Moral Hazard and Long-Run Incentives*

Yuliy Sannikov

October 1, 2013

Abstract

This paper considers dynamic moral hazard settings, in which the agent's actions have consequences over a long horizon. To maintain incentives, the optimal contract defers the agent's compensation and ties it to future performance. Some of the agent's compensation is deferred past termination. Termination occurs when the value of deferred compensation becomes insufficient to maintain adequate incentives. The target pay-performance sensitivity provided by deferred compensation builds up during the agent's tenure, but decreases after termination.

1 Introduction.

This paper studies dynamic agency problems, in which the agent's actions affect future outcomes. These situations are common in practice. CEO's actions have long-term impact on firm performance. The success of private equity funds is not fully revealed until they sell their illiquid investments. The quality of mortgages given by a broker is not known until several years down the road. There has been a lot of informal discussion of these situations. The issues of deferred compensation and clawback provisions come up frequently.

*This is a working (not final) draft of the paper. Comments are welcome, and can be e-mailed to sannikov@gmail.com. I am grateful to Andy Skrzypacz, whose insights and ideas have greatly helped this paper. I would also like to thank Bruno Strulovici and seminar participants at UCLA, Harvard, IAS and TSE for helpful comments and the Alfred P. Sloan Foundation for financial support.

However, it has been difficult to design a formal tractable framework to analyze these issues.

In my model, the agent continuously puts in effort for the duration of employment and continuously consumes a compensation flow. Current effort affects observable output over the entire future. I consider all history-dependent contracts, including contracts that make the agent's pay after termination contingent on the firm's performance history, and characterize the optimal contract.

The optimal contract has interesting features. First, the agent's pay-performance sensitivity generally grows towards a target with tenure. With time, the agent has greater opportunity to affect the project, and so the optimal contract exposes the agent to more project risk. Second, while gains and losses affect the value of the agent's deferred compensation immediately, they become reflected in payments only slowly. Rewards for good performance are banked in to offset potential future losses. Third, the agent's limited liability constraint places a bound on pay-performace sensitivities and the provision of incentives. Following bad performance, the agent's employment is terminated. However, in this case the agent does not receive his deferred compensation immediately. Rather, deferred compensation is tied to performance and is paid out gradually even after termination.

For the case of managerial compensation, the qualitative features of the optimal contract can be implemented through an incentive account. The firm requires the agent to hold stock in a deferred incentive account, and lends him money towards a target stake. Account balance affects the agent's flow compensation, e.g. the agent may be allowed to receive a percentage of surplus, or be required to pay a percentage of account deficit. For example, a CEO whose target annual pay is \$3 million may start with an account that contains \$100 million in stock held against a \$100 million loan (so the account balance is 0). If the stock drops by 20%, then the account balance drops by \$20 million. The agent may be responsible for only 5% of the balance per year: a deduction of \$1 million from his annual pay. Then the agent receives \$2 million in the first year, but he is on the hook to keep covering the shortfall in the future. If the stock recovers, the agent is paid more next year, but if the stock continues falling, the agent may be fired. Critically, even after termination, the stock in the account continues vesting, and the agent may receive some money if the account recovers. This feature ensures that the agent has some incentives even before he is fired.

Let me make several remarks about this implementation. First, the risk in the incentive account is significant, and expected market return of the firm's stock is by no means sufficient to compensate the agent for this risk. Therefore, the firm may want to add compensating transfers to the account (e.g. \$5 million in year 1). Second, the target level of stock in the account may vary over time: my model suggests that it has to reflect the opportunity to agent had to affect the current stock return. In particular, it makes sense to require lower stock holdings initially (less than \$100 million in year 1), but gradually raise the required risk exposure towards a target. One way to adjust required risk exposure is through *rebalancing*, a concept introduced by Edmans, Gabaix, Sadzik and Sannikov (2012). Importantly, new stock is not given to the CEO for free, but rather against a loan provided by the firm: the account balance does not change due to rebalancing. Third, target pay may grow over time to reflect the backloading of the agent's payments. This would be particularly relevant if the agent can employ hidden savings to self-insure against future risk exposure.[1]

Formally, the *principal's problem* of finding the optimal contract is a constrained optimization problem. The objective is the principal's profit, or firm value net of the cost of compensating the agent. The agent's incentives are crucial to determining the principal's profit. Therefore, the first part of this paper analyzes the agent's incentives in an arbitrary contract. When the agent's actions affect future outcomes, then his current incentives depend on the sensitivity of his future pay not only to current, but also future performance.

In a fully history-dependent contract, the agent's compensation c_t at time t can depend on the entire history of past performance from time 0 to time t. When choosing effort at time t, the agent will take into account how effort affects performance (e.g. stock return) at each point of time $t+s$ in the future, and how his compensation is tied to performance. On the margin, the agent's incentives are determined by his *information rent* Φ_t, the derivative of his continuation value with respect to unobservable fundamentals, which he can affect with effort.

[1]In this case, the Euler equation requires that the drift of the agent's consumption is positive whenever his utility function is CRRA. If the agent's relative risk aversion is γ, then the drift of his consumption has to be

$$\mu^c = \frac{\gamma + 1}{2} (\sigma^c)^2$$

when σ^c is the volatility of consumption.

Under the assumption that the impact of the agent's effort on future outcomes is exponentially decaying, the principal's problem, subject to only the agent's first-order incentive constraints, reduces to an optimal stochastic control problem with two state variables: Φ_t and the agent's continuation value W_t. These state variables represent the principal's commitments to the agent regarding the expected utility of future compensation and the expected exposure risk. The principal must honor these commitments. As long as the principal accounts properly for his commitments, W_t and Φ_t are "sufficient statistics" that summarize the agent's payoff and incentives. If the principal replaces the agent's continuation contract with another contract that has the same values of W_t and Φ_t, then the agent does not wish he had chosen effort differently in the past (at least on the margin). Of course, absent commitment the contract would not be renegotiation proof: after the agent has sunk effort expecting strong incentives Φ_t, the contracting parties are tempted to renegotiate and lower the agent's risk exposure, reducing Φ_t.

In the setting of Sannikov (2008), in which the agent's effort affects only current output, the principal's control problem has only one state variable, W_t. In contrast, Φ_t is no longer a state variable but a control: the principal directly controls the agent's incentives by setting the sensitivity of W_t to current performance. In that setting, which is a special case of the model in this paper, the agent's incentives are not interlinked across time. In contrast, when the agent's effort affects future outcomes, then performance at time t reflects the agent's effort at all earlier times, and so exposure to performance at time t affects the agent's incentives in all earlier periods. The interlinked incentives across time require a second state variable, Φ_t.

The control problem simplifies greatly in the space of Lagrange multipliers (adjoint variables), ν_t and λ_t on W_t and Φ_t. Multiplier ν_t is the inverse of the agent's marginal utility. Variable λ_t determines the volatility of ν_t, and thus the volatility of the agent's flow compensation.[2] The law of motion of the variable λ_t is *slow*: λ_t has only drift and no volatility. During the agent's employment, λ_t is adjusted towards a target level, which depends on ν_t. After termination, λ_t decays exponentially towards 0 at the rate which depends on the agent's impact on future outcomes.

