem 3 (Duality Relations)**.** *The following holds:*

(I) *$G^*(V, Y, \phi)$ is an increasing and concave function of Y. It is differentiable in Y for $Y \neq 0$.*

(II) *Let $Y_0 \geq 0$. Then, the dual cost function is related to the primal cost function by*

$$G^*(V_0, Y_0, \phi_0) = \inf_{\xi' \geq 0} \{J^*(V_0, \xi', \phi_0) + Y_0\xi'\}. \tag{18}$$

Any contract solving the dual problem (17) *delivers information rent $G^*_{Y+}(V_0, Y_0, \phi_0)$ to the agent.*[14] *Moreover, if a contract $\mathcal{C}$ solves the primal problem* (13) *with $\xi_0 = G^*_{Y+}(V_0, Y_0, \phi_0)$, then it also solves the dual problem* (17).

[14] I use $G^*_{Y+}(V, Y, \phi)$ to denote the right derivative of $G^*(V, Y, \phi)$ with respect to Y.

(III) *Let $\xi_0 \leq G^*_{Y^+}(V_0, 0, \phi_0)$. Then, the primal cost function is related to the dual cost function by*

$$J^*(V_0, \xi_0, \phi_0) = \sup_{Y' \geq 0} \{G^*(V,_0 Y', \phi_0) - Y'\xi_0\}. \tag{19}$$

If a contract $\mathcal{C}$ solves the dual problem (17) *with $Y_0 = -J^*_\xi(V_0, \xi_0, \phi_0)$, then it also solves the primal problem* (13)*.*

(IV) *A contract is a relaxed optimal contract given outside option V_0 and beliefs ϕ_0 if an only if it is the solution to the dual problem* (17) *with $Y_0 = 0$. The cost of a relaxed optimal contract for the principal is $G^*(V_0, 0, \phi_0)$ if the agent does not shirk.*

Whereas part (I) of Theorem 3 establishes properties of the dual cost function, parts (II) and (III) establish the crucial connections between the primal and the dual problem. Given an initial multiplier $Y_0 \geq 0$, an optimal contract for the dual problem can obtained as a solution to the primal problem after initializing $\xi_0 = G^*_{Y^+}(V_0, Y_0, \phi_0)$. Similarly, given an initial information rent $\xi_0 \leq G^*_{Y^+}(V_0, 0, \phi_0)$, an optimal contract for the primal problem can be obtained as a solution to the dual problem after initializing $Y_0 = -J^*_\xi(V_0, \xi_0, \phi_0)$. Part (IV) builds on parts (II) and (III) to conclude one can obtain the relaxed optimal contract as a solution to the dual problem when the multiplier is initialized to 0.

Moreover, based on Theorem 3, we obtain a dynamic bound on the information rent in a relaxed optimal contract. Let ξ_t be the information rent implied by a relaxed optimal contract at time t. Then,

$$\xi_t = G_{Y^+}(V_t, Y_t, \phi_t) \leq G_{Y^+}(V_t, 0, \phi_t) = \arg\min_{\xi'} J^*(V_t, \xi', \phi_t). \tag{20}$$

By Theorem 3(IV) and equation (16), in any relaxed optimal contract the multiplier starts at zero and remains positive thereafter, that is $Y_t \geq 0$ for all $t \geq 0$. Using the result in Theorem 3(II), we therefore have $\xi_t = G_{Y^+}(V_t, Y_t, \phi_t)$, where $G_{Y^+}(V_t, Y_t, \phi_t) \leq G_{Y^+}(V_t, 0, \phi_t)$ by concavity of $G^*(V, \cdot, \phi)$. Finally, combining parts (II) and (IV) of Theorem 3, we obtain that $G_{Y^+}(V_t, 0, \phi_t)$ is the initial information rent of any relaxed optimal contract for outside option V_t and beliefs ϕ_t. Such information rent, according to (15), is equal to $\arg\min_{\xi'} J^*(V_t, \xi', \phi_t)$.

Equation (20) provides a formal upper bound to the information rent in a relaxed optimal contract. Moreover, it offers a better intuition for why the dual problem (17) is numerically more tractable than the primal problem (13). In the dual problem, one can initialize Y_0 to a constant (zero). Then, the exogenous bound $Y_t \geq 0$ is mechanically satisfied given the law of motion of Y, (16). By contrast, in the primal problem, one should initialize the information rent to a value that is *a priori* not known. Moreover, the endogenous bound on the information rent is satisfied only when control variables are suitably chosen, causing standard numerical methods to fail.

3.4 Change of Variables and Incentives

Thanks to the homotheticity of the agent's utility function and the linearity of the technology, I simplify the problem with a change of variables. I define the *promised value* $v_t := ((1-\rho)V_t)^{\frac{1}{1-\rho}}$, which is the consumption equivalent of the agent's continuation value. I also consider a scaled version of the *multiplier*, $y_t := v_t^{-\rho} Y_t$. I can then express the dual cost function as[15]

$$G^*(V_0, Y_0, \phi_0) = v_0 g^*(y_0, \phi_0). \tag{21}$$

Throughout the paper, I impose restrictions on parameters to ensure that $g^*(0, \phi)$ is positive for all $\phi_0 \in [0, 1]$.[16]

I then define the agent's *incentives* $\hat{\beta}_t := \frac{\beta_t}{(1-\rho)V_t}$, *compensation ratio* $c_t := \frac{C_t}{v_t}$, and *capital ratio* $k_t := \frac{K_t}{v_t}$. Given a contract $(C_t, \beta_t)_{t\geq 0}$ that solves the dual problem (17), the laws of motion of the promised value v_t and multiplier y_t are thus given by

$$\frac{dv_t}{v_t} = \left(\frac{\delta}{1-\rho} - \frac{c_t^{1-\rho}}{1-\rho} + \frac{1}{2}\rho\hat{\beta}_t^2\right) dt + \hat{\beta}_t dW_t^{\mathrm{c}} \tag{22}$$

and

$$dy_t = \phi_t \frac{\eta^2}{\lambda} c_t^{\rho} dt + y_t \left(r - \frac{\delta}{1-\rho} + \rho\frac{c_t^{1-\rho}}{1-\rho} + \frac{1}{2}\rho\hat{\beta}_t^2\right) dt - y_t \rho \hat{\beta}_t dW_t^{\mathrm{c}}, \tag{23}$$

where I omitted the increasing process $(U_t)_{t\geq 0}$ because, as discussed earlier, $U_t = 0$ for all $t \geq 0$ for any contract that solves (17).

The quantity $\hat{\beta}_t$ offers an easy-to-interpret measure of the agent's incentives at time t. Because the promised value v_t is expressed in units of consumption,[17] $\hat{\beta}_t$ represents the percentage change of the agent's continuation value, measured in units of consumption, for a one-standard deviation return. In this formulation, the (scaled) multiplier y_t measures the principal's commitment to limit incentives $\hat{\beta}_t$. In fact, the principal's flow cost in the dual problem can be written as $v_t[c_t - \frac{\eta}{\lambda}\phi_t\hat{\beta}_t c_t^{\rho} + y_t\hat{\beta}_t\eta\phi_t(1-\phi_t)]$.

[15] I provide a formal proof in Lemma O.3 in the online appendix.

[16] Similar to Di Tella and Sannikov (2021), a sufficient condition for a positive dual cost function is (O.4) in Lemma O.6 in the online appendix. Note that, when $g^*(0, \phi_0)$ is positive, the agent's participation constraint binds. That is, the principal does not provide continuation value larger than the agent's outside option, V_0. If she did, costs would increase. When $g^*(0, \phi)$ is negative, an optimal contract does not exist, since the principal could obtain infinite profits by promising an infinite continuation value to the agent.

[17] The promised value v_t can be interpreted as the level of constant consumption from t to ∞ that would deliver continuation value V_t/δ to the agent.

Although compensation and capital are linear in the promised value because of the homogeneity and linearity assumptions, we should in general expect a positive relation between the compensation, capital, and promised value with a risk-averse agent. In fact, the principal increases current compensation when the agent's promised value is larger because of the agent's desire to smooth consumption over time. As for capital, the motivation relies on the agent's decreasing absolute risk aversion. With a larger promised value, the agent tolerates larger risk exposures. The principal thus optimally takes advantage of the agent's higher risk tolerance to increase capital under management while still satisfying (8).[18]

3.5 THE HJB EQUATION AND VERIFICATION

To conclude this section, I provide a verification theorem. Thanks to this theorem, I can solve problem (17) using a dynamic-programming approach.

Consider the following partial differential equation for $y \geq 0$ and $\phi \in [0,1]$:

$$
\begin{aligned}
rg(y,\phi) = \inf_{c \geq 0, \hat{\beta} \geq \eta g_{y+}(y,\phi)} \Bigg\{ & c - \eta\phi\hat{\beta}\frac{c^{\rho}}{\lambda} + y\hat{\beta}\eta\phi(1-\phi) + g(y,\phi)\left(\frac{\delta}{1-\rho} - \frac{c^{1-\rho}}{1-\rho} + \frac{1}{2}\rho\hat{\beta}^2\right) \\
& + g_y(y,\phi)\left[\phi\frac{\eta^2}{\lambda}c^{\rho} + y\left(r - \frac{\delta}{1-\rho} + \rho\frac{c^{1-\rho}}{1-\rho} + \frac{1}{2}\rho\hat{\beta}^2\right)\right] \\
& - \rho y\hat{\beta}^2 g_y(y,\phi) + \hat{\beta}\eta\phi(1-\phi)g_{\phi}(y,\phi) \\
& + \frac{1}{2}(y\rho\hat{\beta})^2 g_{yy}(y,\phi) + \frac{1}{2}\eta^2\phi^2(1-\phi)^2 g_{\phi\phi}(y,\phi) - \rho y\hat{\beta}\eta\phi(1-\phi)g_{y\phi}(y,\phi)\Bigg\}.
\end{aligned}
\tag{24}
$$

This is the HJB equation associated with the dual problem (17).

Note that instead of including a penalizing process $(U_t)_{t\geq 0}$, I introduce a lower bound on $\hat{\beta}$ related to the marginal cost of y. Unlike the limited-commitment model of Miao and Zhang (2015), here the increasing process $(U_t)_{t\geq 0}$ penalizes the principal for a choice of a control variable, namely, $\hat{\beta}$. In Miao and Zhang (2015), an increasing process penalizes the principal based on the value of the state variable. Whereas Miao and Zhang (2015) can characterize the optimal contract in terms of reflective barriers and a variation inequality, their approach is not valid in this situation. However, I show shortly that under standard regularity conditions, a solution to (24) provides a valid solution to the dual problem.

[18] An analogous result holds in the dynamic contracting models in Biais et al. (2010) and DeMarzo and Fishman (2007). These papers show investments (and disinvestments) at firm level depend on the agent's promised value. Similar to my model, when the agent has a larger promised value, the principal can more easily incentivize him to exert effort in a larger firm.

To provide an intuitive justification for the constraint on $\hat{\beta}$ in the HJB equation (24), consider two observations. First, as discussed above, in any contract $(C_t, \beta_t)_{t\geq 0}$ that solves the dual problem (17), we have $\beta_t \geq \eta\xi_t$ to avoid the penalty dU_t. Second, according to equation (20), in such a contract the agent's information rent ξ_t coincides with the marginal cost of commitment $G^*_{Y^+}(V_t, Y_t, \phi_t)$. Therefore, $\beta_t \geq \eta G^*_{Y^+}(V_t, Y_t, \phi_t)$. Because of (21), we have $\hat{\beta}_t \geq \eta g^*_{y^+}(y_t, \phi_t)$. Theorem 4 formalizes this intuition.

Starting from the HJB equation (24), I associate a contract $\mathcal{C}^R_{(v_0,y_0,\phi_0)} = (C^R_t, K^R_t)_{t\geq 0}$ with its policy functions, $c(y,\phi)$ and $\hat{\beta}(y,\phi)$, which are the minimizers in (24). Contract $\mathcal{C}^R_{(v_0,y_0,\phi_0)}$ is a contract for which $C^R_t := v_t c(y_t,\phi_t)$, $K^R_t := v_t k(y_t,\phi_t)$ with $k(y,\phi) := (\hat{\beta}(y,\phi) - \eta g_{y^+}(y,\phi))\frac{c(y,\phi)^\rho}{\lambda\sigma}$, and where v_t, y_t, and ϕ_t evolve as in (22), (23), and (3) for some initial v_0, y_0, and ϕ_0, and with $c_t = c(y_t,\phi_t)$, $\hat{\beta}_t = \hat{\beta}(y_t,\phi_t)$, and $m_t = 0$ for all $t \geq 0$. The following theorem verifies whether this contract is optimal.

THEOREM 4 (VERIFICATION). *Let $g\colon \mathbb{R}^+ \times [0,1] \to \mathbb{R}$ be a twice-differentiable solution of* (24)*, such that $g(y,\phi)$ is increasing and concave in y and $0 < j_1 \leq g \leq j_0$ for some constants j_1 and j_0. Let $\mathcal{C}^R_{(v_0,0,\phi_0)}$ be the contract generated by the policy functions of* (24) *with $y_0 = 0$ and $v_0 = ((1-\rho)V_0)^{\frac{1}{1-\rho}}$. Assume an $M > 0$ exists such that $|c(y,\phi)| + |\hat{\beta}(y,\phi)| < M$ for all $y \geq 0$ and $\phi \in [0,1]$. If $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is admissible, the following holds:*

(I) *The agent obtains lifetime utility $\frac{v_0^{1-\rho}}{1-\rho} = V_0$ from contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$ if he does not shirk.*

(II) *$\mathcal{C}^R_{(v_0,0,\phi_0)}$ is a relaxed optimal contract whose cost is $((1-\rho)V_0)^{\frac{1}{1-\rho}} g(0,\phi_0)$ if the agent does not shirk.*

(III) *Suppose the following condition is satisfied for all $y \geq 0$ and $\phi \in [0,1]$:*

$$[(1-\rho)g_{y^+}(y,\phi) - \rho y g_{yy}(y,\phi)]\hat{\beta}(y,\phi) + g_{y^+\phi}(y,\phi)\eta\phi(1-\phi) - \eta(1-2\phi)g_{y^+}(y,\phi) \geq 0. \quad (25)$$

Then $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is incentive compatible with no shirking, and hence, it is an optimal contract.

