

Dynamic Agency with Career Concerns: A Recursive Framework

by

JEREMY BERTOMEU *

October 18, 2005

Abstract

This paper investigates the design of a long-term compensation contract between a principal and an agent in the presence of career concerns and moral hazard. Specifically, the output history must be interpreted simultaneously as a signal on the productivity of the partnership and on the action choices of the agent. I consider a stationary formulation of the problem where in addition to the continuation utility, the beliefs of the principal *and* the agent are used as state variables. A deviation from equilibrium action choices will trigger an incorrect update of the prior of the principal. The continuation utility will correspond to the maximum feasible utility for a given belief of the agent and given that the principal implements a contract based on a possibly wrong prior. The model unifies the existing literature on agency theory and learning.

(JEL: D31, D80, D82)

(Keywords: Dynamic Agency; Learning; Mechanism Design)

The design of a compensation arrangement between a principal and an agent requires a careful examination of the contractual imperfections that may impair efficient risk-sharing. This problem produced two families of models that explain why observed compensation patterns may feature a combination of performance pay, turnover and insurance. The first family, known as matching or learning, postulates that the productivity of the agent is unknown and must be inferred from the output history (Jovanovic 1979). The second family, known as agency theory, assumes that the input of the agent to the production process is unobservable (Holmström 1979) and proper diligence must be elicited through an output-contingent contract. This paper intends to unify both "standard" models into a general framework that incorporates both learning and agency.

*Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh PA 15213. Contact: bertomeu@cmu.edu. Many thanks to Steve Spear and Rick Green for valuable advice. This research was supported by a grant from the William Larimer Mellon fellowship.

Considered in isolation, both families exhibit several limitations. Learning often relies on short-term contracts or asymmetric commitments. These restrictions may seem at odds with common practice. For example, executive compensation exhibits many medium to long-term compensation arrangements such as stock options, retirement benefits, forgivable loans, non-monetary benefits and long-term bonus plans (Murphy 1999). In the context of venture finance, structured loans can mention events that trigger early repayments or changes in the interest rate. At the employee level, firm-specific data exhibits substantial wage variations in real terms (Baker, Gibbs, and Holmström 1994) although the firm should offer downward rigid contracts if unilateral commitments are possible (Harris and Holmström 1982). In response, agency explains why perfect insurance may not be achieved even in the case of unrestricted commitments. However, the theory of turnover that it proposes bears little resemblance with the real world. In a model of agency, personal wealth is by construction the sole determinant of terminations. When wealth is too low, the ability of the principal to punish the agent is reduced and when wealth is too high, the cost of diligence is too high relative to the marginal benefit of compensation (Spear and Wang 2005). A model with moral hazard and learning reconciles performance-pay and terminations with long-term contracts. It also conveys a more realistic set of assumptions for almost any situation involving a contract between a principal and an agent.

The analysis of a contract that jointly considers learning and moral hazard is much more than the sum of its parts. Career concerns raise the attractiveness of diligence because they induce more favorable updates and higher future compensation. The feedback from current performance is known as implicit incentives. A large body of empirical research reports evidence of career concerns in various economic settings: top-executive compensation (Boschen and Smith 1995, Brickley, Coles, and Linck 1999); fund management (Chevalier and Ellison 1999); temporary labor (Engellandt and Riphahn 2005); broker reports (Hong and Kubik 2003). Career concerns also alter the design of a contract that implements particular action choices. Off-equilibrium paths, where the continuation payoff determines whether a given set of actions is incentive-compatible, cannot be represented as a subproblem where the principal and the agent have symmetric information about productivity. Conditional on a deviation from the prescribed actions, the principal will misinterpret output as a signal on productivity given equilibrium actions. The future beliefs of the principal will misrepresent the continuation utility of the agent as well as future action choices. The optimal contract must successfully provide incentives given that implicit incentives may be partial substitutes to explicit incentives but may also induce the agent to disrupt the learning process.

The seminal contribution to the analysis of learning and agency is Holmström (1999). He examines the provision of effort under the constraint that the principal may not offer output-contingent contracts. Career concerns will typically produce effort levels inferior to the optimum in a com-

plete information context. Although it strongly suggests that adopting output-contingent contracts may induce preferable outcomes, the paper does not offer an analysis of compensation packages in the presence of output-contingent pay. Gibbons and Murphy (1992) propose an analysis of the optimal contract given short-term contracts and linear performance pay. The authors show that career concerns are partial substitutes for performance pay. Related work on problems of career concerns with moral hazard include Hirao (1993), Bergemann and Hege (1998) and Bertomeu (2004). In various settings, these papers show that output-contingent pay may be weakened, or even vanish, in the presence of career concerns. The paper is also related to Fernandes and Phelan (2000) and Aysegul and Mukoyama (2005) who examine a multi-period moral hazard contract with persistent shocks and action choices. As in a model with learning, deviations from equilibrium in such models will break down common knowledge.

This paper considers a general model of learning and moral hazard and shows that it can be simplified to a tractable recursive problem. The reduction is applied in three steps. Section 1 presents the history-dependent problem and derives the existence of an optimal contract. Section 2, considers the reduction of the original to a dual static variational problem. In section 3, I show that this recursive formulation can be restated as a contraction.

History-Dependent Problem

The model closely follows the exposition of the dynamic model of agency in Spear and Srivastava (1987). The model of career concerns, where both the agent and the principal are uncertain about the current productivity, is nested in the framework of repeated agency.

Time is discrete, indexed by $t=1,2,\dots$. Production is achieved through the partnership of a principal and an agent whose actions are not observable. The output of the partnership is a random outcome whose distribution is unknown, possibly time-varying and depends on the action of the agent. The action of the agent at time t is denoted $a_t \in A$, where $A \equiv [\underline{a}, \bar{a}]$ is a bounded interval. The output at time t is denoted y_t . At each period, the principal offers a contingent compensation to the agent I_t . The remaining share of the output $y_t - I_t$ corresponds to the compensation of the principal. Let d_t denote an indicator variable corresponding to the occupation choice, equal to one if the partnership has not been terminated yet at date t . Whenever d_t is zero, output will be zero for all future periods. Let $\sigma_t = \{y_1, \dots, y_t\}$ denote a history of outputs until time t . If d_t is one, the current action choice a_t , a sequence of Bayesian updates computed from σ_{t-1} and the previous action choices yield a current prior of the agent $F(y_t|a_t, \sigma_{t-1}, \{a_1, \dots, a_{t-1}\})$.

At each period, the agent consumes c_t and saves $I_t - c_t$. The wealth of the agent at time t is denoted z_t .

At time t , conditional on a realization of y_t , the instantaneous utility of the agent is denoted $H(c_t, a_t, d_t)$, defined for any feasible action, any positive consumption and any occupation choice. Future consumption is discounted at rate β . H is assumed to be separable in d_t : $H(c_t, a_t, d_t) = h(c_t, a_t) + (1 - d_t)h_0$, where $h(c_t, a_t)$ is a continuous function, strictly increasing and concave in c_t , strictly decreasing in a_t and normalized to $h(1, \underline{a}) = 0$. Further, I assume that $h_0 > 0$ so that the minimum level of utility for the agent is zero. The agent and the principal have perfect access to capital markets where savings and borrowings bear an interest rate $(1 - \alpha)/\alpha$. The principal is a firm maximizing the net present value of its contract; it is risk-neutral and discounts the future at rate α , where $0 < \beta \leq \alpha < 1$.