One component of the optimal contract, the map from λ_t to the volatility

[2]If the agent's utility function has relative risk aversion γ at current consumption level, and if the volatility of the multiplier ν_t is $x\%$, then the volatility of the agent's flow compensation is $x/\gamma\%$.

of the agent's compensation, is determined explicitly up front. To determine two other components: the drift of λ_t and the boundary where the agent's employment is terminated, one has to solve a partial differential equation. I also identify a special case, in which the target level of λ_t is constant, and so the law of motion of λ_t is determined explicitly as well. This cases arises in the limit when the signal-to-noise ratio of the agency problem converges to 0, while the benefits of exposing the agent to some risk persist. This case roughly corresponds to the CEO managing a very large firm.

The impact of the horizon over which the agent's actions affect output can be studied by varying κ, the exponential decay rate of the impact of effort on future output, while keeping the expected present value generated by effort fixed. Numerically, I find interesting results. First, the principal's profit is not very sensitive to κ, as long as κ is much larger than the discount rate r. That is, under the optimal contract, it does not matter much whether the information about the agent's effort is observed immediately, or with delay of one or two years. Second, the target volatility of ν_t, at the target value of λ_t, is not very sensitive to parameter values. However, under the optimal contract, the rate at which λ_t approaches the target does depend on κ : it is roughly proportional to κ. Third, one can construct an approximately optimal contract by borrowing the target level of λ_t from the contract of Sannikov (2008) (where $\kappa = \infty$), which can be found by solving an ordinary differential equation. Even for $\kappa = 0.4$, the approximately optimal contract is within 1%-2% of the optimal contract profit. These facts highlight how a simple standard benchmark can be used to understand a much more complicated setting with delayed impact of the agent's effort on output.

While the agent's contract is found by focusing on the first-order incentive constraints, I also find a simple sufficient condition that can be checked ex-post, or imposed ex-ante to derive a robust contract. The condition is a bound on the sensitivity of Φ_t to performance, called Γ_t. Intuitively if Φ_t changes quickly with performance, then the agent's incentives may change fast enough as he deviates to actually make it profitable to deviate further. In contrast, contracts in which Γ_t satisfies the bound are robust. I show analytically that the bound holds in settings with low signal-to-noise ratio, and can verify the bound in other settings. Many applications, such as that of executive compensation, naturally have a low signal-to-noise ratio.

This paper is organized as follows. In Section 2 lays out a basic model. Section 3 analyzes the agent's incentives on the margin. Section 4 characterizes the optimal contract for the *large firm* case, in which it is feasible to give the agent only a small portion of firm equity. Section 5 tackles the case where the impact of the agent's actions on future outcomes decays exponentially. It provides a sufficient second-order condition to guarantee the optimality of the agent's strategy, and characterizes the optimal contract using a variant of the stochastic maximum principle.

Literature Review. This paper is related to the literatures on dynamic contracts and executive compensation. Papers such as Radner (1985), Spear and Srivastava (1987), Abreu, Pearce and Stacchetti (1990) and Phelan and Townsend (1991) provide foundations for the analysis if repeated principal-agent interactions. In these settings, the agent's effort affects the probability distribution of a signal observed by the principal in the same period, and the optimal contract can be presented in a recursive form. That is, in these settings the agent's continuation value completely summarizes his incentives. Using the recursive structure, Sannikov (2008) provides a continuous-time model of repeated agency, in which it is possible to explicitly characterize the optimal contract using an ordinary differential equation.

The model of Sannikov (2008) is a special case of the model in this paper, as the agent's effort affects the probability of outcomes in the future. That is, the agent's current effort affects firm's unobservable fundamentals, which have impact on future cash flows. To summarize incentives, one also has to keep track of the derivative of the agent's payoff with respect to fundamentals, sometimes referred to as the agent's *information rents* (see Kwon (2012), Pavan, Segal and Toikka (2012), Garrett and Pavan (2012) and Eso and Szentes (2013)). This leads to the so-called first-order approach, which has been used recently to analyze a number of environments. Kapicka (2011) and Williams (2011) use the first-order approach in environments where the agent has private information. DeMarzo and Sannikov (2011) and He, Wei and Yu (2012) study environments with learning, where the agent's actions can affect the principal's belief about fundamentals. In general, first-order conditions do not guarantee full incentive compatibility, which has to be verified ex-post (as in Werning (2002) and Farhi and Werning (2012)). This paper provides a different simpler approach to check full incentive compatibility, through a restriction on the contract space. Contracts that happen to satisfy the restriction are fully incentive compatible. When the first-order approach fails, the restriction can be used to construct robust approximately optimal contracts that are fully incentive compatible.

A few papers have looked at what happens when the agent's effort is observed with delay from specific angles. Hopenhayn and Jarque (2010) consider a setting where the agent's one-time effort input affects output over a long horizon. See also Jarque (2011). Likewise, in Varas (2012) the information about a single project is revealed gradually. Edmans, Gabaix, Sadzik and Sannikov (2012) (in a scale-invariant setting) and Zhu (2012) (in a setting where first-best is attainable) allow the agent to manipulate performance over a limited time horizon and do not allow for termination.

One especially attractive feature of this paper is the closed-form characterization of the optimal contract in environments with large noise. Such a clean characterization is rare in contracting environments. Holmstrom and Milgrom (1987) derive a linear contract for a very particular model with exponential utility. Edmans, Gabaix, Sadzik and Sannikov (2012) obtain a tractable contract in a scale-invariant setting. In contrast, we consider a setting that allows for general utility function and for termination.

This paper is also related to literature on managerial compensation. The model predicts that the agent's pay-performance sensitivity under the optimal contract increases gradually during employment. This is consistent with empirical evidence documented by Gibbons and Murphy (1992). At the same time, the model also suggests that some of CEO's compensation should be deferred after termination, a feature observed rarely in practice. DeMarzo and Sannikov (2006) and Biais, Mariotti, Plantin and Rochet (2007) study managerial compensation in the optimal contracting framework with a risk-neutral agent, but allow the agent's actions to have only contemporaneous effect on cash flows. In these settings, it is also optimal to defer some of the agent's compensation, but only until the time of termination. Deferred compensation creates more room to punish the agent in the future in case of bad performance. Backloaded compensation also helps employee retention, a point first made by Lazear (1979).