Thanks to parts (I) and (II) of Theorem 4, I can solve for a relaxed optimal contract by solving the HJB equation associated with the dual problem. Part (III) combines the results from Theorem 2 and equation (20) to establish a sufficient condition for the relaxed optimal contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$ to be incentive compatible and hence to be the optimal contract according to Definition 2. Equation (25) is equivalent to (12) when $\xi_t = v_t g_{y^+}(y_t,\phi_t)$, which is the case for the relaxed optimal contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$. In particular, the left-hand side of (25) coincides with the quantity $(\omega_t - \eta(1-2\phi_t)\xi_t)/v_t^{1-\rho}$.

In the next section, I rely on Theorem 4 to numerically solve for an optimal contract using the HJB equation (24). I then discuss the implications of the optimal contract for the

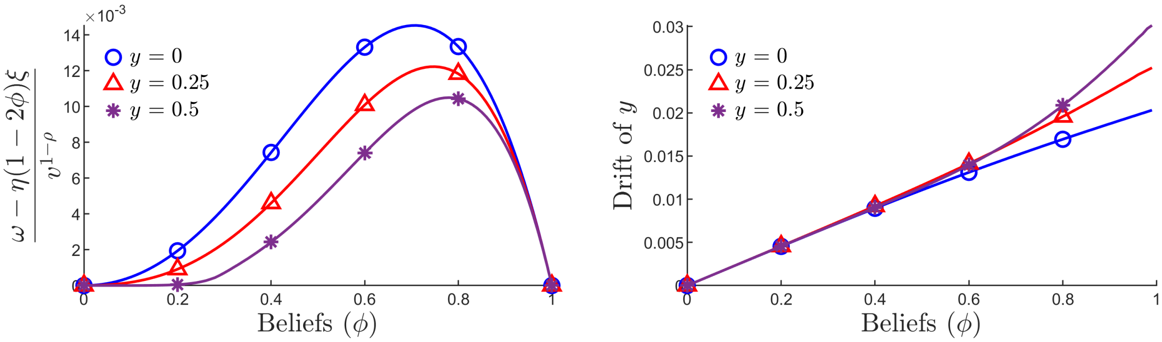

(a) Sufficient condition

(b) Drift of y

Figure 1: Verification of condition (25) and drift of the multiplier y in the optimal contract. Figure (a) plots the left-hand side of equation (25) as a function of the multiplier y and beliefs ϕ. Figure (b) plots the drift of equation (23) as a function of the multiplier y and beliefs ϕ. Results are shown for three values of the multiplier y. The parameters are $r = 0.02$, $\delta = 0.05$, $\rho = 1/3$, $\lambda = 0.95$, $\sigma\eta = 0.02$, and $\sigma = 0.18$.

dynamics of the agent's incentives.

4 Results and Dynamic Trade-Offs

I numerically solve the HJB equation (24), and I verify the numerical solution satisfies the conditions of Theorem 4. In particular, the numerical solution is twice differentiable, increasing and concave in y, and bounded between two positive constants (Figure O.1(a) in the online appendix provides an illustration.) Moreover, it satisfies (25). As an illustration, Figure 1(a) plots the left-hand side of (25) for three values of the multiplier y, showing (25) is satisfied for these values. The same result holds on the entire state space. See Figure O.1(b) in the online appendix for an illustration.

Thanks to the dual formulation of the contract-design problem, I can characterize the dynamics of the principal's commitment through the dynamics of the multiplier y. Key features of the optimal contract, such as the long-run path of incentives, depend on the dynamics of the principal's commitment.

In equation (23), we immediately observe y_t has a negative loading on the shocks $dW_t^{\mathfrak{C}}$, indicating the principal relaxes (reinforces) her commitment after good (bad) performance. After good performance, the agent is rewarded with higher promised value v_t. The principal thus optimally relaxes her commitment in order to be more "unconstrained" in the future, when larger compensation is due to the agent. As a result, the principal optimally establishes a negative correlation between changes in her commitment and changes in the agent's promised value.

Moreover, on average, the principal tightens her commitment over time. In Figure 1(b) and Figure O.1(c) in the online appendix, we observe the drift of y_t is positive, indicating the multiplier y tends to increase over time. This feature is common in dynamic contracting models in which future risk exposures distort the agent's current incentives (Di Tella and Sannikov, 2021; He et al., 2017; Prat and Jovanovic, 2014).

Next, I study how the agent's incentives, $\hat{\beta}(y_t, \phi_t)$, change with the agent performance. To study the relation between performance and compensation, I define the pay-performance sensitivity as the percentage change in compensation for a 1% return. Because in the optimal contract, compensation takes the form $C_t^R = v_t c(y_t, \phi_t)$, the pay-performance sensitivity can be written as

$$\epsilon_C(y_t, \phi_t) := \frac{dC_t^R / C_t^R}{dR_t} = \frac{1}{\sigma} \left(\hat{\beta}(y_t, \phi_t) + \frac{\sigma_c(y_t, \phi_t)}{c(y_t, \phi_t)} \right), \tag{26}$$

where $\sigma_c(y_t, \phi_t)$ is the volatility of $c(y_t, \phi_t)$. Because $\epsilon_C(y, \phi)$ directly depends on the agent's incentives $\hat{\beta}(y, \phi)$, the pay-performance sensitivity reflects the strength of the agent's incentives.

4.1 Performance and Incentives in the Optimal Contract

I now study changes in the agent's incentives, $\hat{\beta}(y_t, \phi_t)$, triggered by the agent's performance. In the optimal contract, incentives are functions of two state variables: beliefs ϕ_t and the multiplier y_t. Figure 2(a) shows how incentives vary for all values of beliefs ϕ and for three values of the multiplier y. Figure O.1(d) in the online appendix shows incentives for a continuum of values for the multiplier.

We immediately notice incentives increase with beliefs and decline with the multiplier. This result is robust across all the numerical solutions that I obtained under a variety of parameters. Moreover, by comparing Figures 2(a) and 2(b), we also observe the capital ratio $k(y, \phi)$ follows a similar pattern.

To explain these results, recall that in contract $\mathcal{C}^R_{(v_0, 0, \phi_t)}$, capital under management is given by $K_t^R = v_t k(y_t, \phi_t)$, where

$$k(y, \phi) = \frac{\hat{\beta}(y, \phi) - \eta g_{y^+}(y, \phi)}{\lambda \sigma}. \tag{27}$$

This expression corresponds to the incentive-compatibility condition (8), holding as an equality in the optimal contract $\mathcal{C}^R_{(v_0, 0, \phi_0)}$, after using the results in Theorems 3 and 4. Incentives, $\hat{\beta}(y, \phi)$, and the capital ratio, $k(y, \phi)$, are intimately connected by this incentive-compatibility condition. In fact, the primary purpose of giving incentives to the agent is to delegate capital while preventing shirking.

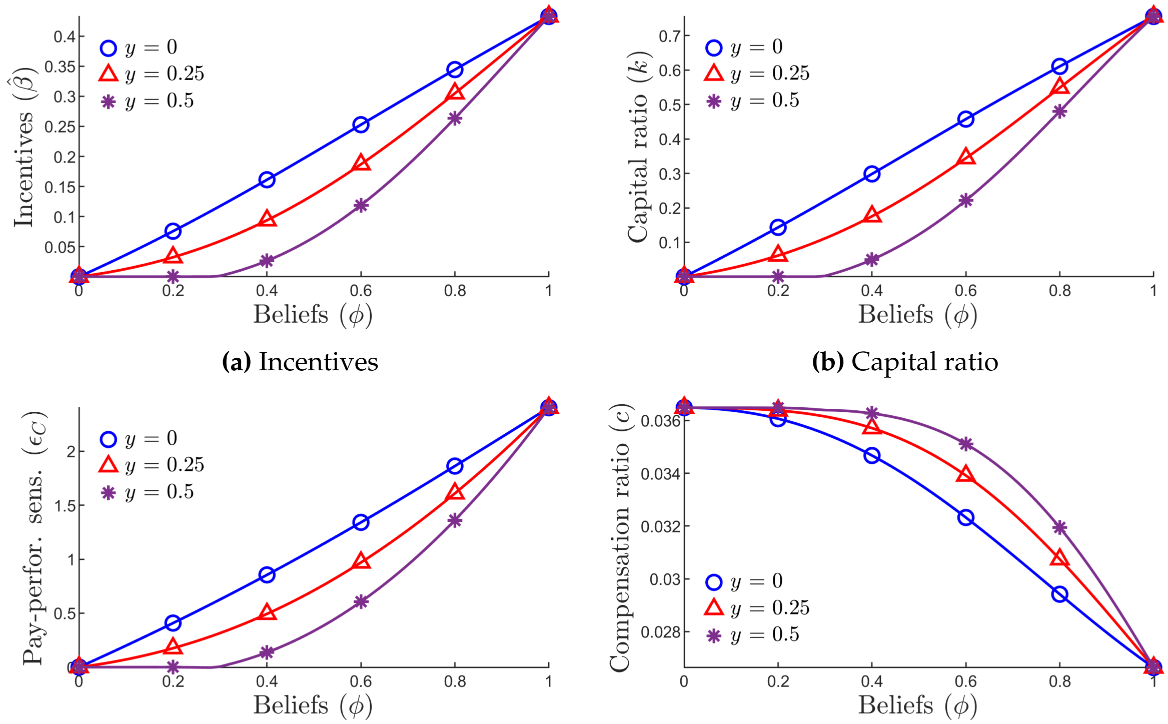

(a) Incentives **(b)** Capital ratio

(c) Pay-performance sensitivity **(d)** Compensation ratio

Figure 2: Incentives, $\hat{\beta}(y, \phi)$, capital ratio, $k(y, \phi)$, pay-performance sensitivity, $\varepsilon_C(y, \phi)$, and compensation ratio, $c(y, \phi)$, in the optimal contract as functions of the multiplier y and beliefs ϕ. Results are shown for three values of the multiplier y. The parameters are $r = 0.02$, $\delta = 0.05$, $\rho = 1/3$, $\lambda = 0.95$, $\sigma\eta = 0.02$, and $\sigma = 0.18$.

To study the trade-offs the principal faces when choosing incentives $\hat{\beta}(y, \phi)$, consider the first-order condition for incentives in (24) for an interior minimizer:

$$\underbrace{\lambda^{-1}\eta\phi c(y,\phi)^{\rho}}_{\text{A: Benefit from IC}} \underbrace{-\eta\phi(1-\phi)g_{\phi}(y,\phi)}_{\text{B: Benefit from Cov}(V,\phi)} = \underbrace{y\eta\phi(1-\phi)}_{\text{C: Cost from commitment}} + \underbrace{\rho\hat{\beta}(y,\phi)\left[g(y,\phi) - yg_y(y,\phi) + y^2\rho g_{yy}(y,\phi)\right] - \rho y\eta\phi(1-\phi)g_{y\phi}(y,\phi)}_{\text{D: Cost from risk aversion}}.$$

On the one hand, the principal benefits from higher incentives because, by the incentive-compatibility condition (27), they allow more capital to be delegated to the agent, thus increasing current expected cash flows. Term A captures this marginal benefit. Moreover, with higher incentives, the principal increases the covariance between promised value and beliefs, thus promising larger (lower) compensation when the agent's expected productivity is higher (lower.) The marginal benefit of such covariance is measured by term B.

On the other hand, the principal suffers costs from higher incentives because of her commitment and because of the agent's risk aversion. In the optimal contract, the principal commits to reduce the agent's future incentives in order to lower ex-ante information rents. In the dual formulation, the commitment of the principal is captured by the multiplier y, which imposes an ex-post marginal penalty for incentives equal to term C. Moreover, the risk-averse agent requires larger future compensation when exposed to more risk, generating the marginal cost of incentives in by term D. Note this term is zero if the agent is risk neutral and $\rho = 0$.

The principal sets the agent's incentives after optimally trading off marginal benefits and marginal costs. When beliefs increase (decline), expected productivity increases (decline), along with the marginal benefit of investment delegation. When the multiplier declines (increases), the commitment of the principal is less (more) binding, and the ex-post penalty for incentives is smaller (larger.) Hence, incentives $\hat{\beta}$ increase with beliefs and decline with the multiplier.

As a result, incentives increase after good performance and decline after bad performance. After good performance, the expected productivity of the agent increases and the principal relaxes her commitment. The principal thus increases the agent's incentives in order to take advantage of higher expected returns and lower commitment. I summarize this observation in the following prediction.

PREDICTION. *Incentives increase (decline) after good (bad) performance; that is,*

$$\frac{d\hat{\beta}_t}{dR_t} \geq 0.$$

For this prediction, both learning and moral hazard play a key role, even in the absence of commitment. Without moral hazard, the principal would fully insure the agent, and incentives would be zero. With moral hazard but no learning, the agent's productivity would be known and constant. The agent's incentives would be increasing in the agent's productivity, but they would be constant over time.[19] With learning and moral hazard, incentives change with performance. Because of learning, the agent's expected productivity increases after good performance, and the principal thus wants to allocate more capital $k(y_t, \phi_t)$ to the agent. Because of moral hazard and the incentive compatibility condition (27), the principal has to increase the agent's incentives, $\hat{\beta}(y, \phi)$, along with capital in order to preventing shirking.