The following timing is considered. At time 1, the true productivity of the partnership is unknown. The principal and the agent share a common prior. The agent chooses an action a_1 and the output y_1 is realized from the distribution $F(y|a_1)$, the distribution of output conditional on a the initial prior and an action choice a_1 . Then, the agent is compensated $I_1(y_1)$ and chooses a consumption c_1 . The remaining share is added to the wealth of the agent so that $z_2 = (z_1 + I_1 - c_1)/\alpha$. At time 2, the principal and the agent decide whether to continue their partnership and set d_2 . Then the agent chooses an action a_2 . The output y_2 is realized from the distribution $F(y|a_2, y_1, a_1)$, the new prior of the agent is updated using the information from (y_1, a_1) and conditional on action a_2 . The agent is compensated $I_2(y_1, y_2)$ and chooses c_2 . The game is extended to the following periods in the same manner.

At each date, the set of possible priors $F(y_t|a_t, \sigma_{t-1}, \{a_1, \dots, a_{t-1}\})$, is restricted to a family of distributions indexed by two real¹ variables $F(y|\rho_t, l_t)$ and such that all priors have a common support S . I assume that for $\rho'_t > \rho_t$, the distribution $F(y|\rho'_t, l_t)$ first-order stochastically dominates $F(y|\rho_t, l_t)$. This parameter is decomposed into $\rho_t = \rho(a_t, m_t)$, where $\rho_a > 0$ and $\rho_m > 0$ and m_t (resp. l_t) take values in Ω_m (resp. Ω_l). The variable m_t represents the current update on the productivity of the partnership. In the model, two different reasons may account for high output. First, better skill corresponds to higher priors and raise the output level. Second, more effort increases output for a given prior. The function ρ translates a given output and a given prior into the distribution of output. Assume that for any (y_t, a_t) , there exists $\theta(y_t, a_t)$ such that $\rho(\theta(y_t, a_t), a_t) = y_t$; $\theta(y_t, a_t)$ is the prior corresponding to a particular output signal y_t considering that an action a_t has been chosen.

The problem will be solved and analyzed for any possible starting prior $(m, l) \in (\Omega_m, \Omega_l)$. I assume that as long as the partnership continues the Bayesian update can be written recursively.

¹This assumption is made for notational tractability; it can be shown that the framework presented here carries over to a larger number of state variables. As we show later, this presentation is chosen to nest various commonly used learning models.

$$m_{t+1} = \mathcal{M}(m_t, \theta(y_t, a_t), l_t) \quad (1)$$

$$l_{t+1} = \mathcal{L}(l_t) \quad (2)$$

Let \mathcal{M} satisfy $\mathcal{M}_m > 0$, $\mathcal{M}_\theta > 0$, $\mathcal{M}_{ml} < 0$ and $\mathcal{M}_{\theta l} > 0$, except possibly at the boundaries of Ω_m and Ω_l if they are attained with zero probability, and assume that \mathcal{L} is an increasing function. Both updating functions are assumed to yield values in their respective domain space. The productivity update is the combination of the previous prior and the new signal disentangled from the action choice. The second variable l_t represents the speed of adjustment of the prior to new information. I assume that the speed of information acquisition is a deterministic process independent of the output sequence. Let $\pi(y; a, m, l)$ denote the probability distribution $F(y|\rho(a, m), l)$.

The model nests various cases of learning previously studied in the literature. Suppose for example that $y|\rho, l$ follows a normal distribution with mean ρ and precision l . If $\rho(m, a) = m + a$ and $\mathcal{M}(m, \theta, l) = (lm + h_\epsilon \theta)/(l + h_\epsilon)$, where $h_\epsilon > 0$, the model corresponds to the usual normal learning model presented in Holmström (1999). If $\mathcal{L}(l) = l + h_\epsilon$, the partnership learns on a constant productivity parameter from noisy observations with precision h_ϵ . In the more general case where $\mathcal{L}(l) = h_\delta(l + h_\epsilon)/(h_\delta + l + h_\epsilon)$, where $h_\delta > 0$, the productivity follows a random walk whose innovations have a precision h_δ . Similarly, the model can be adapted to the case when the output process may take only two values. Suppose now that $y \in \{0, 1\}$ and $y|\rho, l$ follows a Bernoulli random variable with mean ρ . Next, set $\rho(m, a) = ma$, $\mathcal{M}(m, \theta, l) = (ml + \theta)/(l + \theta)$ and $\mathcal{L}(l) = l + 1$. Assume that $\Omega_m = [0, 1]$. In the special case where $A = \{1\}$, the model corresponds to pure learning from a Beta prior studied in Bertomeu (2004). Using the framework developed here, the model can be generalized to the case $A = [0, 1]$.

We now describe a contract. Denote $\{\mathbf{I}, \mathbf{a}, \mathbf{c}, \mathbf{z}, \mathbf{d}\} = \{I(\sigma_t), a(\sigma_{t-1}), c(\sigma_t), z(\sigma_{t-1}), d(\sigma_{t-1})\}$ a sequence of output-contingent compensation, action, consumption and termination choices². At each date, the decision depends only on current information and not on the actions chosen by the agent. Let $\pi(\sigma_{t+\tau}; \sigma_t, \mathbf{a}, m, l)$ denote the probability distribution of $\sigma_{t+\tau}$ given σ_t for an initial prior (m, l) . The distribution $\pi(\sigma_t; \mathbf{a}, m, l)$ denotes the probability distribution of σ_t at time 1.

with:
$$d\pi(\sigma_{t+\tau}; \sigma_t, \mathbf{a}, m, l) = d\pi(y_{t+\tau}; a_{t+\tau}, m_{t+\tau}, l_{t+\tau})d\pi(\sigma_{t+\tau-1}; \sigma_t, \mathbf{a}, m, l) \quad (3)$$

$$d\pi(\sigma_{t+1}; \sigma_t, \mathbf{a}, m, l) = d\pi(y_t; a_t, m_t, l_t) \quad (4)$$

$$m_{t+\tau} = \mathcal{M}(m_{t+\tau-1}, \theta(y_{t+\tau-1}, a_{t+\tau-1}), l_{t+\tau-1}) \quad (5)$$

$$l_{t+\tau} = \mathcal{L}(l_{t+\tau-1}) \quad (6)$$

²In short-hand, the variables are indexed by t : $I(\sigma_t) = I_t, a(\sigma_{t-1}) = a_t, c(\sigma_t) = c_t, z(\sigma_{t-1}) = z_t, d(\sigma_{t-1}) = d_t$.