Do managerial incentives matter? Yes, according to Adam Smith, who writes, "managers rather of other peoples money... it cannot well be expected, that they should watch over it with the same anxious vigilance with which the partners in a private copartnery frequently watch over their own." Empirically, it is hard to design a clean test that points to the benefit of managerial incentives, but many papers point to the fact. Ang, Cole and Lin (2000) study small private firms, and find that the ratio of expenses to sales is lower, and the ratio of sales to assets is higher, for firms with better-aligned incentives (i.e. when the primary owner has a greater equity stake, and when the manager is a shareholder). Jensen and Murphy (1990) estimate that for public companies, the wealth of a typical CEO goes up by \$3.25 for each \$1000 of shareholder value created. Morck, Shleifer and Vishny (1988) find that there are efficiency benefits to management ownership, as measured by Tobin's Q, at least up to a 5% stake (with mixed results thereafter, which the authors attribute to entrenchment).

2 The Model.

Consider a continuous-time environment, in which the agent's action $a_s \in [0, \bar{a}]$ at time $s \geq 0$ adds $f(t-s)a_s$ to output at time $t \geq s$. The output flow is given by

$$dX_t = \mu_t \, dt + \sigma \, dZ_t, \quad \text{where} \quad \mu_t = \int_0^t f(t-s) a_s \, ds \tag{1}$$

and Z_t is a Brownian motion. Thus, effort a_s at time s adds $f(t-s)a_s$ to output flow at each time $t \in [s, \infty)$. To fix the present value of output generated by unit effort, function $f : [0, \infty) \to [0, \infty)$ is assumed to satisfy

$$\int_0^\infty e^{-rt} f(t) \, dt = 1, \tag{2}$$

where r is the discount rate.

Consider contracts that specify the agent's compensation flow $c_t \geq 0$ at time $t \geq 0$ and termination time $\tau \leq \infty$, as functions of the output history $\{X_s, s \in [0, t]\}$. The agent chooses effort a_t from time 0 until time τ, and receives utility flow

$$u(c_t) - h(a_t)$$

until time τ and

$$u(c_t)$$

thereafter. The utility of consumption $u : [0, \infty) \to [0, \infty)$ and cost of effort $h : [0, \bar{a}] \to [0, \infty)$ are C^2 functions that satisfy $u(0) = 0$, $u' > 0$, $u'' < 0$, $h(0) = 0$, $h'(0) = 0$ and $h'' > 0$.

Employment termination is irreversible. If the agent leaves, the principal continues receiving output dX_t, which is influenced by the past effort of the outgoing agent as well as the effort of the new agent. Assume that the expected present value from hiring the new agent at time τ (net of the costs of compensating the new agent) is L, so that the principal's total expected profit is[3]

$$F_0 = E^a\left[\int_0^\infty e^{-rt}\,(dX_t - c_t\,dt)\right] =$$

$$E^a\left[\int_0^\tau e^{-rt}a_t\,dt + e^{-r\tau}L - \int_0^\infty e^{-rt}c_t\,dt\right], \qquad (3)$$

where E^a is the expectation given the agent's strategy a.

The optimal contract (c,τ) together with the agent's effort strategy a have to maximize (3) subject to a set of constraints: the participation constraint

$$W_0 = E^a\left[\int_0^\infty e^{-rt}(u(c_t) - 1_{t\le\tau}\,h(a_t))\,dt\right], \qquad (4)$$

where W_0 is the agent's required utility at time 0, and a set of *incentive constraints*

$$W_0 \ge E^{\hat{a}}\left[\int_0^\infty e^{-rt}(u(c_t) - 1_{t\le\tau}\,h(\hat{a}_t))\,dt\right] \qquad (5)$$

for all alternative strategies $\hat{a}$. For simplicity, assume that the agent's outside option after time 0 is zero, and that after termination the agent simply consumes payments from the principal.

Solving the *principal's problem* of maximizing (3), subject to (4) and (5), involves finding not only the optimal contract (c,τ) but also the agent's optimal strategy a under this contract, since the strategy enters both the objective function and the constraints.

The principal's problem has the agent's optimization problem embedded in it. To solve it, we analyze the agent's problem first in Section 3, in order to understand the relationship between contract design and the agent's effort. This analysis allows us to reduce the principal's problem to an optimal stochastic control problem. The principal can control the agent's effort by designing the agent's incentives appropriately. We analyze the principal's problem in Sections 4 and 5.

[3] I assume that the effort of the new agent can be inferred given his contract, and can be filtered out. Thus, the new agent adds no noise to the signal about the old agent's effort.

2.1 The Exponential Case.

A particular case involves exponentially decaying impact of the agent's effort on future output, i.e.

$$f(t) = (r + \kappa)e^{-\kappa t},$$

where $\kappa > 0$ is a constant. If so, then the expected rate of output μ_t satisfies

$$\mu_0 = 0, \quad d\mu_t = (r + \kappa)a_t\, dt - \kappa\mu_t\, dt. \tag{6}$$

As $\kappa \to \infty$, function f converges to the Dirac delta function. This leads to the standard principal-agent model, in which the agent's effort adds only to *current* output. In this case, the output is given by

$$dX_t = a_t\, dt + \sigma\, dZ_t, \tag{7}$$

and the model becomes identical to that studied in Sannikov (2008).

3 First-Order Incentive Constraints.

Before investigating the principal's problem, it is important to understand the agent's incentives first. This section focuses on the agent's incentives on the margin, and identifies a key variable Φ_t that can be used to express necessary first-order conditions for the agent's strategy to be optimal. This paper addresses sufficient incentive conditions later.

Before presenting formal results, let me summarize the key findings of this section and draw parallels to the standard case, in which output follows (7). Any contract can be thought of as an option that pays continuously in the units of the agent's utility, rather than money. The value of the option is the agent's *agent's continuation value*,

$$W_t \equiv E_t^a \left[\int_t^\infty e^{-r(s-t)} u_s\, ds \right], \quad \text{where} \quad u_s \equiv u(c_s) - 1_{s \le \tau}\, h(a_s). \tag{8}$$

In the standard case, it is the sensitivity of W to X (which is analogous to option *Delta*, its sensitivity to the underlying)

$$\Delta_t \equiv \frac{dW_t}{dX_t}$$

that determines the agent's incentives at time t. Sannikov (2008) shows that the agent's effort strategy a maximizes the agent's utility if, after all histories, a_t maximizes

$$\Delta_t a_t - h(a_t).$$

In contrast, when effort a_t has impact on future output given by the function $f(s)$, then, as shown below, the agent's incentives depend on

$$\Phi_t \equiv E_t^a \left[\int_t^\infty e^{-r(s-t)} f(s-t) \Delta_s \, ds \right], \tag{9}$$

and a_t has to maximize

$$\Phi_t a_t - h(a_t).$$

For the exponential case, Φ_t can be interpreted as the agent's *information rent*: the derivative of his payoff with respect to the accumulated stock of effort

$$\Phi_t = \frac{dW_t}{dA_t}, \quad \text{where} \quad A_t \equiv \int_0^t e^{-\kappa(t-s)} a_s \, ds = \frac{\mu_t}{r+\kappa}.$$

In the standard case, in which f is the Dirac delta function, (9) reduces to $\Phi_t = \Delta_t$.