[19]See the discussion in section S.2 of the supplemental material. Even in the absence of learning, the principal wants to allocate more capital to more productive agents. To prevent shirking, more productive agents are thus offered steeper incentive contracts.

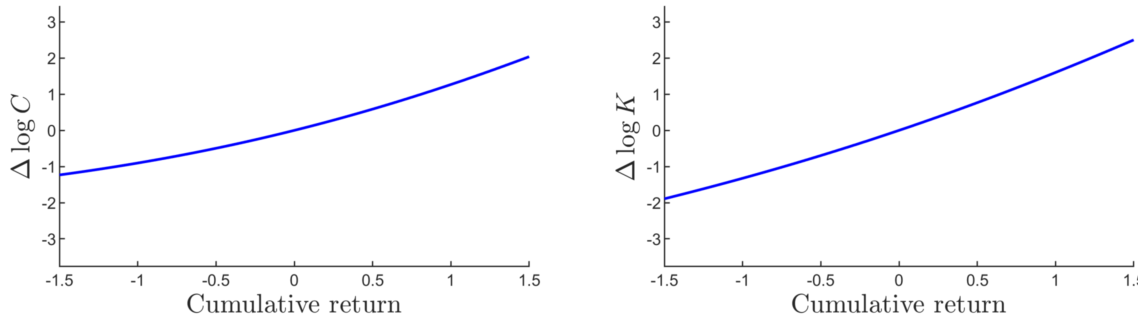

(a) Performance and compensation

(b) Performance and capital

Figure 3: Relation between cumulative performance, change in compensation, and change in capital under management. The curves represent the change in log-compensation (Figure (a)) and log-capital (Figure (b)) as functions of cumulative performance. Curves are shifted to represent changes relative to an agent that has a zero cumulative performance. Performance, change in compensation, and change in capital are computed while assuming that returns are realized uniformly over time during the course of one year. Figures are drawn for initial beliefs $\phi_0 = 0.5$ and initial multiplier $y_0 = 0$. The parameters are $r = 0.02$, $\delta = 0.05$, $\rho = 1/3$, $\lambda = 0.95$, $\sigma\eta = 0.02$, and $\sigma = 0.18$.

Convex and Back-Loaded Compensation. As Figure 2(c) illustrates, the pay-performance sensitivity $\varepsilon_C(y, \phi)$, defined in (26), increases with beliefs ϕ and declines with the multiplier y, similar to incentives $\hat{\beta}(y, \phi)$. Consequently, agents with better (worse) past performance face steeper (flatter) performance-based compensation schemes.

Because of this mechanism, the change in log-compensation appears convex in cumulative performance, as Figure 3(a) illustrates. As the agent keeps performing well, incentives become increasingly steeper and compensation becomes increasingly more sensitive to performance, thus generating a convex compensation scheme for the agent.

When beliefs converge to either zero or one, the agent's compensation scheme converges to a linear one, where log-compensation increases linearly with performance. Figure O.2 in the online appendix provides an illustration. As beliefs converge to zero or one, their volatility, $\eta\phi_t(1 - \phi_t)$, converges to zero. Hence, players stop updating their beliefs based on performance and the marginal cost of commitment, $y_t\eta\phi_t(1 - \phi_t)$, vanishes. This situation resembles a model with moral hazard but without learning, such as the model in section S.2 of the supplemental material, where the pay-performance sensitivity is constant.

However, even in the limit with no learning, the steepness of the compensation scheme still depends on the agent's perceived productivity. With an optimal contract for moral hazard, the principal offers steep incentives to an agent with high perceived productivity in order to delegate a substantial amount of capital while preventing shirking. Agents with low perceived productivity, on the other hand, face flat compensation schemes because the principal delegates less capital to them and provides more insurance.

After good performance, the principal also back-loads the performance-based compensation of the agent. Figure 2(d) shows the compensation ratio $c(y, \phi)$ declines with beliefs and increases with the multiplier. Hence, the compensation ratio declines after good performance, thus indicating the principal increases future promised consumption more than current compensation.

To understand why the principal back-loads the agent's performance-based compensation, consider the dynamic trade-off captured by the first-order condition for the compensation ratio $c(y, \phi)$:

$$\underbrace{1}_{\text{A: direct cost}} = \underbrace{\lambda^{-1}\eta\phi(\hat{\beta}(y,\phi) - \eta g_y(y,\phi))\rho c(y,\phi)^{\rho-1}}_{\text{B: benefit from IC}} + \underbrace{(g(y,\phi) - \rho y g_y(y,\phi))c(y,\phi)^{-\rho}}_{\text{C: benefit from lower } V}.$$

Compensation is a direct cost for the principal, with unit marginal cost (term A). However, compensation is also beneficial for the principal because, first, it relaxes the incentive-compatibility constraint, thus allowing more investment delegation (term B); second, it allows the agent's continuation value to be reduced, thus reducing future compensation costs (term C).

When expected productivity increases or the principal relaxes her commitment, the trade-off tilts in favor of a lower $c(y, \phi)$. With higher expected productivity and weaker commitment, the principal enjoys a larger marginal benefit of relaxing the incentive-compatibility constraint in term B. However, she also can deliver any given continuation value more cheaply, thus lowering the marginal benefit in term C. Because $\rho < 1/2$, the second effect dominates the first, and the compensation ratio declines.

CAPITAL FLOWS. Figure 3(b) shows that, similar to compensation, capital under management increases with performance. In the optimal contract, capital under management, $K_t = v_t k(y_t, \phi_t)$, increases with the agent's promised value and with the capital ratio. As discussed above, the capital ratio $k(y_t, \phi_t)$ increases with past performance because the principal wants to delegate more capital when the agent is perceived to be more productive and when commitment is less binding. The promised value increases with performance in order to deter shirking. Therefore, the total amount of capital under management also increases in performance.

INCENTIVES IN THE LONG RUN Over time, the principal tends to be increasingly committed to reduce the agent's incentives. This increasing commitment is reflected in a positive drift of the multiplier y. Because incentives $\hat{\beta}(y, \phi)$ decline when the multiplier y increases, incentives tend to decline over time.

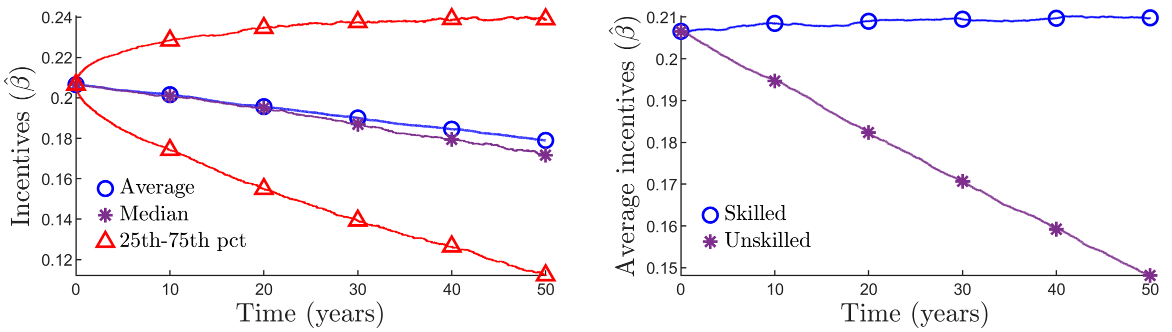

(a) Incentives over time

(b) Average incentives by type

Figure 4: Incentives, $\hat{\beta}(y_t, \phi_t)$, over time for $\phi_0 = 0.5$. Figure (a) shows the unconditional distribution of incentives at each point in time. Figure (b) shows the average incentives conditional on the agent's type. The distributions are obtained from a sample of 10,000 independent simulations in which the fraction of skilled agents is equal to the prior ϕ_0. The parameters are $r = 0.02$, $\delta = 0.05$, $\rho = 1/3$, $\lambda = 0.95$, $\sigma\eta = 0.02$, and $\sigma = 0.18$.

Figure 4(a) shows the unconditional distribution of incentives over time for the entire sample. Here, we observe incentives tend to decline over time. To study the distribution of incentives over time, I generate 10,000 simulated histories of 50 years. In the simulations, the initial prior is $\phi_0 = 0.5$. Half of the paths are simulated using $h = 1$, whereas the remaining half is simulated using $h = 0$. Figure O.3 in the online appendix plots the distribution of incentives over time for different priors (and hence different shares of skilled agents in the 10,000 simulations), namely, $\phi_0 = 0.1$ and $\phi_0 = 0.9$.

One should expect a declining path of incentives when future risk exposures exacerbate current moral-hazard frictions. The connection between future risk exposures and current frictions may be introduced by the agent's information rent in models with learning, as in my model, He et al. (2017), and Prat and Jovanovic (2014), or by agent's precautionary saving motive in models with private savings, as in Di Tella and Sannikov (2021). Similar to Di Tella and Sannikov (2021), in my model the principal also restricts future access to capital to mitigate ex-ante frictions. In fact, both here and in Di Tella and Sannikov (2021), capital and risk exposures are crucially connected by a necessary incentive-compatibility condition.

Although incentives decline on average, agents experience different paths of incentives depending on their history of performance and their hidden type. Figure 4(b) shows average incentives for the two types of agents starting from prior $\phi_0 = 0.5$. In the figure, expected incentives increase slowly over time for skilled agents, and decline steeply for unskilled agents. Suppose an agent is skilled. By Girsanov's theorem, beliefs acquire a positive drifts because of the agent's high expected performance. For analogous reasons, the multiplier acquires a negative term in its drift. Thus, if players' beliefs are pessimistic enough, the skilled agent's superior performance might offset the principal's increasing commitment , causing an upward trend in incentives to be observed.[20] By contrast, unskilled agents experience a declining path of expected incentives because of the combined effect of poor performance and increasing commitment.

Full commitment by the principal is crucial for a declining path of incentives to be observed. Ex post, the principal is tempted to renegotiate the contract and relax her commitment by "resetting" y to zero. In the supplemental material, I study contracts in which the principal cannot fully commit, as in the "commitment and renegotiation" model of Laffont and Tirole (1990). In this case, no declining path of expected incentives is observed.

5 CONCLUSIONS

In this paper, I study how a principal delegates the management of capital to an agent with uncertain productivity. The principal dynamically learns about the agent's productivity and allocates her commitment based on past performance. As a result, the agent's incentives and capital under management also change in response to past performance.

The model highlights the dynamic nature of incentives, compensation, and capital under management. In particular, the model predicts that, after a history of good (bad) performance, the agent's incentives increase (decline.) The agent's incentives are reflected in his pay-performance sensitivity. As the pay-performance sensitivity increases (declines) after a history of good (bad) performance, the agent's compensation scheme appears convex in cumulative performance.

I also make a methodological contribution to overcome the mathematical and computational challenges of the model. These challenges regularly appear in continuous-time contracting models with state variables. Such models often involve an endogenous boundary in the state space of the optimal contract. Thanks to the dual formulation of the contract-design problem, I sidestep the computation of the endogenous boundary. This methodology can be readily applied to a large class of continuous-time contracting models with state variables.

[20] Figure O.3(d) in the online appendix shows that, when $\phi_0 = 0.9$, incentives tend to decline even conditional on $h = 1$. When the principal already expects large returns because of optimistic beliefs, even a skilled agent may fail to produce good enough returns to offset the principal's increasing commitment.

A Proofs

A.A Proof of Lemma 1

I proceed by contradiction. Assume $\mathcal{C}$ is an optimal contract and that it is incentive compatible with a shirking process $(m_t)_{t\geq 0}$ such that $\mathrm{E}\left[\int_0^\infty e^{-rt} m_t K_t \, dt\right] > 0$.

Because the contract is $\mathcal{F}_t$-adapted, the time-t allocation (C_t, K_t) and shirking action m_t in equilibrium can be expressed as functions of the history of returns up to time t, $(R_s)_{0\leq s\leq t}$. That is, $C_t = \mathbf{C}^t((R_s)_{0\leq s\leq t})$, $K_t = \boldsymbol{K}^t((R_s)_{0\leq s\leq t})$, and $m_t = \mathbf{m}^t((R_s)_{0\leq s\leq t})$, for some functions $\mathbf{C}^t$, $\boldsymbol{K}^t$, and $\mathbf{m}^t$.

Now consider an alternative contract $\hat{\mathcal{C}} = (\hat{C}_t, \hat{K}_t)_{t\geq 0}$. This contract is designed in the following way. If the history of returns up to time t is $(R_s)_{0\leq s\leq t}$, $\hat{\mathcal{C}}$ specifies capital at time t equal to

$$\hat{K}_t := \boldsymbol{K}^t\left(\left(R_s - \int_0^s \hat{m}_i \, di\right)_{0\leq s\leq t}\right), \quad \text{with } \hat{m}_t := \mathbf{m}^t\left(\left(R_s - \int_0^s \hat{m}_i \, di\right)_{0\leq s\leq t}\right).$$

Contract $\hat{\mathcal{C}}$ then specifies consumption as

$$\hat{C}_t = \mathbf{C}^t\left(\left(R_s - \int_0^s \hat{m}_i \, di\right)_{0\leq s\leq t}\right) + \lambda \hat{m}_t \hat{K}_t.$$

Suppose the agent's best response to contract $\hat{\mathcal{C}}$ is the shirking process $(m'_t)_{t\geq 0}$. After any history up to time t, with contract $\hat{\mathcal{C}}$, the agent receives capital equal to the amount he would have received if he had chosen a shirking strategy $(m_s + m'_s)_{0\leq s\leq t}$ under the original contract $\mathcal{C}$. Similarly, he receives compensation equal to the combined amount of compensation and private benefits he would have enjoyed if he had chosen a shirking strategy $(m_s + m'_s)_{0\leq s\leq t}$ under the original contract $\mathcal{C}$. In particular,

$$\mathrm{E}^{(m'_s)_{0\leq s\leq t}}[u(\hat{C}_t)|\mathcal{F}_0] = \mathrm{E}^{(m_s+m'_s)_{0\leq s\leq t}}[u(C_t + \lambda m_t K_t)|\mathcal{F}_0],$$

where I explicitly express the dependence of the probability measure of returns on the shirking process.