For a particular initial prior (m, l) and a positive constant W , the optimal contract is defined as a solution to the following problem.

$$(P_L) \quad \max \sum_{t=1}^{\infty} \alpha^{t-1} \int \{d(\sigma_{t-1})y_t - I(\sigma_t)\} d\pi(\sigma_t; \mathbf{a}, m, l)$$

s.t. (5), (6)

$$\sum_{t=1}^{\infty} \beta^{t-1} \int H(c(\sigma_t), a(\sigma_{t-1}), d(\sigma_{t-1})) d\pi(\sigma_t; \mathbf{a}, m, l) \geq W \quad (7)$$

and: $\forall \sigma_t$

$$d(\sigma_t) \in \{0, d(\sigma_{t-1})\} \quad (8)$$

$$\mathbf{a} \in \arg \max_{\tilde{\mathbf{a}}} \sum_{\tau=1}^{\infty} \beta^{\tau-1} \int H(c(\sigma_{t+\tau}), \tilde{a}(\sigma_{t+\tau-1}), d(\sigma_{t+\tau-1})) d\pi(\sigma_{t+\tau}; \sigma_t, \tilde{\mathbf{a}}, m, l) \quad (9)$$

$$\lim_{t \rightarrow \infty} \beta^t \int z(\sigma_{t-1}) d\pi(\sigma_t; \mathbf{a}, m, l) = 0 \quad (10)$$

Although the model allows for returns from experience, for example if \mathcal{M} includes a trend, I assume that growth that may originate from experience is bounded by the interest rate on capital markets. The following sufficient condition guarantees that the net present value of the project remains finite. For any \mathbf{a} , for any (m, l) ,

$$\sum_{t=1}^{\infty} \alpha^{t-1} \int \max(h_0, y_t) d\pi(\sigma_t; \mathbf{a}, m, l) \text{ is finite.}$$

The Pareto frontier is obtained as the set of contracts that maximize the utility of the principal under constraints (5), (6), (7), (8), (9) and (10). The productivity of the partnership is updated at each period according to the updating rule (5) and (6). Equation (8) prescribes permanent terminations. I assume that learning stops whenever the partnership is not operating. Therefore, once termination occurs once, the prior will remain fixed and thus the termination decision will remain optimal for all future periods. Equation (7) is an individual rationality constraint. The lifetime utility of the agent must be greater than a lifetime reservation W . Equation (9) is the incentive-compatibility constraint. The action sequence \mathbf{a} announced to the principal must maximize the utility of the agent so that truth-telling remains a best response for all possible output histories. Equation (10) is a transversality condition. The future debt of the agent must have a zero net present value.

In the next Lemma, I show that in the case of long-term contracts presented in (P_L) , the model can be simplified by ignoring the consumption and savings choices. The principal will integrate the consumption and savings choices to the compensation contract.

Lemma 1. *Suppose that $\{\mathbf{I}, \mathbf{a}, \mathbf{c}, \mathbf{z}, \mathbf{d}\}$ is optimal and define \mathbf{I}' such that $I'_1 = c_1 - z_1$ and for $t > 1$, $I'_t = c_t$ and $z'_t = 0$. Then, the contract $\{\mathbf{I}', \mathbf{a}, \mathbf{c}, \mathbf{z}', \mathbf{d}\}$ is also optimal.*

Following Lemma 1, I will ignore the consumption and savings choices in the determination of the long-term contract. An optimal long-term contract will be defined as a sequence $\{\mathbf{I}, \mathbf{a}\}$ such that wealth remains zero and the agent consumes the compensation at each period. Payments after a termination correspond to pure consumption and savings choices of the agent. Alternatively, the agent will be offered a parachute or retirement payment equal to the net present value of the consumption stream.

In the standard formulation of dynamic agency, the same incentive problem is repeated over several periods. After a single deviation, the contracting problem is analogous to the initial problem, in that it can be formulated as a problem of the form (P_L) . Off-equilibrium branches following a deviation from the prescribed action path will be analogous to branches on the equilibrium. This property has two consequences. First, the contract $\{\mathbf{I}, \mathbf{a}\}$ proposed by the principal will remain on the Pareto frontier. Second, the truth-telling strategy will remain optimal even if the agent deviated once. The argument breaks down in the presence of learning. Following a deviation at date t from \mathbf{a} to $\tilde{\mathbf{a}}$, the incentive mechanism in Equation (9) will be designed for the distributions $\{d\pi(\sigma_{t+\tau}; \sigma_t, \mathbf{a}, m_1, l_1)\}$ and not $\{d\pi(\sigma_{t+\tau}; \sigma_t, \tilde{\mathbf{a}}, m_1, l_1)\}$. Further, since off-equilibrium actions are not anticipated by the principal, the evaluation of the continuation utility of the agent may be biased. The new subproblem characterizing the off-equilibrium branches will be a contracting choice between the principal and the agent given that the agent has superior information on the distribution of output.

For $W \in \mathcal{R}$, let $\mathbf{V}(W, m, l, e)$ be the set of all feasible expected utilities of the principal when W is promised to the agent and the initial prior is (m, l) when $e = 1$, whereas the partnership is closed at time 1 when $e = 0$.

$$\mathbf{V}(W, m, l, e) \equiv \{V \in \mathcal{R} \mid \exists \{\mathbf{I}, \mathbf{a}, \mathbf{d}\} \text{ s.t. (5)-(9) are verified and } d_1 \leq e\}$$

The following Proposition guarantees that an optimal contract exists.

Proposition 1. *If $\exists \hat{I}$ s.t. $H(\hat{I}, \underline{a}, 0) \geq W(1 - \beta)$, then there exists a solution to (P_L) .*

Before stating the recursive problem, it will be useful to introduce an accessory function to measure the utility of the agent off the equilibrium path. Suppose that the current prior of the principal is (m, l) and fix $v \in \mathbf{V}(W, m, l, e)$, a lifetime utility level for the principal. Denote $\{\mathbf{I}^*, \mathbf{a}^*, \mathbf{d}^*\}$ the optimal policy in (P_L) that yields a utility v to the principal, under the constraint

that $d_1 \leq e$. Consider the problem of finding the maximum utility that can be achieved by the agent given that the true prior is (\tilde{m}, l) and the principal implements $\{\mathbf{I}^*, \mathbf{a}^*, \mathbf{d}^*\}$.

$$(\tilde{P}_L) \quad \max_{\tilde{\mathbf{a}}} \sum_{t=1}^{\infty} \beta^{t-1} \int H(I(\sigma_t), \tilde{a}(\sigma_{t-1}), d(\sigma_{t-1})) d\pi(\sigma_t; \tilde{\mathbf{a}}, \tilde{m}, l) \quad (11)$$

Define the following set:

$$\mathbf{W}(v, m, \tilde{m}, l, e) \equiv$$

$$\{w \in \mathcal{R} \mid \exists \tilde{\mathbf{a}} \text{ s.t. } w = \sum_{t=1}^{\infty} \beta^{t-1} \int H(I^*(\sigma_t), \tilde{a}(\sigma_t), d^*(\sigma_{t-1})) d\pi(\sigma_t; \tilde{\mathbf{a}}, \tilde{m}, l)\}$$

$\mathbf{W}(v, m, \tilde{m}, l, e)$ represents the set of utilities that can be attained by the agent given that the prior of the principal is (m, l) but the prior of the agent is (\tilde{m}, l) .