There are two convenient expressions for the key process Φ_t. Besides (9), there is an alternative equivalent expression (see Proposition 3) that is derived by looking at the effect of the agent's actions on the probability distribution over future paths. Girsanov Theorem implies that the effect of effort a_t on the likelihood of the path $\{X_s, s \in [t, t']\}$ is given by

$$\zeta_t^{t'} = \int_t^{t'} f(s-t) \frac{dX_s - \mu_s \, ds}{\sigma^2}. \tag{10}$$

Higher effort a_t increases the likelihood of those paths, for which the increments dX_s exceed their expected trend. Expression (10) integrates over future increments, with weight $f(s-t)$, to determine if higher effort would make a given path $\{X_s, s \in [t, t']\}$ more or less likely.

The following expression for Φ_t is equivalent to (9):

$$\Phi_t = E_t^a \left[\int_t^\infty e^{-r(s-t)} \zeta_t^s u_s \, ds \right]. \tag{11}$$

Process Φ_t is useful for analyzing the principal's problem using methods from optimal stochastic control. However, unlike in standard case, where the principal's problem has a single state variable W_t and $\Phi_t = \Delta_t$ is a control, now both W_t and Φ_t are state variables.

I will present relevant formal statements below, before moving on.

3.1 Formal Statements.

Proposition 1 identifies the first-order conditions for a strategy a to be optimal under a contract (c, τ).

Proposition 1 *A necessary condition for a strategy a_t to be optimal under the contract (c, τ) is that for all t,*

$$a_t \ \textit{maximizes} \ \Phi_t a - h(a), \tag{12}$$

where Φ_t is defined by (11).

From now on, denote by $a(\Phi_t)$ the effort that solves the problem (12). Then, $a(h'(a)) = a$ for $a \geq 0$, and $a(\Phi) = 0$ for $\Phi \leq h'(0)$.

Proof. To identify the first-order incentive-compatibility constraint, consider a deviation away from the strategy a towards an alternative strategy $\hat{a}$. Formally, for $\epsilon \in [0, 1]$, let the strategy $(1-\epsilon)a + \epsilon\hat{a}$ assign effort $(1-\epsilon)a_t + \epsilon\hat{a}_t$ to each history of output $\{X_s,\ s \in [0, t]\}$. Then, if the strategies a and $\hat{a}$ lead to expected output rates of μ_t and $\hat{\mu}_t$ respectively, then the strategy $\{(1-\epsilon)a + \epsilon\hat{a}\}$ leads to the expected output rate of $(1-\epsilon)\mu_t + \epsilon\hat{\mu}_t$.

If the agent follows the strategy $(1-\epsilon)a + \epsilon\hat{a}$ rather than a, then by Girsanov's Theorem, he changes the underlying probability measure over output paths by the relative density process

$$\xi_t(\epsilon) = \exp\left(-\frac{1}{2}\int_0^t \frac{\epsilon^2(\hat{\mu}_s - \mu_s)^2}{\sigma^2}\, ds + \int_0^t \frac{\epsilon(\hat{\mu}_s - \mu_s)}{\sigma}\, \frac{dX_s - \mu_s\, ds}{\sigma}\right),$$

where $(dX_s - \mu_s\, ds)/\sigma$ represents increments of a Brownian motion under the strategy a, and $\epsilon(\hat{\mu}_s - \mu_s)/\sigma$ is the rate at which the agent's deviation changes the drift of this Brownian motion.

The agent's utility from deviating to the strategy $(1-\epsilon)a + \epsilon\hat{a}$ is

$$E^a\left[\int_0^\infty e^{-rt}\xi_t(\epsilon)(u(c_t) - 1_{t\leq\tau}\, h((1-\epsilon)a_t + \epsilon\hat{a}_t))\, dt\right]. \tag{13}$$

We would like to differentiate this expression with respect to ϵ at $\epsilon = 0$. Note that

$$\left.\frac{d\xi_t(\epsilon)}{d\epsilon}\right|_{\epsilon=0} = \int_0^t \frac{\hat{\mu}_s - \mu_s}{\sigma} \frac{dX_s - \mu_s\, ds}{\sigma} = \int_0^t (\hat{a}_s - a_s)\zeta_s^t\, ds, \tag{14}$$

where the last equality is obtained using

$$\hat{\mu}_t - \mu_t = \int_0^t f(t-s)(\hat{a}_s - a_s)\, ds.$$

changing the order of integration, and then using the definition (10) of ζ_s^t.

The derivative of the agent's utility (13) with respect to ϵ, at $\epsilon = 0$, is

$$E^a\left[\int_0^\infty e^{-rt}\left(\left.\frac{d\xi_t(\epsilon)}{d\epsilon}\right|_{\epsilon=0}\right)(u(c_t) - 1_{t\le\tau}\, h(a_t))\ dt - \int_0^\tau e^{-rt}(\hat{a}_t - a_t)h'(a_t)\, dt\right] =$$

$$E^a\left[\int_0^\tau e^{-rt}(\hat{a}_t - a_t)\left(\int_t^\infty e^{-r(s-t)}\zeta_t^s u_s\, ds - h'(a_t)\right)\ dt\right] =$$

$$E^a\left[\int_0^\tau e^{-rt}(\hat{a}_t - a_t)(\Phi_t - h'(a_t))\, dt\right], \tag{15}$$

where the second line is obtained using (14) and changing the order of integration. The last line follows from the definition of Φ_t, (11). These expressions represent the Malliavin derivative of the agent's payoff taken in the direction from the strategy a towards the strategy $\hat{a}$.

If the strategy a does not satisfy (12) on a set of positive measure, then let us choose $\hat{a}_t > a_t$ when $\Phi_t > h'(a_t)$, $\hat{a}_t < a_t$ when $\Phi_t < h'(a_t)$ and $a_t > 0$, and otherwise let $\hat{a}_t = a_t = 0$. Then

$$E^a\left[\int_0^\tau e^{-rt}(\hat{a}_t - a_t)(\Phi_t - h'(a_t))\, dt\right] > 0,$$

and so a deviation to the strategy $(1-\epsilon)a + \epsilon\hat{a}$ for sufficiently small ϵ is profitable. ■

One may wonder why the proof of Proposition 1 proceeds by differentiating the agent's payoff in the direction as the agent deviates from one strategy a to another strategy $\hat{a}$. One natural guess to get condition (12) would be to try to look at instantaneous deviations with effort at particular moments of time, like one-shot deviations in discrete time. In continuous time, this method does not work, because individual time points have measure 0. Another approach involves looking at deviations to the strategy $\hat{a}$, as in Sannikov (2008). Unfortunately, that does not work either, because the agent's actions have delayed consequences. In the setting of Sannikov (2008), the agent's continuation value W_t and its sensitivity to output Δ_t depend on the agent's past actions only through the history of past output $\{X_s, s \in [0,t]\}$. In contrast, here past actions influence future output, and affect the agent's continuation value and incentives. Therefore, if condition (12) is violated, then there is no guarantee that the deviation strategy $\hat{a}$ that maximizes $\Phi_t a - h(a)$ is superior to a, as such a deviation induces a different process $\hat{\Phi}_t$. Hence, Proposition 1 proceeds by taking a Malliavin derivative in the direction of the deviation strategy.