The best response of the agent to contract $\hat{\mathcal{C}}$ thus solves the following problem:

$$\max_{(m'_t)_{t\geq 0}} \mathrm{E}^{(m'_t)_{t\geq 0}}\left[\int_0^\infty e^{-\delta t} u(\hat{C}_t)\, dt \middle| \mathcal{F}_0\right] = \max_{(m'_t)_{t\geq 0}} \mathrm{E}^{(m_t+m'_t)_{t\geq 0}}\left[\int_0^\infty e^{-\delta t} u(C_t + \lambda m_t K_t)\, dt \middle| \mathcal{F}_0\right].$$

Because $(m_t)_{t\geq 0}$ is the best response of the agent to contract $\mathcal{C}$, the right-hand side is maximized by $m'_t = 0$ for all $t \geq 0$. Hence, no shirking is the agent's best response to contract $\hat{\mathcal{C}}$.

Note that with the new contract $\hat{\mathcal{C}}$, the agent receives the same lifetime utility as in contract $\mathcal{C}$. However, the costs for the principal change by $\mathrm{E}\left[\int_0^\infty e^{-rt}(\lambda - 1) m_t K_t \, dt\right] < 0$, because she pays $\lambda m_t K_t$ as an additional compensation at time t, but enjoys an additional cash flow $m_t K_t$ by preventing shirking. Because the principal bears lower costs with contract $\hat{\mathcal{C}}$, the assumption that $\mathcal{C}$ is an optimal contract is contradicted. □

A.B Proof of Theorem 1

I use the stochastic maximum principle (Yong and Zhou, 1999, Chapter 3 and references therein) to derive necessary conditions for incentive compatibility, as in Williams (2011) and DeMarzo and Sannikov (2017).

Instead of working with posterior beliefs ϕ_t as a state variable, it is convenient to work with the log-likelihood ratio

$$x_t := \frac{\eta}{\sigma}\int_0^t (dR_s + m_s ds) - \frac{1}{2}\eta^2 t.$$

By Girsanov's theorem, e^{x_t} represents the ratio between the likelihood that the path $(R_s)_{0\le s\le t}$ is generated by a skilled agent ($h = 1$) and the likelihood that the same path is generated by an unskilled agent ($h = 0$.) That is, $e^{x_t} = \frac{\mathrm{E}[h|\mathcal{F}_t]}{1-\mathrm{E}[h|\mathcal{F}_t]}$.

We can then express beliefs as a function of the log-likelihood ratio,

$$\phi_t = \psi(x_t) := \frac{\phi_0 e^{x_t}}{1-\phi_0+\phi_0 e^{x_t}}.$$

Let $\mathcal{C}$ be an incentive-compatible contract with no shirking, and let $P^{\mathcal{C}}$ be the probability measure for which $W_t^{\mathcal{C}}$ is a standard Brownian motion. Consider an arbitrary shirking process $(m_t)_{t\ge 0}$. Let x_t be the principal's log-likelihood ratio and let $x_t + \Delta_t^x$ be the agent's log-likelihood ratio. The laws of motion of x_t and Δ_t^x are given by

$$\begin{aligned} dx_t &= \left(\psi(x_t) - \frac{1}{2}\right)\eta^2 dt + \eta dW_t^{\mathcal{C}} \\ d\Delta_t^x &= \frac{m_t}{\sigma}\eta dt. \end{aligned}$$

Let P^m be a measure for which $W_t^m := W_t^{\mathcal{C}} - \int_0^t \frac{m_s}{\sigma}\, ds$ is a standard Brownian motion. The continuation value of the agent, given $\mathcal{C}$ and $(m_t)_{t\ge 0}$, can be written as

$$\mathrm{E}\left[\int_0^\infty \Gamma_t^m e^{-\delta t} u(C_t + m_t\lambda K_t)\, dt \Big| \mathcal{F}_0\right],$$

where $\Gamma^m := \frac{dP^m}{dP^{\mathcal{C}}}$ is a density process representing the change of measure for the path of returns induced by the shirking strategy $(m_t)_{t\ge 0}$. By Girsanov's theorem, Γ^m evolves as

$$d\Gamma_t^m = \left(-\frac{m_t}{\sigma} + \eta(\psi(x_t+\Delta_t^x) - \psi(x_t))\right)\Gamma_t dW_t^{\mathcal{C}}.$$

Consider the following Hamiltonian,

$$H(\Gamma, x, \Delta^x; m; \mathcal{B}, \xi) := \Gamma u(C + m\lambda K) + \left(-\frac{m}{\sigma} + \eta(\psi(x+\Delta^x) - \psi(x))\right)\Gamma\mathcal{B} + \frac{m}{\sigma}\eta\xi,$$

If $m_t = 0$ is optimal for all t, the stochastic maximum principle (Williams, 2011; Yong and Zhou, 1999) implies that $m = 0$ must maximize the Hamiltonian above with $\Gamma^m = 1$, $\Delta x = 0$, and with $(\mathcal{V}_t, \mathcal{B}_t)_{t\geq 0}$ and $(\xi_t, \omega_t)_{t\geq 0}$ solving the following backward stochastic differential equations (BSDEs):

$$d\mathcal{V}_t = (\delta\mathcal{V}_t - H_\Gamma(1, x_t, 0; 0; \mathcal{B}_t, \xi_t))dt + \mathcal{B}_t dW_t^{\mathcal{C}}$$
$$d\xi_t = (\delta\xi_t - H_x(1, x_t, 0; 0; \mathcal{B}_t, \xi_t))dt + \omega_t dW_t^{\mathcal{C}},$$

with limit conditions $\lim_{T\to\infty} E[e^{-\delta T}\mathcal{V}_T|\mathcal{F}_0] = 0$ and $\lim_{T\to\infty} E[e^{-\delta T}\xi_T|\mathcal{F}_0] = 0$ (Konlack Socgnia and Menoukeu-Pamen, 2015).

We must therefore have

$$u'(C_t)\lambda K_t \leq \frac{\mathcal{B}_t}{\sigma} - \frac{\eta\xi_t}{\sigma}$$
$$d\mathcal{V}_t = (\delta\mathcal{V}_t - u(C_t))dt + \mathcal{B}_t dW_t^{\mathcal{C}}$$
$$d\xi_t = (\delta\xi_t - \mathcal{B}_t\eta\phi_t(1-\phi_t))dt + \omega_t dW_t^{\mathcal{C}},$$

Solving the BDSE for $\mathcal{V}_t$, we obtain

$$\mathcal{V}_t = \mathrm{E}\left[\int_t^\infty e^{-\delta(s-t)}u(C_s)\,ds\middle|\mathcal{F}_t\right],$$

Hence, $\mathcal{V}_t = V_t$ and $\mathcal{B}_t = \beta_t$, which concludes the proof. □

A.C Proof of Theorem 2

Define $\zeta_t := \frac{\xi_t}{\phi_1(1-\phi_t)}$, which, by Ito's lemma, evolves as

$$d\zeta_t = \left(\delta\zeta_t - \eta\beta_t + \eta^2\phi_t(1-\phi_t)\zeta_t - \eta(1-2\phi_t)\omega_{\zeta t}\right)dt + \omega_{\zeta t} dW_t^{\mathcal{C}},$$

where

$$\omega_{\zeta t} := \frac{\omega_t - \xi_t(1-2\phi_t)\eta}{\phi_t(1-\phi_t)}.$$

Let ϕ_t^A be the agent's posterior at time t after a deviation $(m_s)_{0\leq s\leq t}$, while ϕ_t is the principal's. Because $m_s \geq 0$, $\phi_t^A \geq \phi_t$. I want to show that if the conditions of the theorem are satisfied, then $V_t + (\phi_t^A - \phi_t)\zeta_t$ is an upper bound on the agent's future expected utility at time t. Because $\phi_0^A = \phi_0$, this upper bound proves the agent has no (strictly) better strategy than choosing $m_t = 0$ for all $t \geq 0$ and receive the continuation value V_0.

Consider an arbitrary deviation up to time t and let

$$S_t := \int_0^t e^{-\delta s}u(C_s + m_s\lambda K_s)\,ds + e^{-\delta t}\left(V_t + (\phi_t^A - \phi_t)\zeta_t\right).$$

It suffices to show S_t is a supermartingale because in this case,

$$S_t \geq \mathrm{E}[S_T|\mathcal{F}_t] = \mathrm{E}\left[\int_0^T e^{-\delta s}u(C_s + m_s\lambda K_s)\,ds + e^{-\delta T}(V_T + (\phi_T^A - \phi_T)\zeta_T)\middle|\mathcal{F}_t\right],$$

which would then imply

$$V_t + (\phi_t^A - \phi_t)\zeta_t \geq \mathrm{E}\left[\int_t^\infty e^{-\delta(s-t)}u(C_s + m_s\lambda K_s)\,ds\middle|\mathcal{F}_t\right] \quad \forall t \geq 0,$$

because V_u, $\phi_u^A - \phi_u$, and ζ_u are positive.

To show S_t is a supermartingale, it is sufficient to prove the drift of dS_t is non-positive. Using Ito's lemma, and after some algebra, the drift of dS_t simplifies to

$$e^{-\delta t}\left[u(C_t + m_t\lambda K_t) - u(C_t) - \beta_t\frac{m_t}{\sigma} + \eta\phi_t(1-\phi_t)\zeta_t\frac{m_t}{\sigma} - (\phi_t^A - \phi_t)\omega_{\zeta t}\frac{m_t}{\sigma}\right].$$

Recall $\xi_t = \phi_t(1-\phi_t)\zeta_t$. By (8), (12), and the concavity of $U(\cdot)$,

$$u'(C_t + m_t\lambda K_t) - \frac{\beta_t - \eta\phi_t(1-\phi_t)\zeta_t}{\sigma} - (\phi_t^A - \phi_t)\frac{\omega_{\zeta t}}{\sigma} \leq 0 \quad \forall m_t \geq 0$$

Hence, the drift is maximized for $m_t = 0$. For $m_t = 0$, the drift of dS_t is zero, and it is negative for $m_t > 0$. Therefore, for an arbitrary deviation $(m_t)_{t\geq 0}$, S_t is a supermartingale, thus concluding the proof. □

A.D Proof of Theorem 3

Notation. Before proceeding to the proof, I introduce some useful notation. I use $\mathcal{P}_{(V,\xi,\phi)}$ to denote the set of admissible contracts that solve the primal problem (13) for $(V_0, \xi_0, \phi_0) = (V, \xi, \phi)$. I call such contracts *optimal primal contracts*. I use $\mathcal{D}_{(V,Y,\phi)}$ to denote the set of admissible contracts that solve the dual problem (17) for $(V_0, Y_0, \phi_0) = (V, Y, \phi)$. I call such contracts *optimal dual contracts*.

Let $\mathcal{C} = (C_t, \beta_t)_{t\geq 0}$ be an arbitrary admissible contract that delivers expected lifetime utility V(if the agent does not shirk) and information rent ξ. I use $J(V, \xi, \phi|\mathcal{C})$ to denote the value of the objective function of primal problem (13) when the principal chooses contract $\mathcal{C}$. Similarly, I use $G(V, Y, \phi|\mathcal{C})$ to denote value of the objective function of the dual problem (17) when the principal chooses contract $\mathcal{C}$. In particular, if $\mathcal{C}^P \in \mathcal{P}_{(V,\xi,\phi)}$ and $\mathcal{C}^D \in \mathcal{D}_{(V,Y,\phi)}$, we must have $J^*(V, \xi, \phi) = J\big(V, \xi, \phi|\mathcal{C}^P\big)$ and $G^*(V, Y, \phi) = G\big(V, Y, \phi|\mathcal{C}^D\big)$.

Given a contract $\mathcal{C} = (C_t, \beta_t)_{t\geq 0}$, I denote with $\tilde{\xi}(\mathcal{C}, \phi)$ the information rent implied by the contract $\mathcal{C}$ when beliefs are ϕ; that is,

$$\tilde{\xi}(\mathcal{C}, \phi) = \mathrm{E}\left[\int_0^\infty e^{-\delta t}\beta_t\eta\phi_t(1-\phi_t)\,dt\middle|\mathcal{F}_0\right], \quad \text{with } \phi_0 = \phi.$$

Before proving the theorem, I provide a preliminary lemma. The proof of the lemma is informative about the relation between the primal and dual problem; thus, I provide it here.

Lemma A.1. *Consider an admissible contract $\mathcal{C}$ and a process $(U_t)_{t\geq 0} \in \mathcal{I}$. Let $\xi_t := \tilde{\xi}(\mathcal{C}_t, \phi_t)$, where $\mathcal{C}_t$ is the continuation contract of $\mathcal{C}$ at time t. Then,*

$$\mathrm{E}\left[\int_0^\infty e^{-rt} Y_t \beta_t \eta \phi_t (1-\phi_t)\, dt \middle| \mathcal{F}_0\right] = Y_0 \tilde{\xi}(\mathcal{C}, \phi_0) + \mathrm{E}\left[\int_0^\infty e^{-rt}\left(\eta^2 \phi_t \frac{\eta \xi_t}{\lambda} C_t^\rho dt + \eta \xi_t dU_t\right)\middle| \mathcal{F}_0\right]. \tag{A.1}$$

Moreover, if $\mathcal{C} \in \mathcal{D}_{(V_0, Y_0, \phi_0)}$, it must satisfy $\beta_t \geq \xi_t$ for all $t \geq 0$.