The next Corollary will be stated only for a set of states (v, m, \tilde{m}, l, e) that can be attained with non-zero probability starting from the original problem (m_1, l_1, W_1) . In the problem at hand, any set that can be attained with zero probability is irrelevant for the purpose of the analysis.

Corollary 1. *There exists a solution to (\tilde{P}_L) .*

The optimum of (\tilde{P}_L) is denoted $U^*(v, m, \tilde{m}, l, e)$. Note that the special case $U^*(v, m, m, l, e)$ is the optimum of (P_L) .

Reduction to a Stationary Problem

The history-dependent will be rewritten in terms of the stationary analogue of the dual problem. The constraints (5)-(10) must first be adapted to the stationary analogue for the dual of the history-dependent problem. Suppose that at a certain date t , the principal expects a continuation utility v , the current prior of the principal is (m, l) , the current prior of the agent is (\tilde{m}, l) and the status of the partnership is an indicator variable e , equal to one if the partnership has not been terminated. The state variables of the stationary problem are (v, m, \tilde{m}, l, e) . Denote $d(v, m, l, e)$, the employment decision at date t .

$$d(v, m, l, e) \in \{0, e\} \quad (12)$$

Equation (12) correspond to the analogue of Equation (8). If termination has occurred prior to the current period, then the partnership cannot be operated. Denote $a(v, m, \tilde{m}, l, e)$, the action chosen by the agent. The principal will assume common knowledge of both distributions and thus anticipate a choice $a(v, m, m, l, e)$.

$$m' = \mathcal{M}(m, \theta(y, a(v, m, m, l, e)), l) \quad (13)$$

$$l' = \mathcal{L}(l) \quad (14)$$

The updates of the principal are stated in Equations (13) and (14), where (m', l') is the new prior of the principal. In Equation (13), the update will follow the current prior m and the anticipated action choice $a(v, m, m, l, e)$. They allow the principal to modify the optimal contract to track variations in the productivity of the partnership. In comparison, the agent will update the distribution with respect to the correct action choice $a(v, m, \tilde{m}, l, e)$ and the correct prior \tilde{m} .

$$\tilde{m}' = \mathcal{M}(\tilde{m}, \theta(y, a(v, m, \tilde{m}, l, e)), l) \quad (15)$$

Equation (15) states the updating rule for the agent, where \tilde{m}' is the new prior of the agent. The recursive structure of the problem must make sure to be internally consistent in that the principal obtains the level of utility v prescribed as a state variable. This condition is also known as the promise-keeping constraint. Denote $I(v, m, l, y, e)$, the contingent compensation of the agent.

$$\int \{d(v, m, l, e)y - I(v, m, l, y, e) + \beta V(v, m, l, y, d(v, m, l, e))\} d\pi(y; a', m, l) \geq v \quad (16)$$

where: $a' \equiv a(v, m, m, l, e)$.

The promise-keeping constraint is stated in Equation (16). Note that the utility promised to the principal is measured here with respect to the equilibrium prior. The principal needs only believe that the promise-keeping is satisfied although it may not be off equilibrium. The contract chosen by the principal must be a solution to the optimal contract given common knowledge of the distribution of output. Similarly, the action anticipated by the principal must a best response to a problem with common knowledge. Next, $U(v, m, \tilde{m}, e)$ will denote the maximum utility that can be achieved by the agent for a set of states (v, m, \tilde{m}, e) . Define also $V(v, m, l, y, e)$, the utility received by the principal when, after a state (v, m, m, l, e) , output y has been observed; V is the continuation utility of the principal and corresponds to future payments. For a particular contract $\{V, I, a, d\}$, define the following function:

$$U(V, I, a, d; v, m, \tilde{m}, l, e) \equiv$$

$$\int \{H(I(v, m, l, y, e), \tilde{a}, d(v, m, l, e)) + \beta U(v, m', \tilde{m}', d(v, m, l, e))\} d\pi(y; \tilde{a}, \tilde{m}, l)$$

$$\begin{aligned}
\text{where: } \quad \tilde{a} &= a(v, m, \tilde{m}, l, e) \\
m' &= \mathcal{M}(m, \theta(y, a(v, m, \tilde{m}, l, e))), l) \\
\tilde{m}' &= \mathcal{M}(\tilde{m}, \theta(y, a(v, m, \tilde{m}, l, e))), l)
\end{aligned}$$

$U(V, I, a, d; v, m, \tilde{m}, l)$ is the utility achieved by the agent given a contract (V, I, a, d) , for a given set of states (v, m, \tilde{m}, l) .

$$\forall (\mathbf{v}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e}) \quad a(\mathbf{v}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e}) \in \arg \max U(V, I, a, d; \mathbf{v}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e}) \quad (17)$$

The temporary incentive-compatibility constraint (17) is written with respect to the beliefs of the principal. It states that the action path announced by the agent must also maximize the utility of the agent in the absence of the promise-keeping constraint of the principal; that is, it must coincide with the choices of the agent for any possible states. Incentive-compatibility takes into account the profit that can be achieved when the agent deviates to a different action profile and common knowledge breaks down.

$$\forall (\mathbf{v}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e})$$

$$(V(\mathbf{v}, \mathbf{m}, \mathbf{l}, y, \mathbf{e}), I(\mathbf{v}, \mathbf{m}, \mathbf{l}, y, \mathbf{e}), d(\mathbf{v}, \mathbf{m}, \mathbf{l}, \mathbf{e})) \in \arg \max U(V, I, a, d; \mathbf{v}, \mathbf{m}, \mathbf{m}, \mathbf{l}, \mathbf{e}) \quad (18)$$

The continuation utility, the current compensation and the termination choices must be an optimal contract solution for the current equilibrium beliefs of the principal (18), that is given that the principal believes that the current prior is (m, l) and the announced action path is effectively realized. After a deviation, the agent will be compensated according to the wrong updates and the contract may no longer be on the Pareto frontier.

The continuation utility of the agent is computed with respect to the correct prior $d\pi(y; \tilde{a}, \tilde{m}, l)$ and is maximized over the current action of the agent \tilde{a} . In the complete program, the agent must maximize the feasible utility anticipating that the learning process of the principal is disrupted off the equilibrium branch. Define Ω_Γ as the set where the constraints will have a feasible solution.

$$\Omega_\Gamma = \{(v, m, \tilde{m}, l, e) \text{ s.t. } v \in \mathbf{V}(0, m, l, e), (m, \tilde{m}, l, e) \in \Omega_m \times \Omega_m \times \Omega_l \times \{0, 1\}\}$$

Define an operator Γ mapping the set of continuous functions from Ω_Γ to \mathcal{R} bounded with the Sup Norm into itself as follows:

$$\begin{aligned}
\Gamma(U)(v, m, \tilde{m}, l, e) &= \max \int \{H(\tilde{a}, I(v, m, l, y, e), d(v, m, l, e)) \\
&\quad + \beta U(V(v, m, l, y, d(v, m, l, e)), m', \tilde{m}', l', y, d(v, m, l, e))\} d\pi(y; \tilde{a}, \tilde{m}, l) \\
\text{s.t.} \quad &\quad (12) - (18) \text{ and } \tilde{a} = a(v, m, \tilde{m}, l, e)
\end{aligned}$$

Note that on the equilibrium path, when $\tilde{m} = m$, the choice of the optimal contract can be reincorporated as the arguments of the optimization problem and Equations (15) and (18) can be taken out of the program.

$$\begin{aligned} \Gamma(U)(v, m, m, l, e) = \max \int \{ & H(I(v, m, l, y, e), a(v, m, l, e), d(v, m, l, e)) \\ & + \beta U(V(v, m, l, y, d(v, m, l, e)), m', \tilde{m}', l', y, d(v, m, l, e)) \} d\pi(y; a(v, m, \tilde{m}, l, e), \tilde{m}, l) \\ \text{s.t. } & (12) - (14) \quad (16) - (17) \end{aligned}$$

The following Proposition shows that this stationary formulation delivers an optimal contract for (P_L) .