Next, I will show that the representation (11) of Φ_t is equivalent to (9). First, the following standard proposition provides the law of motion of W_t and introduces formally the sensitivity Δ_t of W_t to X_t.

Proposition 2 *Fix a contract (c, τ) and a strategy a, with finite expected payoff to the agent. Then the processes W_t corresponds to the definition (8) if and only if for some Δ in L^2,*

$$dW_t = (rW_t - u_t)\,dt + \Delta_t \underbrace{(dX_t - \mu_t\,dt)}_{\sigma\,dZ_t} \tag{16}$$

and the transversality condition $E_t^a[e^{-rt}W_t] \to 0$ holds.[4]

Proof. See, for example, Proposition 1 in Sannikov (2013). ■

Using the options analogy, the drift of W_t in (16) follows from the requirement that the expected return on the option has to be r. The return equals to the sum of the dividend yield u_t/W_t and the capital gains rate $(rW_t - u_t)/W_t$.

Proposition 3 *The expressions (11) and (9) for Φ_t are equivalent.*

Proof. For simplicity, I will write dZ_t in place of $(dX_t - \mu_t\,dt)/\sigma$ below, as the agent's strategy a is fixed. Consider Φ_t defined by (11), and note that

$$\zeta_t^s = \int_t^s f(s'-t)\frac{dZ_{s'}}{\sigma} = \zeta_t^{t'} + \int_{t'}^s f(s'-t)\frac{dZ_{s'}}{\sigma}$$

[4] A process Δ is in L^2 if $E\left[\int_0^t \Delta_s^2\,ds\right] < \infty$ for all t.

when $t \le t' \le s$. Then

$$\Psi_{t'} \equiv E^a_{t'}\left[\int_t^\infty e^{-r(s-t)}\zeta_t^s u_s\, ds\right] = \int_t^{t'} e^{-r(s-t)}\zeta_t^s u_s\, ds + e^{-r(t'-t)}\zeta_t^{t'} W_{t'} + \Phi^t_{t'}$$

is a martingale, where $\Phi^t_{t'}$ is the contribution to Φ_t of the compensation after time t',

$$\Phi^t_{t'} \equiv E^a_{t'}\left[\int_{t'}^\infty e^{-r(s-t)}\left(\int_{t'}^s f(s'-t)\frac{dZ_{s'}}{\sigma} ds'\right) u_s\, ds\right].$$

Then, using Ito's lemma and (16), we have

$$d\Psi_{t'} = e^{-r(t'-t)} f(t'-t)\Delta_{t'}\, dt' + e^{-r(t'-t)}\left(\frac{f(t'-t)}{\sigma} W_{t'} + \zeta_t^{t'}\Delta_{t'}\sigma\right) dZ_{t'} + d\Phi^t_{t'}.$$

Integrating over $[t, t']$ and taking expectation at time t,

$$0 = E^a_t\left[\int_t^s e^{-r(t'-t)} f(t'-t)\Delta_{t'}\, dt'\right] + \underbrace{E^a_t\left[\int_t^s d\Phi^t_{t'}\right]}_{E^a_t[\Phi^t_s]-\Phi_t}.$$

Taking the limit $s \to \infty$, also assuming that the transversality condition $E^a_t\left[\Phi^t_s\right] \to 0$ as $s \to \infty$ holds, leads to (9). ■

Already using the expressions that have been derived so far, in particular the expression (11), it is possible to find the optimal contract fairly quickly under special assumptions. The following section derives the optimal contract for the particular class of empirically relevant environments, which I call the *large-firm* case. After that, I consider a more general environments and analyze the principal's problem using methods from stochastic control.

Remark. When dealing with the principal's problem, it is more convenient to think about the probability distribution over the paths of W_t and Φ_t instead of explicitly capturing how these processes depend on the path of output X controlled by the agent. Thus, from now on I express the processes W_t and Φ_t in terms of

$$dZ_s = \frac{dX_s - \mu_s\, ds}{\sigma} \tag{17}$$

rather than X directly. The principal removes expected output that is determined by the agent's effort, as he can infer effort from the agent's incentive constraints and can control it through contract, and takes into account only surprise innovations in output. From the point of view of the principal, on the equilibrium path, Z defined by (17) is a Brownian motion.

4 The Large-Firm Case.

This section presents a special *large-firm* case. This case leads to a tractable solution, which provides a convenient benchmark for the general case. It leads to an optimal contract in which compensation is determined by performance in closed form, and the optimal termination time is determined by a problem similar to the optimal exercise time of an American option.

Specifically, consider settings in which it is possible to expose the agent to only a small fraction of project risk due to noise, yet the benefits of giving the agent even small exposure to project risk can be significant. One situation in practice that matches these assumptions well is executive compensation. The informational problem is large, but a well-designed contract can have a strong impact to shareholder value. If a CEO of a \$10-billion dollar firm can add \$500 million a year (i.e. 5%) to shareholder value by increasing effort, this value is economically significant. However, if the volatility of firm value is 30%, then effort is extremely difficult to identify. Over t years, it matters how the incremental 5% return compares with the standard error of $30\%/\sqrt{t}$. Effort identification leads to a large probability of type I and II errors, and to motivate effort, the contract has to expose the agent to a significant amount of risk.[5] This makes incentive provision difficult. Jensen and Murphy (1990) and Murphy (1999) estimate that average CEO wealth increases by only \$3.25 to \$5 for each \$1000 increase in shareholder value.

Assume a quadratic cost of effort

$$h(a) = \theta a^2/2, \qquad a \in [0, \bar{a}],$$

where the bound $\bar{a}$ is large enough so that it never binds. Then the incentive constraint (12) reduces to $\theta a_t = \Phi_t$, i.e. the agent's effort is linear in Φ_t. The principal's profit and the agent's payoff can be expressed conveniently as follows.