Proof. Consider an admissible contract $\mathcal{C} = (C_t, \beta_t)_{t\geq 0}$ and let $\tilde{Y}_t := e^{(\delta - r)t} Y_t$. For a finite $T > 0$, we can use integration by parts to obtain

$$\begin{aligned}
&\mathrm{E}\left[\int_0^T e^{-rt} Y_t \beta_t \eta \phi_t (1-\phi_t)\, dt \middle| \mathcal{F}_0\right] \\
&= \mathrm{E}\left[\int_0^T e^{-\delta t}\left(\tilde{Y}_t \beta_t \eta \phi_t (1-\phi_t)\right) dt \middle| \mathcal{F}_0\right] \\
&= -\mathrm{E}\left\{\left[\tilde{Y}_t \int_t^T e^{-\delta s} \beta_s \eta \phi_s (1-\phi_s)\, ds\right]_0^T + \int_0^T \left(\int_t^T e^{-\delta s} \beta_s \eta \phi_s (1-\phi_s)\, ds\right) d\tilde{Y}_t \middle| \mathcal{F}_0\right\} \\
&= \tilde{Y}_0 \mathrm{E}\left[\int_0^T e^{-\delta s} \beta_s \eta \phi_s (1-\phi_s)\, ds \middle| \mathcal{F}_0\right] + \mathrm{E}\left[\int_0^T e^{-\delta t}\left(\int_t^T e^{-\delta(s-t)} \beta_s \eta \phi_s (1-\phi_s)\, ds\right) d\tilde{Y}_t \middle| \mathcal{F}_0\right] \\
&= \tilde{Y}_0 \mathrm{E}\left[\int_0^T e^{-\delta s} \beta_s \eta \phi_s (1-\phi_s)\, ds \middle| \mathcal{F}_0\right] + \mathrm{E}\left[\int_0^T e^{-rt}\left(\int_t^T e^{-\delta(s-t)} \beta_s \eta \phi_s (1-\phi_s)\, ds\right)\left(\frac{\eta^2}{\lambda}\phi_t C_t^\rho dt + \eta dU_t\right)\middle| \mathcal{F}_0\right].
\end{aligned} \tag{A.2}$$

By definition of admissible contract, $\mathrm{E}\big[\int_0^\infty e^{-rt}\,|Y_t \beta_t \eta \phi_t (1-\phi_t)|\, dt \big| \mathcal{F}_0\big] < \infty$. Hence, by the monotone convergence theorem and the law of iterated expectations, we can conclude

$$\begin{aligned}
\mathrm{E}\left[\int_0^\infty e^{-rt} Y_t \beta_t \eta \phi_t (1-\phi_t)\, dt \middle| \mathcal{F}_0\right] &= \lim_{T\to\infty} \mathrm{E}\left[\int_0^T e^{-rt}\left(Y_t \beta_t \eta \phi_t (1-\phi_t)\right) dt \middle| \mathcal{F}_0\right] \\
&= \tilde{Y}_0 \tilde{\xi}(\mathcal{C}, \phi_0) + \mathrm{E}\left[\int_0^\infty e^{-rt}\left(\eta^2 \phi_t \frac{\xi_t}{\lambda} C_t^\rho dt + \eta \xi_t dU_t\right)\middle| \mathcal{F}_0\right].
\end{aligned}$$

Because $\tilde{Y}_0 = Y_0$, this proves the first part of the lemma.

To prove the last part, I use (A.1) to rewrite the objective function in (17) as

$$\begin{aligned}
&G(V_0, Y_0, \phi_0 | \mathcal{C}) = \\
&\quad \sup_{(U_t)_{t\geq 0} \in \mathcal{I}} \mathrm{E}\left[\int_0^\infty e^{-rt}\left\{\left(C_t - \frac{\eta \phi_t}{\lambda}(\beta_t - \eta \xi_t) C_t^\rho\right) dt - (\beta_t - \eta \xi_t) dU_t\right\}\right] + Y_0 \tilde{\xi}(\mathcal{C}, \phi).
\end{aligned} \tag{A.3}$$

Because dU_t can be arbitrarily large, if $\mathcal{C}$ is an optimal dual contract, it must satisfy $\beta_t - \eta \xi_t \geq 0$ for all t. $\square$

Proof of Theorem 3(i). Consider Y^0 and Y^1 such that $Y^0 \leq Y^1$. To prove $G^*(V, \cdot, \phi)$ is increasing, consider $\mathcal{C}_1 \in \mathcal{D}_{(V, Y_1, \phi)}$ and note

$$G^*(V, Y^0, \phi) \leq G\big(V, Y^0, \phi | \mathcal{C}_1\big) = G^*(V, Y^1, \phi) + (Y^0 - Y^1)\tilde{\xi}\big(\mathcal{C}_1, \phi\big) \leq G^*(V, Y^1, \phi),$$

where the equality follows from equation (A.3).

To prove $G^*(V,\cdot,\phi)$ is concave, define $Y^\nu := \nu Y^0 + (1-\nu)Y^1$ for $\nu \in [0,1]$. Let $\mathcal{C}_\nu \in \mathcal{D}_{(V,Y^\nu,\phi)}$. Then,

$$\begin{aligned}
&\nu G^*(V,Y^0,\phi) + (1-\nu)G^*(V,Y^1,\phi) \\
&\leq \nu G\big(V,Y^0,\phi|\mathcal{C}_\nu\big) + (1-\nu)G\big(V,Y^1,\phi|\mathcal{C}_\nu\big) \\
&= G^*(V,Y^\nu,\phi),
\end{aligned}$$

where the last equality follows from equation (A.3).

Finally, I prove $G^*(V,Y,\phi)$ is differentiable in Y for $Y \neq 0$ in Lemma O.5 of the online appendix. I delegate this proof to the online appendix because it is a purely technical proof involving lengthy arguments based on viscosity solution concepts.[21]

PROOF OF THEOREM 3(II). I begin by proving equation (18). Let $\mathcal{C}$ be an admissible contract that satisfies (8) and $\beta_t \geq \eta\tilde{\xi}(\mathcal{C}_t,\phi_t)$ for all $t \geq 0$ and that delivers continuation value V_0 and information rent ξ_0 to the agent. Then,

$$\begin{aligned}
J(V_0,\xi_0,\phi_0|\mathcal{C}) &= \mathrm{E}\left[\int_0^\infty e^{-rt}\left(C_t - \eta\phi_t\frac{\beta_t-\eta\xi_t}{\lambda}C_t^\rho\right)dt\Big|\mathcal{F}_0\right] \\
&= G(V,Y,\phi|\mathcal{C}) + \mathrm{E}\left[\int_0^\infty e^{-rt}\frac{\eta^2\phi_t}{\lambda}\xi_t C_t^\rho\,dt\Big|\mathcal{F}_0\right] - \mathrm{E}\left[\int_0^\infty e^{-rt}Y_t\beta_t\eta\phi_t(1-\phi_t)\,dt\Big|\mathcal{F}_0\right] \\
&= G(V_0,Y_0,\phi_0|\mathcal{C}) - Y_0\xi_0,
\end{aligned} \tag{A.4}$$

where the last equality follows from Lemma A.1 and the fact that when $\beta_t \geq \eta\tilde{\xi}(\mathcal{C}_t,\phi_t)$ for all $t \geq 0$, the supremum in (A.3) is achieved by a process $(U_t^{\mathcal{C}})_{t\geq 0} \in \mathcal{I}$ such that $(\beta_t - \eta\xi_t)dU_t^{\mathcal{C}} = 0$ for all $t \geq 0$.

Consider $\mathcal{C}^P \in \mathcal{P}_{(V,\xi,\phi)}$ and $\mathcal{C}^D \in \mathcal{D}_{(V,Y,\phi)}$. Using (A.4), we have

$$J^*(V,\xi,\phi) + Y\xi = J\left(V,\xi,\phi|\mathcal{C}^P\right) + Y\xi = G\left(V,Y,\phi|\mathcal{C}^P\right) \geq G^*(V,Y,\phi) \tag{A.5}$$

and

$$J^*\left(V,\tilde{\xi}\left(\mathcal{C}^D,\phi\right),\phi\right) \leq J\left(V,\tilde{\xi}\left(\mathcal{C}^D,\phi\right),\phi|\mathcal{C}^D\right) = G^*(V,Y,\phi) - Y\tilde{\xi}\left(\mathcal{C}^D,\phi\right). \tag{A.6}$$

Equation (18) follows directly from (A.5) and (A.6). In particular, the infimum in (18) is achieved by $\xi = \tilde{\xi}\big(\mathcal{C}^D,\phi\big)$; that is,

$$\tilde{\xi}\left(\mathcal{C}^D,\phi\right) \in \arg\inf_{\xi\geq 0}\{J^*(V,\xi,\phi) + Y\xi\}. \tag{A.7}$$

Next, I show that, for any optimal dual contract $\mathcal{C}^D \in \mathcal{D}_{(V,Y,\phi)}$, $\tilde{\xi}\big(\mathcal{C}^D,\phi\big) = G^*_{Y+}(V,Y,\phi)$.

[21] See Crandall et al. (1992) and Pham (2009) for the notion of viscosity solution of a partial differential equation.

Let $\mathcal{C}^D \in \mathcal{D}_{(V,Y,\phi)}$. Using (A.5) and (A.6), we obtain

$$G^*(V,Y,\phi) = J^*\left(V, \tilde{\xi}\left(\mathcal{C}^D,\phi\right),\phi\right) + Y\tilde{\xi}\left(\mathcal{C}^D,\phi\right)$$
$$G^*(V,Y+\varepsilon,\phi) \leq J^*\left(V, \tilde{\xi}\left(\mathcal{C}^D,\phi\right),\phi\right) + (Y+\varepsilon)\tilde{\xi}\left(\mathcal{C}^D,\phi\right).$$

Together, they imply

$$\frac{G^*(V,Y+|\varepsilon|,\phi) - G^*(V,Y,\phi)}{|\varepsilon|} \leq \tilde{\xi}\left(\mathcal{C}^D,\phi\right) \quad \text{if } \varepsilon > 0$$
$$\frac{G^*(V,Y,\phi) - G^*(V,Y-|\varepsilon|,\phi)}{|\varepsilon|} \geq \tilde{\xi}\left(\mathcal{C}^D,\phi\right) \quad \text{if } \varepsilon < 0.$$

For $Y > 0$, the function G^* is differentiable in Y, and hence, the left and right derivatives with respect to Y exist and coincide. For $Y \geq 0$, the right derivatives exist because of the concavity of $G^*(V,\cdot,\phi)$. After taking limits for $\varepsilon \to 0^+$ and $\varepsilon \to 0^-$, we obtain

$$\tilde{\xi}\left(\mathcal{C}^D,\phi\right) = G^*_Y(V,Y,\phi) \quad \text{for} \quad Y > 0 \tag{A.8}$$

and

$$\tilde{\xi}\left(\mathcal{C}^D,\phi\right) \geq G^*_{Y^+}(V,0,\phi) \quad \text{for} \quad Y = 0.$$

I then prove $\tilde{\xi}\big(\mathcal{C}^D,\phi\big) = G^*_{Y^+}(V,0,\phi)$ when $\mathcal{C}^D \in \mathcal{D}_{(V,0,\phi)}$.[22] For $n > 0$, define $\tau_n := \sup\{t \in (0,1/n] : \phi_t = \phi_0\}$. By the Blumenthal Zero-One Law, $\tau_n > t$ exists because $\phi_t - \phi_0$ changes sign infinitely many times in any time interval $[0,1/n]$ (Karatzas and Shreve, 1998, Chapter 2.7.C). By Lemma O.3 in the online appendix, $G^*(V_0,Y_0,\phi_0) = v_0 g^*(y_0,\phi_0)$ for a function g^* where $v_t := ((1-\rho)V_t)^{\frac{1}{1-\rho}}$ and $y_t := v_t^{-\rho}Y_t$. The function g^* inherits some of the properties of G^*. In particular, $g^*(y,\phi)$ is increasing and concave in y and differentiable in y for $y \neq 0$ with $G^*_{Y^+}(V,Y,\phi) = (1-\rho)V g^*_{y^+}(y,\phi)$.

For any $\mathcal{C}^D \in \mathcal{D}_{(V_0,0,\phi_0)}$, let $\mathcal{C}^D_t$ be its continuation contract at time t. Then,

$$\begin{aligned}
\tilde{\xi}(\mathcal{C}^D,\phi_0) &= \mathrm{E}\left[\int_0^{\tau_n} e^{-\delta t}\beta_t\eta\phi_t(1-\phi_t)\,dt\Big|\mathcal{F}_0\right] + E\left[e^{-\delta\tau_n}\tilde{\xi}(\mathcal{C}^D_{\tau_n},\phi_0)|\mathcal{F}_0\right] \\
&= \mathrm{E}\left[\int_0^{\tau_n} e^{-\delta t}\beta_t\eta\phi_t(1-\phi_t)\,dt\Big|\mathcal{F}_0\right] + E\left[e^{-\delta\tau_n}(1-\rho)V_{\tau_n}g^*_y(y_{\tau_n},\phi_0)|\mathcal{F}_0\right] \\
&\leq \mathrm{E}\left[\int_0^{\tau_n} e^{-\delta t}\beta_t\eta\phi_t(1-\phi_t)\,dt\Big|\mathcal{F}_0\right] + (1-\rho)g^*_{y^+}(0,\phi_0)E\left[e^{-\delta\tau_n}V_{\tau_n}|\mathcal{F}_0\right].
\end{aligned}$$

The first equality is because of (10). The second equality is because, by Bellman's optimality principle, $\mathcal{C}^D_{\tau_n} \in \mathcal{D}_{(V_{\tau_n},Y_{\tau_n},\phi_{\tau_n})}$ and because, with $Y_{\tau_n} > 0$ by (16), equation (A.8) implies $\tilde{\xi}\big(\mathcal{C}^D_{\tau_n},\phi_0\big) = G^*_{Y^+}(V_{\tau_n},Y_{\tau_n},\phi_0) = (1-\rho)V_{\tau_n}g^*_{y^+}(y_{\tau_n},\phi_0)$. The final inequality follows from $y_{\tau_n} > 0$ and the concavity of $g^*(\cdot,\phi)$.