Proposition 2. $U^*(v, m, \tilde{m}, l, e)$ is a fixed point of Γ .

Although the program cannot be shown easily to be a contraction, which will be shown later by reformulating the problem, it is checked in the next Corollary that for any possible starting “guess” of U , the resulting $\Gamma(U)$ remains finite.

Corollary 2. $\Gamma(U)(v, m, \tilde{m}, l, e)$ is bounded.

Proposition 2 shows that an optimal contract is given by five functions $\{U, V, I, a, d\}$ that measure respectively the continuation utility of the agent, the continuation utility of the principal, the current compensation, the current action choice and the current termination decision. As in dynamic agency models, the function I measures the current compensation of the agent and V represents future payments to the principal. The function a corresponds to the effort policy and d measures the termination decision. The main difference with Spear and Srivastava (1987) is that the deviation utility must be explicitly considered when measuring the utility of the agent and the announced actions. Although the principal will not observe the deviation, the utility achievable from the deviation can be inferred from the equilibrium contract. Implicit incentives require the principal to consider more carefully the utility of the agent at each step of the contract.

Reduction to a Contraction

The reformulation of (P_L) as a recursive static optimization problem allows to introduce a recursive procedure to derive the optimal contract from a sequence of approximations. Specifically, the recursive problem presented here can be reformulated as a contraction. In order to proceed further, it will be necessary to convert the dual problem considered previously into a problem where the utility of the principal is maximized.

Given the transformation of the program, an optimal contract is given by four functions $\{\bar{U}, \bar{V}, \bar{I}, \bar{a}, \bar{d}\}$, taking now the utility of the agent as a state variable. The constraints (12)-(18) are now transformed into the following constraints:

$$\bar{d}(w, m, l, e) \in \{0, e\} \quad (19)$$

$$m' = \mathcal{M}(m, \theta(y, \bar{a}(w, m, m, l, e)), l) \quad (20)$$

$$\tilde{m}' = \mathcal{M}(\tilde{m}, \theta(y, \bar{a}(w, m, \tilde{m}, l, e)), l) \quad (21)$$

$$l' = \mathcal{L}(l) \quad (22)$$

$$\int \{H(\bar{I}(w, m, l, y, e), \bar{a}(w, m, m, l, y, e), \bar{d}(w, m, l, y, e)) + \beta \bar{U}(w, m, l, y, \bar{d}(w, m, l, e))\} d\pi(y; a, m, l) \geq w \quad (23)$$

$$\forall (\mathbf{w}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e}) \quad \bar{a}(\mathbf{w}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e}) \in \arg \max \int \{H(\bar{I}(\mathbf{w}, \mathbf{m}, \mathbf{l}, y, \mathbf{e}), \tilde{a}, \bar{d}(\mathbf{w}, \mathbf{m}, \mathbf{l}, y, \mathbf{e})) + \beta \bar{U}(\mathbf{w}, \mathbf{m}', \tilde{\mathbf{m}}', \mathbf{l}', d(\mathbf{w}, \mathbf{m}, \mathbf{l}, \mathbf{e}))\} d\pi(y; \tilde{a}, \tilde{\mathbf{m}}, \mathbf{l}) \quad (24)$$

$$\text{s.t.} \quad \tilde{\mathbf{m}}' = \mathcal{M}(\tilde{\mathbf{m}}, \theta(y, \tilde{a}), l)$$

$$\mathbf{m}' = \mathcal{M}(\mathbf{m}, \theta(y, a(\mathbf{w}, \mathbf{m}, \mathbf{m}, \mathbf{l}, \mathbf{e})), l)$$

$$\mathbf{l}' = \mathcal{L}(\mathbf{l})$$

Define $\Omega_{\bar{\Gamma}} \equiv \mathcal{R}^+ \times \Omega_m \times \Omega_l \times \{0, 1\}$ and define an operator $\bar{\Gamma}$ mapping the set of continuous functions from $\Omega_{\bar{\Gamma}}$ to \mathcal{R} bounded with the Sup Norm into itself as follows:

$$\bar{\Gamma}(V)(w, m, l, e) = \max \int \{\bar{d}(w, m, l, e)y - \bar{I}(w, m, l, y, e)$$

$$+ \alpha(V(\bar{U}(w, m, m, l, y, e), m, l, y, \bar{d}(v, m, l, e)), m', \tilde{m}', l', y, \bar{d}(v, m, l, e))\} d\pi(y; a', m, l)$$

$$\text{s.t.} \quad (19) - (24) \text{ and } a' = a(v, m, m, l, e)^3$$

where: $\forall (\mathbf{w}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e})$

$$U(\mathbf{w}, \mathbf{m}, \tilde{\mathbf{m}}, \mathbf{l}, \mathbf{e}) = \max_{\tilde{a}} \int \{H(\tilde{a}, \bar{I}(\mathbf{w}, \mathbf{m}, \mathbf{l}, y, \mathbf{e}), d(\mathbf{w}, \mathbf{m}, \mathbf{l}, \mathbf{e}))$$

$$+ \beta \bar{U}(\mathbf{w}, \mathbf{m}', \tilde{\mathbf{m}}', \mathbf{l}', d(\mathbf{w}, \mathbf{m}, \mathbf{l}, \mathbf{e}))\} d\pi(y; \tilde{a}, \tilde{\mathbf{m}}, \mathbf{l}) \quad (25)$$

$$\begin{aligned}
\text{s.t. } \tilde{\mathbf{m}}' &= \mathcal{M}(\tilde{\mathbf{m}}, \theta(y, \tilde{a}), \mathbf{l}) \\
\mathbf{m}' &= \mathcal{M}(\mathbf{m}, \theta(y, a(\mathbf{w}, \mathbf{m}, \mathbf{m}, \mathbf{l}, \mathbf{e})), \mathbf{l}) \\
\mathbf{l}' &= \mathcal{L}(\mathbf{l})
\end{aligned}$$

This direct formulation of the contraction $\bar{\Gamma}$ corresponds to the standard recursive formulation of dynamic agency. The utility of the principal is maximized under the constraint that it provides a certain level of utility to the agent (23) and achieves incentive-compatibility (24). As previously argued, this formulation also introduces an accessory function (25). It is clear from Corollary 1 that the program considered in (25) is well-defined. This function corresponds to the problem initially studied in the dual $U(v, m, \tilde{m}, l, e)$ except that the utility of the agent is now used as a state variable. Two more results are needed to conclude the analysis of the optimal contract. First, I will show that this new operator is a contraction and thus a fixed point exists and may be recovered by repeated iterations of the operator. Second, I will show that Pareto-frontier obtained with this operator corresponds to the Pareto-frontier of the dual problem. Combining the two results, it follows that the optimal contract can be recovered from the recursive problem. The next Proposition is a straightforward application of Blackwell's sufficient conditions for a contraction.