Proposition 4

$$F_0 + \nu_0 W_0 = E\left[\int_0^\infty e^{-rt}(\hat{\nu}_t u_t - c_t)\, dt + e^{-r\tau} L\right], \tag{18}$$

[5] It is possible get a more precise signal about CEOs effort by measuring firm performance relative to an industry benchmark, but even then the estimation error remains significant.

where

$$\hat{\nu}_t = \hat{\nu}_0 + \int_0^t \hat{\lambda}_s \, dZ_s \quad \text{and} \quad \hat{\lambda}_t = \frac{1}{\theta\sigma} \int_0^{\min(t,\tau)} f(t-s)\, ds. \tag{19}$$

Proof. Substituting the incentive constraint

$$a_t = \frac{\Phi_t}{\theta} = \frac{1}{\theta\sigma} E_t \left[\int_t^\infty e^{-r(s-t)} \left(\int_t^s f(s'-t)\, dZ_{s'} \right) u_s \, ds \right] \tag{20}$$

(where we used the expression (11) for Φ_t) into the profit expression (3), we get

$$E^a \left[\int_0^\tau e^{-rt} \left(\int_t^\infty e^{-r(s-t)} \left(\int_t^s f(s'-t)\, dZ_{s'} \right) u_s \, ds \right) dt - \int_0^\infty e^{-rt} c_t \, dt + e^{-r\tau} L \right].$$

After changing the order of integration twice we obtain (18). ■

Expression (18) suggests that in general there are two channels, through which the compensation c_t after a particular history of performance at time t, affects the principal's profit. The first channel is direct: utility $u(c_t)$ affects the agent's incentives along the path $\{X_s,\ s \in [0,t]\}$; the value of incremental output created along this path is measured by $\hat{\nu}_t u(c_t)$. The second channel is indirect: c_t affects the agent's cost of effort $h(a_s)$ along the entire path $\{X_s,\ s \in [0,t]\}$. The second channel is complicated, as $h(a_s)$ affects u_s and feeds into the agent's incentives before time s, for all $s < t$. It is also not clear the extent to which the second channel matters empirically. The particular assumptions of the "large-firm" case eliminate the second channel, and simplify the problem.

To capture the "large-firm" case, consider the limit

$$\sigma \to \infty \quad \text{and} \quad \theta \to 0, \quad \text{with } \sigma\theta = \psi \in (0,\infty). \tag{21}$$

While the noise grows, so does the value that the agent can create for any given cost of effort. Then for any effort strategy

- the agent's actions have *marginal* impact on output, relative to noise
- the first-order conditions are sufficient for the *optimality* of the agent's strategy, and

- the cost of effort is negligible, i.e. we can replace u_t with $u(c_t)$ in (18).

Then to maximize (18), consumption c_t after each history $\{Z_s, s \in [0,t]\}$ must directly solve

$$c_t = \arg\max_c \hat{\nu}_t u(c) - c. \tag{22}$$

Thus, in the "large-firm" case, the values of $\hat{\nu}_t$ defined by (55) coincide with the values of the multiplier ν_t on the agent's utility.

For the exponential case, the law of motion of $\hat{\lambda}$, the volatility of the multiplier $\nu = \hat{\nu}$ on the agent's utility, reduces to

$$\hat{\lambda}_0 = 0, \quad d\hat{\lambda}_t = \frac{1_{t \le \tau}(r+\kappa)}{\theta\sigma}\, dt - \kappa\hat{\lambda}_t\, dt. \tag{23}$$

This expression is intuitive: it reflects the extent to which the agent could have affected output at time t with past effort. It only makes sense to make the agent's compensation sensitive to output at time t to the extent that the agent has control over output through effort choice. In particular, we see from (55) and (23) that the agent's exposure to risk increases over the agent's tenure but falls after termination (assuming that function f is decreasing) as outcomes start to depend less and less on the agent's past effort.

Note also that $\hat{\nu}_t$ is a martingale[6], i.e. any output surprise shifts the agent's compensation over the entire future. Thus, while the contract exposes the agent to risk, it smoothes consumption to minimize the negative impact on the agent's utility. Over time, some of the risks cancel each other out.

Importantly, the characterization of the optimal contract given by (22) and (55) is remarkably simple! The specification of the optimal termination time τ aside, these formulas describe exactly how output paths map into the agent's compensation. Payments to the agent are continuously increasing in ν_t and fall to zero whenever $\hat{\nu}_t$ becomes less than $1/u'(0)$. The volatility $\hat{\lambda}_t$ of $\hat{\nu}_t$ is deterministic both before and after time τ.

It is remarkable that, in a setting with delayed information revelation about the agent's actions, such a simple characterization of the optimal contract exists, even for arbitrary impact profiles $f(t)$.

Optimal Termination. Termination is an optimal stopping problem, and its solution is standard. The trade-off is that termination generates an instantaneous liquidation value L but it allows the volatility of $\hat{\nu}_t$ to

[6]This property is sometimes referred to as the Inverse Euler equation.

decay (thus harming the objective function (18)). Proposition 11 in the Appendix characterizes the continuation value function $\underline{G}(\nu_\tau, \tau)$ for (18) after the termination time τ. Then the value function $G(\nu, t)$ before termination must solve the parabolic equation[7]

$$rG(\nu, t) = \max_c \; \nu u(c) - c + \; G_2(\nu, t) + \frac{\hat{\lambda}_t^2}{2} \, G_{11}(\nu, t), \tag{24}$$

in the region of employment $\mathcal{R} \subset \mathbb{R} \times [0, \infty)$, and must satisfy the smooth-pasting conditions

$$G(\nu, t) = \underline{G}(\nu, t) + L \quad \text{and} \quad \nabla G(\nu, t) = \nabla \underline{G}(\nu, t)$$

on the boundary of $\mathcal{R}$.[8] Since ν_t is the multiplier on the agent's utility, it follows immediately that the agent's continuation payoff is given by

$$W_t = \begin{cases} G_\nu(\nu_t, t) \text{ for } t < \tau \\ \underline{G}_\nu(\nu_\tau, \tau) \text{ at time } \tau. \end{cases} \tag{26}$$

Figure 1 illustrates the optimal contract for the case when $u(c) = \sqrt{c}$, $r = 0.05$, $\kappa = 0.4$, $\sigma\theta = 1$, and $L = 100$. The volatility of ν_t, displayed on the vertical axis, converges exponentially to its target level of $(r + \kappa)/(\theta\sigma\kappa)$ during employment. Termination occurs only when ν_t becomes sufficiently negative, when the pair $(\nu_t, \hat{\lambda}_t)$ hits the boundary displayed on the figure. As L increases, termination occurs sooner.

---

[7]These value functions are defined as the following expectations, under the appropriate laws of motion of ν_t and $\hat{\lambda}_t$:

$$\underline{G}(\nu_\tau, \tau) = E_\tau \left[\int_\tau^\infty e^{-r(s-\tau)} \chi(\nu_s) \, ds \right] \quad \text{and}$$

$$G(\nu_t, t) = \max_\tau \; E_t \left[\int_t^\tau e^{-r(s-t)} \chi(\nu_s) \, ds + e^{-r(\tau - t)} (L + \underline{G}(\nu_\tau, \tau)) \right],$$

where $\chi(\nu) \equiv \max_c \; \nu u(c) - c$.

[8]In addition, to ensure that the stopping time τ is optimal, function $\underline{G}(\nu, t)$ must satisfy

$$r\underline{G}(\nu, t) \geq \max_c \; \nu u(c) - c + \; \underline{G}_2(\nu, t) + \frac{\hat{\lambda}_t^2}{2} \, \underline{G}_{11}(\nu, t), \tag{25}$$

outside $\mathcal{R}$, where

$$\hat{\lambda}_t = \int_0^t \frac{f(t-s)}{\psi} ds.$$

Condition (25) is needed for the verification argument, to ensure that continuing employment outside the region $\mathcal{R}$ is suboptimal.