Taking the limit for $n \to \infty$, the first expectation in the right-hand side converges to

[22] Note $G^*_{Y^+}(V,0,\phi)$ is finite because of the global concavity of $G^*(V,\cdot,\phi)$.

zero by the monotone convergence theorem. As for the second expectation, we have

$$V_0 = \mathrm{E}\left[\int_0^{\tau_n} e^{-\delta t}\frac{C_t^{1-\rho}}{1-\rho}\,dt\Big|\mathcal{F}_0\right] + E\left[e^{-\delta\tau_n}V_{\tau_n}|\mathcal{F}_0\right].$$

Taking the limit for $n\to\infty$ in this expression, $\mathrm{E}\big[\int_0^{\tau_n} e^{-\delta t}\frac{C_t^{1-\rho}}{1-\rho}\,dt\big|\mathcal{F}_0\big]\to 0$ by the monotone convergence theorem. Hence, $E[e^{-\delta\tau_n}V_{\tau_n}|\mathcal{F}_0]\to V_0$.

Therefore, for any $\mathcal{C}^D\in\mathcal{D}_{(V_0,0,\phi_0)}$,

$$\tilde{\xi}\left(\mathcal{C}^D,\phi_0\right)\le(1-\rho)V_0 g^*_{y^+}(0,\phi_0)=G^*_{Y^+}(V_0,0,\phi_0).$$

Because V_0 and ϕ_0 were arbitrary and because, from the previous results, $\tilde{\xi}\big(\mathcal{C}^D,\phi\big)\ge G^*_{Y^+}(V,0,\phi)$ when $\mathcal{C}^D\in\mathcal{D}_{(V,0,\phi)}$, we conclude $\tilde{\xi}\big(\mathcal{C}^D,\phi\big)=G^*_{Y^+}(V,0,\phi)$ for any $\mathcal{C}^D\in\mathcal{D}_{(V,0,\phi)}$.

To conclude the proof, it remains to show that, given $\xi'=\tilde{\xi}\big(\mathcal{C}^D,\phi\big)=G^*_{Y^+}(V,Y,\phi)$, if $\mathcal{C}^P\in\mathcal{P}_{(V,\xi',\phi)}$, then $\mathcal{C}^P\in\mathcal{D}_{(V,Y,\phi)}$. To show it, I combine (A.5) and (A.6) to obtain

$$J^*(V,\xi',\phi)+Y\xi'=G\left(V,Y,\phi|\mathcal{C}^P\right)\ge G^*(V,Y,\phi)\ge J^*(V,\xi',\phi)+Y\xi',$$

which thus implies $G\big(V,Y,\phi|\mathcal{C}^P\big)=G^*(V,Y,\phi)$, and that $\mathcal{C}^P\in\mathcal{D}_{(V,Y,\phi)}$. □

Proof of Theorem 3(iii). Consider $\xi\le G^*_{Y^+}(V,0,\phi)$ and let $\bar{Y}\in\arg\sup_{Y'\ge 0}\{G^*(V,Y',\phi)-Y'\xi\}$. Because $G^*(V,\cdot,\phi)$ is concave and differentiable for $Y>0$, we have $G^*_{Y^+}(V,\bar{Y},\phi)=\xi$. Therefore, given information rent $\xi\le G^*_{Y^+}(V,0,\phi)$, any optimal dual contract $\bar{\mathcal{C}}^D\in\mathcal{D}_{(V,\bar{Y},\phi)}$ delivers information rent ξ. Then,

$$J\left(V,\xi,\phi|\bar{\mathcal{C}}^D\right)\ge J^*(V,\xi,\phi)\ge G^*(V,\bar{Y},\phi)-\bar{Y}\xi=G^*(V,\bar{Y},\phi)-\bar{Y}\tilde{\xi}\left(\bar{\mathcal{C}}^D,\phi\right)=J\left(V,\xi,\phi|\bar{\mathcal{C}}^D\right),$$

where the first inequality follows because $\bar{\mathcal{C}}^D$ cannot be strictly better than the optimal (primal) contract, and the second inequality follows from (A.5). The subsequent equality is because $\tilde{\xi}\big(\bar{\mathcal{C}}^D,\phi\big)=G^*_{Y^+}(V,\bar{Y},\phi)=\xi$, and the last equality is a consequence of (A.4).

Because the first and the last terms coincide, we conclude $J^*(V,\xi,\phi)=J\big(V,\xi,\phi|\bar{\mathcal{C}}^D\big)$; that is, $\bar{\mathcal{C}}^D\in\mathcal{P}_{(V,\xi,\phi)}$. Because of this observation and (A.5), we therefore have $J^*(V,\xi,\phi)=\sup_{Y\ge 0}\{G^*(V,Y',\phi)-Y'\xi\}$, where the supremum is achieved by $\bar{Y}$ such that $G^*_{Y^+}(V,\bar{Y},\phi)=\xi$. By the envelope theorem (Milgrom and Segal, 2002), $\bar{Y}=-J^*_\xi(V,\xi,\phi)$. □

Proof of Theorem 3(iv). Consider $\mathcal{C}_0^D\in\mathcal{D}_{(V,0,\phi)}$. By equation (A.6), $G^*(V,0,\phi)\ge J^*\big(V,\tilde{\xi}\big(\mathcal{C}_0^D,\phi\big),\phi\big)$. Moreover by (A.5), we also have $G^*(V,0,\phi)\le J(V,\xi,\phi)$ for any ξ. Therefore,

$$J^*\left(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi\right)\le J(V,\xi,\phi)\quad\text{for any }\xi,$$

proving $\tilde{\xi}(\mathcal{C}_0^D,\phi)$ is a global minimizer of $J(V,\xi,\phi)$ with respect to ξ. Hence, any $\mathcal{C}^P \in \mathcal{P}_{(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi)}$ is a relaxed optimal contract whose cost is

$$J^*\big(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi\big) = G^*(V,0,\phi),$$

by (A.5) and (A.6). Note also that, by Theorem 3(II), $\tilde{\xi}(\mathcal{C}_0^D,\phi) = G^*_{Y+}(V,0,\xi)$ for any $\mathcal{C}_0^D \in \mathcal{D}_{(V,0,\phi)}$.

Combining the results from part (II) and the proof of part (III) of this theorem, we have $\mathcal{C} \in \mathcal{P}_{(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi)}$ if and only if $\mathcal{C} \in \mathcal{D}_{(V,0,\phi)}$. Thus, any optimal dual contract $\mathcal{C}_0^D \in \mathcal{D}_{(V,0,\phi)}$ is a relaxed optimal contract. However, there might be multiple minimizers of $J^*(V,\xi,\phi)$ with respect to ξ. That is, there could be a relaxed optimal contract $\mathcal{C}$ such that $\mathcal{C} \notin \mathcal{P}_{(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi)}$. I rule out this possibility by showing $\tilde{\xi}(\mathcal{C}_0^D,\phi)$ is the unique minimizer of $J^*(V,\xi,\phi)$ with respect to ξ.

I proceed by contradiction. Suppose $\mathcal{C}^{P''}$ and $\xi'' \neq \tilde{\xi}(\mathcal{C}_0^D,\phi)$ exist such that $\mathcal{C}^{P''} \in \mathcal{P}_{(V,\xi'',\phi)}$ and $J^*(V,\xi'',\phi) = J^*\big(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi\big)$. Using (A.5), we have

$$J^*\left(V,\tilde{\xi}\left(\mathcal{C}_0^D,\phi\right),\phi\right) = J^*(V,\xi'',\phi) = J\left(V,\xi'',\phi|\mathcal{C}^{P''}\right) = G\left(V,0,\phi|\mathcal{C}^{P''}\right) \geq G^*(V,0,\phi).$$

Because $J^*\big(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi\big) = G^*(V,0,\phi)$, the inequality above holds as an equality and $\mathcal{C}^{P''} \in \mathcal{D}_{(V,0,\phi)}$. By assumption, $\mathcal{C}^{P''}$ delivers information rent $\xi'' \neq \tilde{\xi}(\mathcal{C}_0^D,\phi) = G_{Y+}(V,0,\phi)$. However, by Theorem 3(II), any optimal dual contract $\mathcal{C} \in \mathcal{D}_{(V,0,\phi)}$ delivers information rent $G_{Y+}(V,0,\phi)$, thus creating a contradiction.

Therefore, for any $\mathcal{C}_0^D \in \mathcal{D}_{(V,0,\phi)}$, $\tilde{\xi}(\mathcal{C}_0^D,\phi) = G_{Y+}(V,0,\phi)$ is the unique minimizer of $J^*(V,\cdot,\phi)$. Hence, if $\mathcal{C}$ is a relaxed optimal contract, $\mathcal{C} \in \mathcal{P}_{(V,\tilde{\xi}(\mathcal{C}_0^D,\phi),\phi)}$, thus concluding the proof. □

A.E PROOF OF THEOREM 4

PROOF OF THEOREM 4(I). To begin with, I show $c(y,\phi)$ is uniformly bounded away from zero for all $y \geq 0$ and $\phi \in [0,1]$. Consider (24) and the first-order condition for $c(y,\phi)$:

$$1 - A(y,\phi)c(y,\phi)^{\rho-1} - (g(y,\phi) - \rho y g_y(y,\phi))c(y,\phi)^{-\rho} \geq 0, \tag{A.9}$$

where $A(y,\phi) := \eta\phi\frac{\hat{\beta}(y,\phi)-\eta g_y(y,\phi)}{\lambda}\rho \geq 0$.

Note

$$g(y,\phi) - \rho y g_y(y,\phi) \geq g(y,\phi) - y g_y(y,\phi) \geq g(0,\phi) \geq j_1 > 0,$$

where the first inequality is because $\rho < 1$ and $y g_y(y,\phi) \geq 0$, the second is because $g(\cdot,\phi)$ is concave, and the third is because $g \geq j_1$. Furthermore, note that $c(y,\phi)$ is interior because, as $c(y,\phi) \to 0$ the left-hand side in (A.9) tends to $-\infty$. Moreover, as $c(y,\phi) \to \infty$, the left-hand side of (A.9) tends to 1. Hence, $0 < c(y,\phi) < \infty$ and (A.9) holds as an equality. Together, these observations also imply $1 - A(y,\phi)c(y,\phi)^{\rho-1} > 0$.

Therefore,

$$c(y,\phi) = (1 - A(y,\phi)c(y,\phi)^{\rho-1})^{-\frac{1}{\rho}}(g(y,\phi) - \rho y g_y(y,\phi))^{\frac{1}{\rho}} \geq j_1^{\frac{1}{\rho}} > 0,$$

proving $c(y,\phi)$ is uniformly bounded away from zero.

Next, I show that for any $t \geq 0$, the contract $\mathcal{C}^R_{(v_0,y_0,\phi_0)}$ provides continuation value $\frac{v_t^{1-\rho}}{1-\rho}$ to the agent if he does not shirk.

Consider the variable $D_t = \frac{v_t^{1-\rho}}{1-\rho}$. By Ito's lemma,

$$dD_t = (1-\rho)D_t\left(\frac{\delta}{1-\rho} - \frac{c(y_t,\phi_t)^{1-\rho}}{1-\rho}\right)dt + (1-\rho)D_t\hat{\beta}(y_t,\phi_t)dW^{\mathcal{C}}_t.$$

Consider a localizing sequence of stopping times $(\tau_n)_{n=0}^{\infty}$ such that $\tau_n \to \infty$ as $n \to \infty$. Then, using the Dynkin's formula (Øksendal, 2003, Chapter 7.4) and taking the limit for $n \to \infty$, we have

$$D_t = \mathrm{E}\left[\int_t^{\infty} e^{-\delta(s-t)}\frac{(C^R_s)^{1-\rho}}{1-\rho}\,ds\Big|\mathcal{F}_t\right] + \lim_{n\to\infty}\mathrm{E}[e^{-\delta(\tau_n-t)}D_{\tau_n}|\mathcal{F}_t],$$

by the monotone convergence theorem.

To show $\lim_{n\to\infty}\mathrm{E}[e^{-\delta(\tau_n-t)}D_{\tau_n}|\mathcal{F}_t] = 0$, consider

$$\begin{aligned}
e^{-\delta(\tau_n-t)}D_{\tau_n} &= D_t\exp\left\{\int_t^{\tau_n}\left(-c(y_s,\phi_s)^{1-\rho} - \frac{1}{2}(1-\rho)^2\hat{\beta}(y_s,\phi_s)^2\right)ds + \int_t^{\tau_n}(1-\rho)\hat{\beta}(y_s,\phi_s)\,dW^{\mathcal{C}}_s\right\} \\
&\leq D_t\exp\left\{\int_t^{\tau_n}\left(-j_1^{\frac{1-\rho}{\rho}} - \frac{1}{2}(1-\rho)^2\hat{\beta}(y_s,\phi_s)^2\right)ds + \int_t^{\tau_n}(1-\rho)\hat{\beta}(y_s,\phi_s)\,dW^{\mathcal{C}}_s\right\}.
\end{aligned}$$

The random variable $e^{\int_t^{\tau_n}-\frac{1}{2}(1-\rho)^2\hat{\beta}(y_s,\phi_s)^2\,ds + \int_t^{\tau_n}(1-\rho)\hat{\beta}(y_s,\phi_s)\,dW^{\mathcal{C}}_s}$ satisfies the Novikov condition because $\hat{\beta}(y,\phi)$ is uniformly bounded. Hence, it is a martingale and $\mathrm{E}\left[e^{\int_t^{\tau_n}-\frac{1}{2}(1-\rho)^2\hat{\beta}(y_s,\phi_s)^2\,ds + \int_t^{\tau_n}(1-\rho)\hat{\beta}(y_s,\phi_s)\,dW^{\mathcal{C}}_s}|\mathcal{F}_t\right] = 1$. Therefore, using also $D_{\tau_n} \geq 0$,

$$0 \leq \mathrm{E}[e^{-\delta(\tau_n-t)}D_{\tau_n}|\mathcal{F}_t] \leq D_t e^{-j_1^{\frac{1-\rho}{\rho}}(\tau_n-t)},$$

and, thus, $\lim_{n\to\infty}\mathrm{E}[e^{-\delta(\tau_n-t)}D_{\tau_n}|\mathcal{F}_t] = 0$.