Proposition 3. $\bar{\Gamma}$ is a contraction.

Proposition 3 suggests a convenient algorithm to approximate the unique fixed point of the operator $\bar{\Gamma}$.

- (1) Start with a guess of $V(w, m, l, e)$; for example one could use the solution to the standard dynamic agency problem for a given starting and invariant prior (m, l) .
- (2) Make a second guess for $U(w, m, \tilde{m}, l, e)$; for example, a convenient guess may be:
 $U(w, m, \tilde{m}, l, e) = U(w, m, m, l, e)$.
- (3) Given these two functions, solve for the optimal policy functions $\{\bar{I}, a, d\}$ ignoring constraint (25), where a is obtained only for values of the form $a(w, m, m, l, e)$.
- (4) Plug in these policies into program (25), solve for the remaining $a(w, m, \tilde{m}, l, e)$ and update the function $U(w, m, \tilde{m}, l, e)$.
- (5) Until $U(w, m, \tilde{m}, l, e)$ has converged, return to Step (3). When it has converged, update $V(w, m, l, e)$ and return to Step (3). Stop when both functions have converged.

Define

$X(m, l, e) \equiv \{(U^*(v, m, m, l, e), v)\}$. It is the Pareto-frontier for each possible initial choice of the prior. Define $\bar{X}(m, l, e) \equiv \{w, V^*(w, m, l, e)\}$, where V^* is the fixed point of $\bar{\Gamma}$.

Proposition 4. $X(m, l, e) = \bar{X}(m, l, e)$

Proposition 4 shows that this new formulation, shown to be a contraction, can be substituted for the dual.

Conclusion

I analyze a multi-period moral hazard that incorporates a central feature of many real-life contracts, learning about the quality of the partnership from the observed sequence of outputs. The standard formulation of the dynamic agency problem must be reformulated. Future updates on the distribution of outputs will depend on the action choices announced by the agent to the principal. Whenever, the agent deviates to a different action, the principal will follow a contract that will not correspond to the current distribution of output. The incentive-compatibility is rewritten by considering the continuation utility of the agent of any possible continuation utility of the agent. The optimal contract must be solved considering the continuation utility but also the beliefs of the agent and the possibly wrong beliefs of the principal.

The purpose of this paper was to provide a simple representation of the optimal contract and present a framework to answer several open research questions. To what extent do career concerns explicitly reduce the need for explicit incentives, such as current and future pay-for-performance? Until now, this assertion is proved only for very restrictive learning dynamics. Does the presence of career concerns dampen performance-pay more for low histories? If the contract features excess upside when performance is low, the agent may misreport the effort and thus enter off-equilibrium branches knowing that the principal is underestimating skill. Will the dismissal decision be characterized in terms of not only wealth but also current performance? Given career concerns, the continuation utility will no longer be a sufficient statistic to decide over terminations. To what extent does the optimal long-term contract allow to attain a higher and more efficient level of effort than a model without contracts (Holmström 1999)? How does renegotiation-proofness constraints affect the performance-pay and the effort chosen in the optimal contract? Renegotiation-proofness reduces the insurance chosen by the principal, it will also affect the off-equilibrium continuation utility. Will the resulting contract be transient, in the sense that for a sufficiently long history either the agent or the principal obtains the whole surplus? With perpetual learning (as in Holmström (1999)), the production technology will change through time and thus may also imply changes in the allocation of the surplus.

Appendix:

Proof of Lemma 1: The constraints (7)-(10) are obviously verified by $\{\mathbf{I}', \mathbf{a}, \mathbf{c}, \mathbf{z}', \mathbf{d}\}$. The contract is optimal if it yields the same utility to the principal as $\{\mathbf{I}, \mathbf{a}, \mathbf{c}, \mathbf{z}, \mathbf{d}\}$. For $T > 1$, denote $V(m, l, T)$ (resp. $V'(m, l, T)$), the utility of the principal up to period T in the optimal (resp. new) contract.

$$\begin{aligned}
V'(m, l, T) &= \sum_{t=1}^T \alpha^{t-1} \int \{d(\sigma_{t-1})y_t - I'(\sigma_t)\} d\pi(\sigma_t; \mathbf{a}, m, l) \\
&= z_1 + \sum_{t=1}^T \alpha^{t-1} \int \{d(\sigma_{t-1})y_t - (I(\sigma_t) - \alpha z(\sigma_t) + z(\sigma_{t-1}))\} d\pi(\sigma_t; \mathbf{a}, m, l) \\
&= V(m, l, T) + z_1 + \sum_{t=1}^T \alpha^{t-1} \int \{\alpha z(\sigma_t) - z(\sigma_{t-1})\} d\pi(\sigma_t; \mathbf{a}, m, l) \\
&= V(m, l, T) + \alpha^T \int z(\sigma_T) d\pi(\sigma_T; \mathbf{a}, m, l)
\end{aligned}$$

Letting T go to infinity, by (10), the two contracts yield the same utility to the principal. \square

Proof of Proposition 1: The case $e = 0$ is obvious. Suppose $e = 1$. If $\mathbf{V}(W, m, l, 1)$ is nonempty (i), bounded from above (ii) and closed (iii), then a solution always exists. (i) The set is nonempty. Set $I_t = \hat{I}$, $a_t = \underline{a}$ and $d_t = 0 \forall t$. The principal will achieve a utility $\hat{I}/(1 - \alpha)$, finite. Clearly, the actions are incentive-compatible and yield W to the agent. Therefore, $\hat{I}/(1 - \alpha) \in \mathbf{V}(W, m, l, 1)$. (ii) To prove boundedness, consider the production technology $\tilde{y}_t = \max(h_0, y_t)$ and ignore for now constraints (7) and (9). This new problem will necessarily yield more utility to the principal. Further, the optimal strategy of the principal will be to set $d_t = 1$ and $I_t = 0$. The upper bound on the utility of the principal becomes

$$\sum_{t=1}^{\infty} \alpha^{t-1} \int \max(h_0, y_t) d\pi(\sigma_t; \mathbf{a}, m, l)$$

This expression was assumed earlier to be finite. (iii) Consider now a sequence of contracts $\{\mathbf{I}^k, \mathbf{a}^k, \mathbf{d}^k\}$ verifying (5)-(10), and such that the principal obtains $\{V^k\}$, a sequence converging to V^∞ . Note that a similar argument to (ii) implies that if the principal receives a finite utility, then the maximum utility achievable by the agent must be finite as well. Then at a neighborhood of V^∞ , the utility achieved by the agent must be bounded by some value \overline{W} . Consider the value of \underline{V} that yields \overline{W} in (i). Any $\{\mathbf{I}, \mathbf{a}, \mathbf{d}\}$ such that the utility of the principal falls below \underline{V} is clearly suboptimal and can be ignored. Therefore, any optimal contract must satisfy:

$$I(\sigma_t) \leq \alpha^{-t} \left(\sum_{t=1}^{\infty} \alpha^{t-1} \int y_t d\pi(\sigma_t; \overline{a}, m_1, l_1) - \underline{V} \right)$$