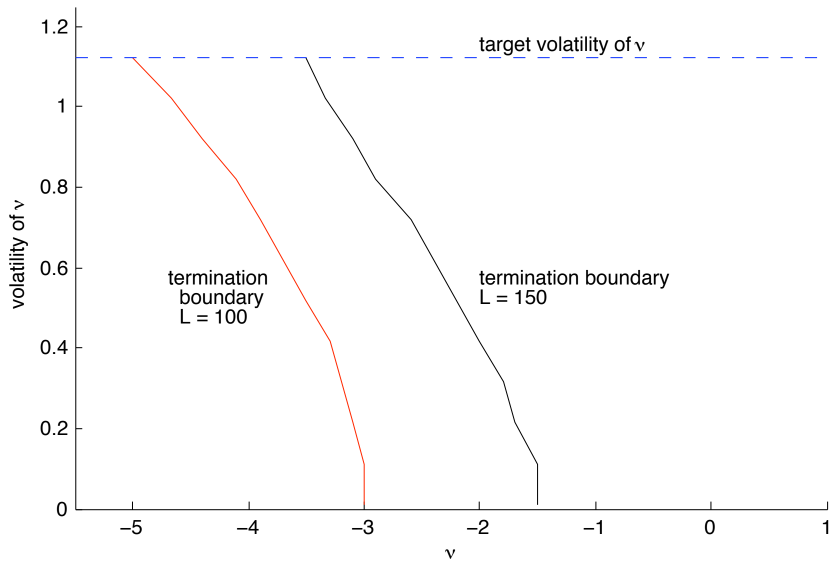

Figure 1: Contract dynamics in the large-firm case.

5 Contract Design via Optimal Control.

This section investigates the general version of principal's problem using methods from optimal stochastic control. To ensure that the principal's problem has a recursive structure, assume that the impact of the agent's actions on future outcomes is exponentially decaying, with the impact function

$$f(t) = (r + \kappa)e^{-\kappa t}.$$

Then the agent's information rent $\Phi_t = dW_t/dA_t$ has a recursive representation characterized by the following proposition.

Proposition 5 *Fix a contract* (c, τ) *and a strategy* a. *Then the finite process* Φ_t *is characterized by (9) (or, equivalently, (11)) if and only if for some* Γ *in* L^2,

$$d\Phi_t = (r + \kappa)(\Phi_t - \Delta_t)\, dt + \Gamma_t\, (dX_t - \mu_t\, dt) \tag{27}$$

and the transversality condition $E_t^a[e^{-(r+\kappa)s}\Phi_s] \to 0$ *as* $s \to \infty$ *holds, where* Δ_t *is the sensitivity of* W_t *to output from (16).*

Proof. See Appendix. ■

The *relaxed problem* of maximizing the principal's profit

$$E^a \left[\int_0^\tau e^{-rt} a(\Phi_t)\, dt + e^{-r\tau} L - \int_0^\infty e^{-rt} c_t\, dt \right] \tag{28}$$

subject to

$$W_0 = E^a \left[\int_0^\infty e^{-rt} u_t\, dt \right] \tag{29}$$

and just the first-order incentive constraints (12) can be framed as an optimal stochastic control problem. Consider the state variables W_t and Φ_t, controls c_t, Δ_t, Γ_t and τ, objective function (28) and the laws of motion of the state variables

$$dW_t = (rW_t - u(c_t) + 1_{t \le \tau} h(a(\Phi_t)))\, dt + \Delta_t\, \sigma\, dZ_t \tag{30}$$

$$\text{and} \quad d\Phi_t = (r + \kappa)(\Phi_t - \Delta_t)\, dt + \Gamma_t\, \sigma\, dZ_t,$$

subject to appropriate transversality conditions. Then by Propositions 2 and 5, this control problem is equivalent to the relaxed problem. The method of solving the relaxed problem, instead of the original problem, and verifying the full set of incentive constraints ex-post, is called the *first-order approach.* See, for example, Kapicka (2012).

The solution to the relaxed problem may or may not satisfy the full set of incentive constrains (5) (i.e. be fully incentive compatible). If it does, then the contract that solves the relaxed problem is in fact the optimal contract. If it does not, then it merely provides a lower bound on the principal's objective function.

There are several methods for checking full incentive compatibility. The most direct way is to compute numerically the agent's optimal strategy under a given contract. That is, one has to solve the agent's optimal control problem, which, in this case, would have three state variables: the recursive variables W_t and Φ_t of the candidate contract and the stock of past effort

$$A_t \equiv \int_0^t e^{-\kappa(t-s)} a_s\, ds,$$

which summarizes the agent's past deviations. This is a laborious approach, which has been implemented in another context, for example, in Werning (2002).

A much quicker test to see if a given contract is fully incentive compatible is to use a simple sufficient condition. The following proposition derives one such condition.

Proposition 6 *Suppose that the agent's cost of effort is quadratic of the form $h(a) = \theta a^2/2$. Then an effort strategy a satisfies (5) if it satisfies (12) and also*[9]

$$\Gamma_t \leq \frac{\theta\kappa}{2}. \tag{31}$$

Proof. See Appendix. ■

The sufficient condition (31) is a bound on the rate Γ_t at which incentives Φ_t change with output X_t. If Γ_t is large, then following a reduction in effort, the agent can benefit by lowering effort further as he faces a lower Φ_t. Note the analogy with options. If the agent's contract is a package of call options on X_t, then the agent's incentives to lower effort depend the downside protection of calls. The more protection the agent gets, the quicker the *Delta*s of the agent's options have to fall as losses occur, i.e. the *Gamma*s of the agent's options are higher.

Condition (31) gives us another benefit: we can use it to derive a good fully incentive-compatible contract recursively (using a control problem) in situations where the first-order approach fails. Specifically, consider the relaxed control problem described above together with the restriction that Γ_t must satisfy (31), and call it the *restricted control problem*. This recursive problem, with the same state variables W_t and Φ_t, leads to a fully incentive-compatible contract that gives a lower bound on the optimal contract profit. This is a fruitful approach, which lets one avoid a dead end in the event that the first-order approach fails. Moreover, we should expect that in many cases, the restriction (31) has minimal impact on efficiency, particularly when condition (31) is violated only in some distant parts of the state space under the solution to the relaxed control problem.

Various methods exist to tackle the relaxed control problem (28)-(30). One can approach it directly using the HJB equation

$$rF(W,\Phi) = \max_{c,\Delta,\Gamma} a(\Phi) - c + [F_W \; F_\Phi] \begin{bmatrix} rW - u(c) + h(a(\Phi)) \\ (r+\kappa)(\Phi - \Delta) \end{bmatrix} +$$

[9] Condition (31) can be weakened to $\Gamma_t \leq \theta\kappa$ for $t \geq \tau$.

$$\frac{\sigma^2}{2}\frac{\partial^2}{\partial\epsilon^2}F(W+\Delta\epsilon,\Phi+\Gamma\epsilon). \tag{32}$$

The solution to the restricted control problem is characterized by the same equation, but with the constraint (31) imposed on the maximization problem.