We therefore conclude

$$\frac{v_t^{1-\rho}}{1-\rho} = D_t = \mathrm{E}\left[\int_t^{\infty} e^{-\delta(s-t)}\frac{(C^R_s)^{1-\rho}}{1-\rho}\,ds\Big|\mathcal{F}_t\right],$$

proving $\frac{v_t^{1-\rho}}{1-\rho}$ is the agent's continuation value at time t if the agent does not shirk. $\square$

Proof of Theorem 4(ii). I begin by proving $v_t g(y_t,\phi_t) \leq G^*(V_t,Y_t,\phi_t)$. To do so, consider an optimal dual contract $\mathcal{C} = (C_t,\beta_t)_{t\geq 0}$ that delivers continuation value $V_0 =$

$\frac{v_0^{1-\rho}}{1-\rho}$ to the agent. Define

$$\begin{aligned}\mathcal{A}[g;(y,\phi);c,\hat{\beta}] \coloneqq g(y,\phi)&\left(-r+\frac{\delta}{1-\rho}-\frac{c^{1-\rho}}{1-\rho}+\frac{1}{2}\rho\hat{\beta}^2\right)\\&+g_y(y,\phi)\left[-\phi\frac{\eta^2}{\lambda}c^\rho+y\left(r-\frac{\delta}{1-\rho}+\rho\frac{c^{1-\rho}}{1-\rho}+\frac{1}{2}\rho\hat{\beta}^2\right)\right]\\&-g_y(y,\phi)y\rho\hat{\beta}^2+g_\phi(y,\phi)\eta\phi(1-\phi)\hat{\beta}\\&+\frac{1}{2}g_{yy}(y,\phi)(y\rho\hat{\beta})^2+\frac{1}{2}g_{\phi\phi}(y,\phi)\eta^2\phi^2(1-\phi)^2-g_{y\phi}(y,\phi)y\rho\hat{\beta}\eta\phi(1-\phi).\end{aligned}$$

Let $c_t \coloneqq C_t/v_t$ and $\hat{\beta}_t \coloneqq \beta_t/v_t^{1-\rho}$. For a process $U \coloneqq (U_t)_{t\ge 0} \in \mathcal{I}$, the multiplier $y_t^U \coloneqq v_t^{-\rho}Y_t$ evolves as

$$dy_t^U = \phi_t\frac{\eta^2}{\lambda}c_t^\rho dt + y_t\left(r-\frac{\delta}{1-\rho}+\rho\frac{c_t^{1-\rho}}{1-\rho}+\frac{1}{2}\rho\hat{\beta}_t^2\right)dt - y_t\rho\hat{\beta}_t dW_t^{\mathcal{C}} + \eta v_t^{-\rho}dU_t.$$

Consider also a process $U^c \coloneqq (U_t^c)_{t\ge 0} \in \mathcal{I}$ with $dU_t^c = I_t dt$, where $I_t \ge 0$ is such that

$$\left[c_t - \frac{\eta}{\lambda}\phi_t\hat{\beta}_t c^\rho + y_t\hat{\beta}_t\eta\phi_t(1-\phi_t)\right] + \mathcal{A}[g;(y_t,\phi_t);c_t,\hat{\beta}_t] - (\hat{\beta}_t - \eta g_{y^+}(y_t,\phi_t))I_t \ge 0. \quad \text{(A.10)}$$

Note such I_t exists because if $\hat{\beta}_t \ge \eta g_{y^+}(y_t,\phi_t)$, $I_t = 0$ satisfies this condition by (24). If, on the other hand, $0 \le \hat{\beta}_t < \eta g_{y^+}(y_t,\phi_t)$, a large enough I_t will satisfy this condition.

Consider a localizing sequence of stopping times $(\tau_n)_{n=0}^\infty$ such that $\tau_n \to \infty$ as $n \to \infty$. Then, by the Dynkin's formula,

$$\mathrm{E}[e^{-r\tau_n}v_{\tau_n}g(y_{\tau_n}^{U^c},\phi_{\tau_n})|\mathcal{F}_0] - v_0 g(y_0,\phi_0) = \mathrm{E}\left[\int_0^{\tau_n} e^{-rt}v_t\left\{\mathcal{A}[g;(y_t^{U^c},\phi_t);c_t,\hat{\beta}_t] + \eta g_y(y_t^{U^c},\phi_t)I_t\right\}dt\Big|\mathcal{F}_0\right].$$

Using inequality (A.10) and the differentiability of $g(\cdot,\phi)$ inside the domain, we obtain

$$\begin{aligned}v_0 g(y_0,\phi_0) &\le \mathrm{E}\left[\int_0^{\tau_n} e^{-rt}v_t\left(c_t - \frac{\eta}{\lambda}\phi_t\hat{\beta}_t c^\rho + y_t^{U^c}\hat{\beta}_t\eta\phi_t(1-\phi_t) - \hat{\beta}_t I_t\right)dt\Big|\mathcal{F}_0\right] + \mathrm{E}[e^{-r\tau_n}v_{\tau_n}g(y_{\tau_n}^{U^c},\phi_{\tau_n})|\mathcal{F}_0]\\&\le \sup_{(U_t)_{t\ge0}\in\mathcal{I}}\mathrm{E}\left[\int_0^{\tau_n} e^{-rt}\left\{\left(C_t - \frac{\eta}{\lambda}\phi_t\beta_t C_t^\rho + Y_t\beta_t\eta\phi_t(1-\phi_t)\right)dt - \beta_t dU_t\right\}\Big|\mathcal{F}_0\right]\\&\qquad + \sup_{(U_t)_{t\ge0}\in\mathcal{I}}\mathrm{E}[e^{-r\tau_n}v_{\tau_n}g(y_{\tau_n}^U,\phi_{\tau_n})|\mathcal{F}_0].\end{aligned} \quad \text{(A.11)}$$

First, I show $\lim_{n\to\infty}\mathrm{E}[e^{-r\tau_n}v_{\tau_n}g(y_{\tau_n}^U,\phi_{\tau_n})|\mathcal{F}_0] = 0$ for any $(U_t)_{t\ge0} \in \mathcal{I}$. Define $j_0^* \coloneqq g^*(0,0)$. By Theorem 3 and equation (21), $v_{\tau_n}j_0^* \ge 0$ is the cost of the cheapest (admissible) contract that delivers continuation value $\frac{v_{\tau_n}^{1-\rho}}{1-\rho}$ to the agent without any investment.[23] I

[23] Recall that throughout the paper, I assume parameters are such that $g^*(0,\phi) > 0$ for all $\phi \in [0,1]$.

then proceed as in Di Tella and Sannikov (2021) by noting

$$\mathrm{E}\left[\int_0^\infty e^{-rt}|C_t|\right] \geq \mathrm{E}\left[\int_0^\infty e^{-rt}C_t\right] \geq \mathrm{E}\left[\int_0^{\tau_n} e^{-rt}C_t + e^{-r\tau_n}v_{\tau_n}j_0^*\right].$$

Because of the admissibility condition (14) and because $v_{\tau_n} \geq 0$, I apply the dominated convergence theorem and conclude $\lim_{n\to\infty}\mathrm{E}[e^{-r\tau_n}v_{\tau_n}|\mathcal{F}_0] = 0$. Because $0 \leq \mathrm{E}[e^{-r\tau_n}v_{\tau_n}g(y^U_{\tau_n},\phi_{\tau_n})|\mathcal{F}_0] \leq j_0\mathrm{E}[e^{-r\tau_n}v_{\tau_n}|\mathcal{F}_0]$, then $\lim_{n\to\infty}\mathrm{E}[e^{-r\tau_n}v_{\tau_n}g(y^U_{\tau_n},\phi_{\tau_n})|\mathcal{F}_0] = 0$.

Because $\mathcal{C} = (C_t, K_t)_{t\geq 0}$ is assumed to be an optimal dual contract, by Lemma (A.1), $\beta_t \geq \eta\tilde{\xi}(\mathcal{C}_t,\phi_t)$. Moreover, for any τ_n, $\tilde{\xi}(\mathcal{C}_t,\phi_t) \geq \tilde{\xi}^T(\mathcal{C}_t,\phi_t,\tau_n) := \mathrm{E}\left[\int_t^{\tau_n} e^{-\delta(s-t)}\beta_s\eta\phi_s(1-\phi_s)\,ds|\mathcal{F}_t\right]$. Using equation (A.2) in the proof of Lemma (A.1), we therefore have

$$\begin{aligned}
&\sup_{(U_t)_{t\geq 0}\in\mathcal{I}}\mathrm{E}\left[\int_0^{\tau_n} e^{-rt}\left\{\left(C_t - \frac{\eta}{\lambda}\phi_t\beta_tC_t^\rho + Y_t\beta_t\eta\phi_t(1-\phi_t)\right)dt - \beta_t dU_t\right\}\Big|\mathcal{F}_0\right]\\
&= \sup_{(U_t)_{t\geq 0}\in\mathcal{I}}\mathrm{E}\left[\int_0^{\tau_n} e^{-rt}\left\{\left(C_t - \frac{\eta\phi_t}{\lambda}(\beta_t - \eta\tilde{\xi}^T(\mathcal{C}_t,\phi_t,\tau_n))C_t^\rho\right)dt - (\beta_t - \eta\tilde{\xi}^T(\mathcal{C}_t,\phi_t,\tau_n))dU_t\right\}\right] + Y\tilde{\xi}^T(\mathcal{C},\phi_0,\tau_n)\\
&= \mathrm{E}\left[\int_0^{\tau_n} e^{-rt}\left\{\left(C_t - \frac{\eta\phi_t}{\lambda}(\beta_t - \eta\tilde{\xi}^T(\mathcal{C}_t,\phi_t,\tau_n))C_t^\rho\right)dt\right\}\right] + Y_0\tilde{\xi}^T(\mathcal{C},\phi_0,\tau_n)\\
&= \mathrm{E}\left[\int_0^{\tau_n} e^{-rt}\left(C_t - \frac{\eta}{\lambda}\phi_t\beta_tC_t^\rho + Y_t\beta_t\eta\phi_t(1-\phi_t)\right)dt\Big|\mathcal{F}_0\right].
\end{aligned}$$

Taking the limit as $n\to\infty$ in (A.11) and applying the dominated convergence theorem, we obtain

$$v_0g(y_0,\phi_0) \leq \mathrm{E}\left[\int_0^\infty e^{-rt}\left(C_t - \frac{\eta}{\lambda}\phi_t\beta_tC_t^\rho + Y_t\beta_t\eta\phi_t(1-\phi_t)\right)dt\Big|\mathcal{F}_0\right] = G^*(V,Y,\phi).$$

Next, I show $v_0g(y_0,\phi_0) \geq G^*(V_0,Y_0,\phi_0)$. To prove this inequality, note

$$\begin{aligned}
&v_t\mathcal{A}[g;(y_{t^-},\phi_t);c(y_{t^-},\phi_t),\hat{\beta}(y_{t^-},\phi_t)]dt\\
&\quad + v_t\left(c(y_{t^-},\phi_t) - \eta\phi\hat{\beta}(y_{t^-},\phi_t)\frac{c(y_{t^-},\phi_t)^\rho}{\lambda} + y_{t^-}\hat{\beta}(y_{t^-},\phi_t)\eta\phi_t(1-\phi_t)\right)dt\\
&\qquad - v_t^{1-\rho}\hat{\beta}(y_{t^-},\phi_t)dU_t + v_t\left[g(y_{t^-} + v_t^{-\rho}\eta dU_t,\phi_t) - g(y_{t^-},\phi_t)\right]\\
&\qquad\quad = -v_t^{1-\rho}\hat{\beta}(y_{t^-},\phi_t)dU_t + v_t\left[g(y_{t^-} + v_t^{-\rho}\eta dU_t,\phi_t) - g(y_{t^-},\phi_t)\right]\\
&\qquad\quad \leq -v_t^{1-\rho}[\hat{\beta}(y_{t^-},\phi_t) - \eta g_{y^+}(y_{t^-},\phi_t)]dU_t\\
&\qquad\quad \leq 0,
\end{aligned}\tag{A.12}$$

with equality if $dU_t = 0$. The equality follows because of the definition of $c(y,\phi)$ and $\hat{\beta}(y,\phi)$ and the HJB equation (24). The first inequality follows from the concavity of g with respect to y. The last inequality follows because $\hat{\beta}(y_{t^-},\phi_t) \geq \eta g_{y^+}(y_{t^-},\phi_t)$.