The set \mathbf{V} can therefore be restricted to compensations such that for each t , $I(\sigma_t)$ is bounded. It must be shown that there exists a contract $\{\mathbf{I}^\infty, \mathbf{a}^\infty, \mathbf{d}^\infty\}$ verifying (7) and (9) and such that the principal receives

V^∞ . The contract is constructed as follows. For each $t \geq 1$, $\{I_t^k\}_{k=1}^\infty$, $\{a_t^k\}_{k=1}^\infty$ and $\{d_t^k\}_{k=1}^\infty$ are bounded sequences. Therefore each sequence has a convergent subsequence. This limit is denoted $\{I_t^\infty\}$, $\{a_t^\infty\}$ and $\{d_t^\infty\}$. The resulting contract is in \mathbf{V} and generates V^∞ . \square

Proof of Proposition 1: Note first that whenever $e = 0$, (\tilde{P}_L) becomes a simple consumption smoothing problem. Next suppose that $e = 1$. If $\tilde{m} = m$, by incentive-compatibility (9) in (P_L) , the actions \mathbf{a} will be a solution to (\tilde{P}_L) and therefore (\tilde{P}_L) can be written as the dual of (P_L) . By monotonicity of the objective of the agent and the principal, as long as there exists a solution in (P_L) delivering v to the principal, there exists a solution in the dual. Consider the case $\tilde{m} \neq m$. Clearly the set is nonempty. Suppose now it is unbounded for a set $\Omega_d \equiv (v, m, \tilde{m}, l) \in \mathcal{R} \times \Omega_m \times \Omega_m \times \Omega_l$. By assumption, in the original problem (m_1, l_1, W_1) , there exists a deviation strategy such that Ω_d is reached with non-zero probability and therefore, there exists t_0 such that the set is reached with non-zero probability at t_0 . Once the set is reached, the optimal contract continued from (m_1, l_1, W_1) must coincide with the optimal contract starting from date t_0 . Therefore, the agent would also achieve infinite deviation utility in the original contract (m_1, l_1, W_1) thus contradicting incentive-compatibility. To prove closedness, consider a sequence of strategies $\{\tilde{\mathbf{a}}^k\}$ that yields W^k to the agent and such that W^k converges to W^∞ . For each t , a deviation strategy \tilde{a}_t^k is constructed as the limit of a converging subsequence of $\{\tilde{a}_t^k\}_{k=1}^\infty$. Clearly, this strategy $\{\tilde{\mathbf{a}}^\infty\}$ must yield W^∞ to the agent. \square

Proof of Proposition 2: For a particular value of (v, m, \tilde{m}, l, e) , $\Gamma(U^*)(v, m, \tilde{m}, l, e)$ must be shown to be smaller (i) and greater (ii) than $U^*(v, m, \tilde{m}, l, e)$.

(i) In the program associated to $\Gamma(U^*)(v, m, \tilde{m}, l, e)$, denote $V^R(v, m, l, y, e)$, $I^R(v, m, l, y, e)$ and $a^R(v, m, \tilde{m}, l, e)$, the solutions corresponding to $\Gamma(U^*)(w, m, \tilde{m}, l, e)$. A feasible contract will be constructed for (\tilde{P}_L) . Set first: $d_1^L = d^R(v, m, l, e)$, $a_1^L = a^R(v, m, \tilde{m}, l, e)$ and $I_1^L(y_1) = I^R(v, m, l, y, e)$. In the following period, the contract a_2^L , I_2^L and d_2^L is constructed as:

$$\begin{aligned} d_2^L &= d^R(V^R(v, m, m, l, y_1, e), m', l', d_1^L) \\ a_2^L(\sigma_1) &= a^R(V^R(v, m, m, l, y_1, e), m', \tilde{m}', l', d_1^L) \\ I_2^L(\sigma_2) &= I^R(V(v, m, m, l, y_1, e), m', l', y_2, d_1^L) \end{aligned}$$

where: $m' = \mathcal{M}(m, \theta(y_1, a^R(v, m, m, l, e)), l)$, $\tilde{m}' = \mathcal{M}(\tilde{m}, \theta(y_1, a_1^L), l)$ and $l' = \mathcal{L}(l)$. A contract $\{\mathbf{a}^L, \mathbf{I}^L, \mathbf{d}^L\}$ is then constructed recursively in the same manner. Suppose at date t , the surplus expected by the principal is $\mathcal{V}(\sigma_{t-1})$, the current employment status is $d(\sigma_{t-1})$, the current prior of the principal is (m_t, l_t) and the current prior of the agent is (\tilde{m}_t, l_t) . The contract $a^L(\sigma_{t-1})$, $I^L(\sigma_t)$ and $d^L(\sigma_{t-1})$ is constructed as follows:

$$\begin{aligned} I^L(\sigma_t) &= I^R(\mathcal{V}(\sigma_{t-1}), m_t, l_t, y_t, d(\sigma_{t-1})) \\ a^L(\sigma_{t-1}) &= a^R(\mathcal{V}(\sigma_{t-1}), m_t, \tilde{m}_t, l_t, d(\sigma_{t-1})) \\ d^L(\sigma_{t-1}) &= d^R(\mathcal{V}(\sigma_{t-1}), m_t, l_t, d(\sigma_{t-1})) \end{aligned}$$

Finally, the following variables are updated:

$$\begin{aligned}
m_{t+1} &= \mathcal{M}(m_t, \theta(y_t, a^R(\mathcal{V}(\sigma_{t-1}), m_t, m_t, l_t, d(\sigma_{t-1}))), l_t) \\
\tilde{m}_{t+1} &= \mathcal{M}(\tilde{m}_t, \theta(y_t, a^L(\sigma_{t-1})), l_t) \\
l_{t+1} &= \mathcal{L}(l_t) \\
V(\sigma_t) &= I^R(\mathcal{V}(\sigma_{t-1}), m_t, l_t, y_t, d(\sigma_{t-1}))
\end{aligned}$$

The contract $\{\mathbf{I}^L, \mathbf{a}^L, \mathbf{d}^L\}$ is readily verified to be feasible in (\tilde{P}_L) .