Therefore,

$$
\begin{aligned}
&v_0 g(y_0,\phi_0) - \mathrm{E}[e^{-r\tau_n} v_{\tau_n} g(y_{\tau_n},\phi_{\tau_n})|\mathcal{F}_0] \\
&\quad = \mathrm{E}\Big[\int_0^{\tau_n} e^{-rt}\Big[-v_t\mathcal{A}[g;(y_{t^-},\phi_t);c(y_{t^-},\phi_t),\hat\beta(y_{t^-},\phi_t)]dt - v_t(g(y_{t^-}+v_t^{-\rho}\eta dU_t,\phi_t)-g(y_{t^-},\phi_t))\Big]\Big|\mathcal{F}_0\Big] \\
&\quad \geq \mathrm{E}\Big[\int_0^{\tau_n} e^{-rt}\Big[v_t\left(c(y_{t^-},\phi_t)-\eta\phi_t\hat\beta(y_{t^-},\phi_t)\frac{c(y_{t^-},\phi_t)^\rho}{\lambda}+y_{t^-}\hat\beta(y_{t^-},\phi_t)\eta\phi_t(1-\phi_t)\right)dt - v_t^{1-\rho}\hat\beta(y_{t^-},\phi_t)dU_t\Big]\Big|\mathcal{F}_0\Big],
\end{aligned}
$$

with equality if $dU_t = 0$ for all t.

Taking the limit for $n\to\infty$ and using $\lim_{n\to\infty}\mathrm{E}[e^{r\tau_n}v_{\tau_n}g(y_{\tau_n},\phi_{\tau_n})|\mathcal{F}_0]=0$, we obtain

$$
v_0 g(y_0,\phi_0) \geq \mathrm{E}\left[\int_0^\infty e^{-rt}\left[v_t\left(c(y_{t^-},\phi_t)-\eta\phi_t\hat\beta(y_{t^-},\phi_t)\frac{c(y_{t^-},\phi_t)^\rho}{\lambda}+y_{t^-}\hat\beta(y_{t^-},\phi_t)\eta\phi_t(1-\phi_t)\right)dt - v_t^{1-\rho}\hat\beta(y_{t^-},\phi_t)dU_t\right]\Big|\mathcal{F}_0\right],
$$

with equality if $dU_t = 0$ for all t.

Therefore, we conclude

$$
\begin{aligned}
v_0 g(y_0,\phi_0) &= \mathrm{E}\left[\int_0^\infty e^{-rt}v_t\left(c(y_{t^-},\phi_t)-\eta\phi_t\hat\beta(y_{t^-},\phi_t)\frac{c(y_{t^-},\phi_t)^\rho}{\lambda}+y_{t^-}\hat\beta(y_{t^-},\phi_t)\eta\phi_t(1-\phi_t)\right)dt\Big|\mathcal{F}_0\right] \\
&= \sup_{(U_t)_{t\geq 0}\in\mathcal{I}}\mathrm{E}\Big[\int_0^\infty e^{-rt}\Big\{v_t\left(c(y_{t^-},\phi_t)-\eta\phi_t\hat\beta(y_{t^-},\phi_t)\frac{c(y_{t^-},\phi_t)^\rho}{\lambda}+y_{t^-}\hat\beta(y_{t^-},\phi_t)\eta\phi_t(1-\phi_t)\right)dt - v_t^{1-\rho}\hat\beta(y_t,\phi_t)dU_t\Big\}\Big|\mathcal{F}_0\Big] \\
&= \sup_{(U)_{t\geq 0}\in\mathcal{I}}\mathrm{E}\Big[\int_0^\infty e^{-rt}\Big\{\left(C_t^R-\eta\phi_t\beta_t^R\frac{(C_t^R)^\rho}{\lambda}+Y_t\beta_t^R\eta\phi_t(1-\phi_t)\right)dt-\beta_t^R dU_t\Big\}\Big|\mathcal{F}_0\Big].
\end{aligned}
$$

Because the contract $\mathcal{C}^R_{(v_0,y_0,\phi_0)}$ delivers utility $V_0 = \frac{v_0^{1-\rho}}{1-\rho}$ to the agent, we conclude $v_0 g(y_0,\phi_0)\geq G^*(V_0,Y_0,\phi_0)$.

Combining the results so far, we therefore have

$$
G^*(V_0,Y_0,\phi_0) = v_0 g(y_0,\phi_0) = \mathrm{E}\left[\int_0^\infty e^{-rt}v_t\left(c(y_t,\phi_t)-\eta\phi_t\hat\beta(y_t,\phi_t)\frac{c(y_t,\phi_t)^\rho}{\lambda}-y_t\hat\beta(y_t,\phi_t)\eta\phi_t(1-\phi_t)\right)dt\Big|\mathcal{F}_0\right].
$$

Hence, the contract $\mathcal{C}^R_{(v_0,y_0,\phi_0)}$ represents a solution to the dual problem. Because the relaxed optimal contract coincides with the solution to the dual problem when $Y_0 = 0$, we conclude $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is a relaxed optimal contract. □

PROOF OF THEOREM 4(III). Because $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is a relaxed optimal contract, the cost $v_0 g(y_0,\phi_0)$ provides a lower bound for the principal's cost function in an optimal contract. Hence, if the contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is incentive compatible, it is also an optimal contract. I use Theorem 2 to verify whether $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is incentive compatible. Because condition (8) holds by construction in a relaxed optimal contract, it remains to verify whether (12) is satisfied.

By Theorem 3 and Theorem 4(II), the agent's information rent at time t with contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$ is given by $\xi_t = G^*_{Y^+}(V_t,Y_t,\phi_t) = v_t^{1-\rho}g_{y^+}(y_t,\phi_t)$. I then use Ito's lemma and express the volatility of the agent's information rent, ω_t, as

$$
\omega_t = (1-\rho)v_t^{1-\rho}g_{y^+}(y_t,\phi_t)\hat\beta(y_t,\phi_t) - \rho v_t^{1-\rho}y_t g_{yy}(y_t,\phi_t)\hat\beta(y_t,\phi_t) + v_t^{1-\rho}g_{y^+\phi}(y_t,\phi_t)\eta\phi_t(1-\phi_t).
$$

With contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$, condition (12) is thus equivalent to

$$v_t^{1-\rho}\left\{\left[(1-\rho)g_{y^+}(y_t,\phi_t)-\rho y_t g_{yy}(y_t,\phi_t)\right]\hat{\beta}(y_t,\phi_t)+g_{y^+\phi}(y_t,\phi_t)\eta\phi_t(1-\phi_t)-\eta(1-2\phi_t)g_{y^+}(y_t,\phi_t)\right\}\geq 0.$$

After dividing by $v_t^{1-\rho}$, I obtain (25). If (25) holds, the relaxed optimal contract $\mathcal{C}^R_{(v_0,0,\phi_0)}$ satisfies (12) and it is therefore an optimal contract. $\square$

References

Agarwal, V., Daniel, N. D., and Naik, N. Y. (2009). Role of Managerial Incentives and Discretion in Hedge Fund Performance. *Journal of Finance*, 64(5):2221–2256.

Arie, G. (2016). Dynamic costs and moral hazard: A duality-based approach. *Journal of Economic Theory*, 166(166):1–50.

Battaglini, M. (2005). Long-term contracting with markovian consumers. *American Economic Review*, 95(3):637–658.

Battaglini, M. and Lamba, R. (2019). Optimal dynamic contracting: The first-order approach and beyond. *Theoretical Economics*, 14(4):1435–1482.

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

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

Biais, B., Mariotti, T., Plantin, G., and Rochet, J.-C. (2007). Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications. *Review of Economic Studies*, 74(2):345–390.

Biais, B., Mariotti, T., Rochet, J.-C., and Villeneuve, S. (2010). Large Risks, Limited Liability, and Dynamic Moral Hazard. *Econometrica*, 78(1):73–118.

Carrasco, V., Fuchs, W., and Fukuda, S. (2019). From equals to despots: The dynamics of repeated decision making in partnerships with private information. *Journal of Economic Theory*, 182:402–432.

Cisternas, G. (2018). Two-Sided Learning and the Ratchet Principle. *Review of Economic Studies*, 85(1):307–351.

Crandall, M. G., Ishii, H., and Lions, P. L. (1992). User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1):1–67.

DeMarzo, P. M. and Fishman, M. J. (2007). Agency and Optimal Investment Dynamics. *Review of Financial Studies*, 20(1):151–188.

DeMarzo, P. M. and Sannikov, Y. (2006). Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model. *Journal of Finance*, 61(6):2681–2724.

DeMarzo, P. M. and Sannikov, Y. (2017). Learning, Termination, and Payout Policy in Dynamic Incentive Contracts. *Review of Economic Studies*, 84:182–236.

Di Tella, S. (2019). Optimal regulation of financial intermediaries. *American Economic Review*, 109(1):271–313.

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

Dittmann, I., Maug, E., and Spalt, O. (2010). Sticks or carrots? Optimal CEO compensation when managers are loss averse. *Journal of Finance*, 65(6):2015–2050.

Edmans, A. and Gabaix, X. (2011). Tractability in Incentive Contracting. *Review of Financial Studies*, 24(9):2865–2894.

Edmans, A., Gabaix, X., and Jenter, D. (2017). Executive Compensation: A Survey of Theory and Evidence.

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

Fama, E. (1980). Agency problems and the theory of the firm. *Journal of Political Economy*, 88(2):270–282.

Farhi, E. and Werning, I. n. (2013). Insurance and taxation over the life cycle. *Review of Economic Studies*, 80(2):596–635.

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

Garrett, D. F. and Pavan, A. (2012). Managerial Turnover in a Changing World. *Journal of Political Economy*, 120(5):879–925.

Gibbons, R. and Murphy, K. J. (1992). Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence. *Journal of Political Economy*, 100(3):468–505.

Golosov, M., Kocherlakota, N., and Tsyvinski, A. (2003). Optimal Indirect and Capital Taxation. *Review of Economic Studies*, 70(3):569–587.

Graham, J. R., Harvey, C. R., and Puri, M. (2015). Capital allocation and delegation of decision-making authority within firms. *Journal of Financial Economics*, 115(3):449–470.

Guo, Y. (2016). Dynamic Delegation of Experimentation. *American Economic Review*, 106(8):1969–2008.

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

Hall, B. J. and Liebman, J. B. (1998). Are CEOs Really Paid Like Bureaucrats? *Quarterly Journal of Economics*, 113(3):653–691.

He, Z. (2012). Dynamic Compensation Contracts with Private Savings. *Review of Financial Studies*, 25(5):1494–1549.

He, Z., Wei, B., Yu, J., and Gao, F. (2017). Optimal Long-Term Contracting with Learning. *Review of Financial Studies*, 30(6):2006–2065.

Holmstrom, B. (1999). Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies*, 66(1):169–182.

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

Kapička, M. (2013). Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach. *Review of Economic Studies*, 80(3):1027–1054.

Kaplan, S. N. and Strömberg, P. (2003). Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. *Review of Economic Studies*.

Karatzas, I. and Shreve, S. E. (1998). *Brownian Motion and Stochastic Calculus*. Graduate Texts in Mathematics. Springer, New York, NY, second edition.

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

Kocherlakota, N. R. (2005). Zero expected wealth taxes: A Mirrlees approach to dynamic optimal taxation. *Econometrica*, 73(5):1587–1621.

Konlack Socgnia, V. and Menoukeu-Pamen, O. (2015). An infinite horizon stochastic maximum principle for discounted control problem with Lipschitz coefficients. *Journal of Mathematical Analysis and Applications*, 422(1):684–711.

Laffont, J.-J. and Tirole, J. (1990). Adverse Selection and Renegotiation in Procurement. *Review of Economic Studies*, 57(4):597–625.

Li, C. W. and Tiwari, A. (2009). Incentive Contracts in Delegated Portfolio Management. *Review of Financial Studies*, 22(11):4681–4714.

Marcet, A. and Marimon, R. (2019). Recursive Contracts. *Econometrica*, 87(5):1589–1631.

Miao, J. and Zhang, Y. (2015). A duality approach to continuous-time contracting problems with limited commitment. *Journal of Economic Theory*, 159:929–988.

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

Myerson, R. B. (1984). Cooperative Games with Incomplete Information. *International Journal of Game Theory*, 13:69–96.

Myerson, R. B. (1985). Bayesian Equilibrium and Incentive Compatibility: An Introduction. In Hurwicz, L., Schmeidler, D., and Sonnenschein, H., editors, *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, chapter 8. Cambridge University Press, Cambridge, UK.

Øksendal, B. (2003). *Stochastic Differential Equations*. Springer, Berlin, Heidelberg.

Pavan, A., Segal, I., and Toikka, J. (2014). Dynamic Mechanism Design: A Myersonian Approach. *Econometrica*, 82(2):601–653.

Pavoni, N., Sleet, C., and Messner, M. (2018). The Dual Approach to Recursive Optimization: Theory and Examples. *Econometrica*, 86(1):133–172.

Pham, H. (2009). *Continuous-time stochastic control and optimization with financial applications*. Springer-Verlag.

Phelan, C. and Stacchetti, E. (2001). Sequential Equilibria in a Ramsey Tax Model. *Econometrica*, 69(6):1491–1518.

Phelan, C. and Townsend, R. M. (1991). Computing Multi-Period, Information-Constrained Optima. *Review of Economic Studies*, 58(5):853.

Prat, J. and Jovanovic, B. (2014). Dynamic contracts when the agent's quality is unknown. *Theoretical Economics*, 9(3):865–914.

Sannikov, Y. (2008). A Continuous-Time Version of the Principal–Agent Problem. *Review of Economic Studies*, 75(3):957–984.

Sannikov, Y. (2014). Moral Hazard and Long-Run Incentives.

Spear, S. E. and Srivastava, S. (1987). On Repeated Moral Hazard with Discounting. *Review of Economic Studies*, 54(4):599–617.

Williams, N. (2009). On Dynamic Principal-Agent Problems in Continuous Time.

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

Yong, J. and Zhou, X. Y. (1999). *Stochastic Controls*. Springer-Verlag, New York.