(ii) The case $e = 0$ is trivial, so I will consider the case $e = 1$. Consider an optimal contract for (\tilde{P}_L) $\{\mathbf{a}^L, \mathbf{I}^L, \mathbf{d}^L\}$ such that the principal obtains a utility v and the initial prior of the principal is (m, l) and the prior of the agent is (\tilde{m}, l) . The following contract is feasible in $\Gamma(U^*)(v, m, \tilde{m}, y, 1)$. First, I^R , d^R and a^R are set to: $I^R(v, m, l, y, 1) = I_1^L(y)$, $d^R(v, m, l, 1) = d_1^L$ and $a^R(v, m, \tilde{m}, l, 1) = a_1^L$. Finally, the continuation utility $V^R(v, m, l, y, 1)$ is set to:

$$V^R(v, m, l, y, 1) = \sum_{t=2}^{\infty} \alpha^{t-2} \int \{d^L(\sigma_{t-1})y_t - I^L(\sigma_t)\} d\pi(\sigma_t; \sigma_1, \mathbf{a}^L, m, l)$$

Clearly, these functions satisfy the constraints of the recursive problem. \square

Proof of Corollary 2: Applying Jensen's inequality in the program of the agent:

$$\Gamma(U)(v, m, \tilde{m}, l, e) \leq H(\int I(v, m, l, y, e) d\pi(y; \tilde{a}, \tilde{m}, l), a(v, m, l, e), d(v, m, l, e))$$

$$+\beta \int U(V(v, m, l, y, d(v, m, l, e)), m', \tilde{m}', l', y, d(v, m, l, e)) d\pi(y; \tilde{a}, \tilde{m}, l)$$

The expected wage can be bounded from Equation (16):

$$\int I(v, m, l, y, e) d\pi(y; a(v, m, l, e), m, l) \leq$$

$$\int (d(v, m, l, e)y + \beta V(v, m, l, y, d(v, m, l, e))) d\pi(y; a, m, l) - v$$

On the right hand-side, the output and the continuation utility of the principal must be bounded Therefore $\Gamma(U)(v, m, m, l, e)$ is bounded. Next, from similar arguments as Corollary 1,

$\Gamma(U)(v, m, \tilde{m}, l, e)$ can be shown to be bounded as well for $\tilde{m} \neq m$. \square

Proof of Proposition 4: Suppose that $(w, v) \in X(m, l, e)$, then I will show that there exists (w', v') in $\bar{X}(m, l, e)$ such that $w' \geq w$ and $v' \geq v$. To do so, it is necessary to construct a contract that verifies (19)-(24), yields v to the principal and at least w to the agent. In the dual, there exists a contract $\{V, I, a, d\}$, where the utility of the principal v is used as a state variable, such that the principal receives v and the agent receives w . For any $\mathbf{w}, \mathbf{m}, \mathbf{l}$ and \mathbf{e} , define $\mathbb{V}(\mathbf{w}, \mathbf{m}, \mathbf{l}, \mathbf{e})$, the unique utility level such that $(\mathbf{w}, \mathbb{V}(\mathbf{w}, \mathbf{m}, \mathbf{l}, \mathbf{e})) \in X(\mathbf{m}, \mathbf{l}, \mathbf{e})$. I will now proceed constructively to obtain the fixed point $V^*(w, m, l, e)$ of the new problem.

Define the following function:

$$\bar{U}(w, m, l, y, e) \equiv$$

$$U^*(v', \mathcal{M}(m, \theta(y, a(v', m, l, e)), l), \mathcal{M}(m, \theta(y, a(v', m, l, e)), l), \mathcal{L}(l), d(v', m, l, e))$$

where: $v' = V(\nabla(w, m, l, e), m, l, y, e)$.

And:

$$\bar{I}(w, m, l, y, e) \equiv I(U^*(\nabla(w, m, l, e), m, l, e), m, l, y, e)$$

$$\bar{a}(w, m, l, e) \equiv a(U^*(\nabla(w, m, l, e), m, l, e), m, l, y, e)$$

$$\bar{d}(w, m, l, e) \equiv d(U^*(\nabla(w, m, l, e), m, l, e), m, l, y, e)$$

This contract is the desired solution:

$$\bar{V}^*(w, m, l, e) \equiv \int \{\bar{d}(w, m, l, e)y - \bar{I}(w, m, l, y, e)$$

$$+ \alpha V(\nabla(w, m, l, e), m, l, y, \bar{d}(w, m, l, e))\} d\pi(y; \bar{a}(w, m, l, e), m, l)$$

By a similar argument, it is easy to recover the fixed point of the dual from the new problem. \square

References

- AYSEGUL, S., AND T. MUKOYAMA (2005): “Repeated Moral Hazard with Persistence,” Economic Theory, 25(4), 831–854.
- BAKER, G., M. GIBBS, AND B. HOLMSTRÖM (1994): “The Wage Policy of a Firm,” Quarterly Journal of Economics, 109(4), 921–955.
- BERGEMANN, D., AND U. HEGE (1998): “Venture capital financing, moral hazard, and learning,” Journal of Banking and Finance, 22(6-8), 703–735.
- BERTOMEU, J. (2004): “CEO Compensation: Agency or Learning ?,” Working Paper.
- BOSCHEN, J. F., AND K. J. SMITH (1995): “You Can Pay Me Now and You Can Pay Me Later: The Dynamic Response of Executive Compensation to Firm Performance,” Journal of Business, 68(4), 577–608.
- BRICKLEY, J. A., J. L. COLES, AND J. S. LINCK (1999): “What Happens to CEOs After They Retire? New Evidence on Career Concerns, Horizon Problems, and CEO Incentives,” Journal of Financial Economics, 52(3), 341–377.
- CHEVALIER, J., AND G. ELLISON (1999): “Career Concerns of Mutual Fund Managers,” Quarterly Journal of Economics, 114(2), 389–432.
- ENGELLANDT, A., AND R. T. RIPHAHN (2005): “Temporary Contracts and Employee Effort,” Labor Economics, 12(3), 281–299.
- FERNANDES, A., AND C. PHELAN (2000): “A Recursive Formulation for Recursive Agency with History Dependence,” Journal of Economic Theory, 91(2), 223–247.

- GIBBONS, R., AND K. J. MURPHY (1992): "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," Journal of Political Economy, 100(3), 468–505.
- HARRIS, M., AND B. HOLMSTRÖM (1982): "A Theory of Wage Dynamics," Review of Economic Studies, 49(3), 315–333.
- HIRAO, Y. (1993): "Learning and Incentive Problems in Repeated Partnerships," International Economic Review, 34(1), 101–119.
- HOLMSTRÖM, B. (1979): "Moral Hazard and Observability," Bell Journal of Economics, 10(1), 74–91.
- (1999): "Managerial Incentive Problems: A Dynamic Perspective," Review of Economic Studies, 66(1), 169–182.
- HONG, H., AND J. D. KUBIK (2003): "Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts," Journal of Finance, 58(1), 313–351.
- JOVANOVIĆ, B. (1979): "Job Matching and the Theory of Turnover," Journal of Political Economy, 87(5), 972–990.
- MURPHY, K. J. (1999): Executive Compensation, vol. 3 of Handbook of labor economics. North Holland, Ashenfelter, Orley and Card, David.
- SPEAR, S. E., AND S. SRIVASTAVA (1987): "On Repeated Moral Hazard with Discounting," Review of Economic Studies, 54(4), 599–617.
- SPEAR, S. E., AND C. WANG (2005): "When to Fire a CEO: Optimal Termination in Dynamic Contracts," Journal of Economic Theory, 120(2), 239–256.
- WANG, C. (1997): "Incentives, CEO Compensation, and Shareholder Wealth in a Dynamic Agency Model," Journal of Economic Theory, 76, 72–105.