Equation (32) is hardly tractable. It is a second-order partial differential equation with a degenerate (parabolic) second-order derivative in an endogenous direction. As shown below, the solution to the control problem derived using the stochastic maximum principle, in the domain of Lagrange multipliers on the variables W_t and Φ_t, is significantly more tractable. The solution is analogous that given by (55) and (22) in the large-firm case, but taking into account the agent's nonnegligible disutility of effort. In fact, the large-firm case, as well as the "standard" model with $\kappa=\infty$, provide good benchmark for interpreting the solution.

The Lagrange multipliers ν_t and λ_t are the principal's marginal costs of the state variables W_t and Φ_t (to draw connection the the HJB equation, $\nu_t=-F_W(W_t,\Phi_t)$ and $\lambda_t=-F_\Phi(W_t,\Phi_t)$). Necessary first-order conditions for the solution of a control problem via the multipliers can be obtained by the following mechanical procedure. One writes the Hamiltonian: payoff flow minus multipliers multiplied by the drift of corresponding state variables, minus the products of the volatilities of the multipliers and the volatilities of the state variables. In our case, the Hamiltonian is

$$H=1_{t\le\tau}a(\Phi)-c-[\nu\ \lambda]\left[\begin{array}{c} rW-u(c)+1_{t\le\tau}h(a(\Phi)) \\ (r+\kappa)(\Phi-\Delta)\end{array}\right]-[\sigma^\nu\ \sigma^\lambda]\left[\begin{array}{c}\Delta\sigma\\ \Gamma\sigma\end{array}\right].$$

The first-order conditions, obtained by differentiating the Hamiltonian with respect to the *controls*, have to hold (corner solutions are also possible).[10] Differentiating with respect to c, we obtain the condition (22). Differentiating with respect to Δ, we find that the volatility of ν_t is $\sigma^\nu=(r+\kappa)/\sigma$. Differentiating with respect to Γ, we find that the volatility of λ_t is 0.

The drifts of the multipliers are obtained by differentiating the Hamiltonian with respect to the *states*. The drift of ν is $H_W+r\nu$, and the drift of

[10]Strangely, optimal controls may not always optimize the Hamiltonian - they may correspond to local minima rather than maxima. Yong and Zhou (1999) introduce an additional process to deal with this issue. However, the benefits of the extra process are not clear, since with or without it, the Hamiltonian gives only necessary first-order conditions for optimal control, and a verification argument is necessary to check that the controls are indeed optimal.

λ is $H_\Phi + r\lambda$, where r stands for discounting. Performing differentiation, we find the laws of motion of multipliers to be

$$d\nu_t = \lambda_t \frac{r+\kappa}{\sigma}\, dZ_t \qquad \text{and} \tag{33}$$

$$d\lambda_t = 1_{t\leq\tau}\, a'(\Phi_t)(1 - h'(a(\Phi_t))\nu_t)\, dt - \kappa\lambda_t\, dt, \quad \lambda_0 = 0.$$

The advantages of the Lagrangian characterization, over that expressed directly in terms of the states W_t and Φ_t, are as follows:

- The agent's compensation is determined directly from ν_t rather than indirectly through a function on the space of W and Φ : c_t maximizes

$$\nu_t u(c) - c.$$

- The joint law of motion of the multipliers ν_t and λ_t is simpler and more explicit than the laws of motion (30) of W_t and Φ_t. In particular, λ_t is a *slow-moving* variable, as it has no volatility. Variable ν_t has no drift,[11] and its volatility is determined explicitly by λ_t.

- Only one ingredient of the joint law of motion of ν_t and λ_t is not explicitly determined: it is the values of Φ on the space of ν and λ. This variable determines the agent's effort, the drift of λ_t and expected output μ_t (needed to calculate the Brownian motion dZ_t from dX_t). I outline the procedure I use to determine the function Φ_t and compute the principal's value function in Appendix B.

The Lagrangian characterization, however, does not explain how the optimal termination time τ can be determined. I explain the determination of τ, as well as the determination of Φ that enters the law of motion of λ, in the next subsection. The next subsection also provides sufficient conditions for the characterization (33) to lead to a solution to the principal's control problem (see Proposition 8). A reader who wishes to see the characterization (33) in practice may wish to skip to Section 6, which illustrates the characterization through its relationship to the large-firm case.

[11] This is the well-known inverse Euler equation, e.g. see Spear and Srivastava (1987).

5.1 The Determination of Φ_t and τ.

This section explains how the process Φ_t and the termination time τ that enter the characterization (33) are determined. I characterize the variables Φ, W and F (the principal's profit given by (32)), as well as the boundary where termination occurs, by a system of partial differential equations over the space of Lagrange multipliers ν and λ. I do so in two steps. First, I derive the form of the optimal contract after time τ. This step describes the range of options available to the principal at termination time: the range of implementable pairs (W, Φ) and cheapest ways of implementing them. Second, I characterize the functions W, Φ and F before termination.

The optimal contract after termination. Some of agent's compensation may be paid out after termination. The form of this compensation influences the agent's incentives during employment. A contract that solves the relaxed problem (28) has to give the agent the desired continuation value W_τ and information rent Φ_τ at time τ in the cheapest possible way. Indeed, if we replace the continuation contract after time τ with another contract with the same values of W_τ and Φ_τ, the agent's marginal incentives during employment remain unchanged.

Formally, the optimal contract after termination has to solve the following problem (where τ is replaced by 0 to simplify notation):[12]

$$\max_c \quad E\left[-\int_0^\infty e^{-rt} c_t \, dt\right] \tag{34}$$

$$\text{s.t.} \quad E\left[\int_0^\infty e^{-rt} u(c_t)\, dt\right] = W_0 \quad \text{and} \quad E\left[\int_0^\infty e^{-rt} \zeta_0^t \, u(c_t)\, dt\right] = \Phi_0.$$

Problem (34) is easy to solve. Letting ν_0 and λ_0 be the multipliers on the

[12] Interestingly, problem (34) also solves a different interesting model, in which the agent puts effort only once at time 0, and his effort determines the unobservable level of fundamentals μ_0. Specifically, suppose the agent's utility is given by

$$F_0 = E\left[\int_0^\infty e^{-rt} u(c_t)\, dt\right] - H(\mu_0),$$

where H is a convex increasing cost of effort, and fundamentals affect output according to $dX_t = \mu_t \, dt + \sigma dZ_t$, where $\mu_t = e^{-\kappa t}\mu_0$. Then the agent's incentive constraint is $H'(\mu_0) = \Phi_0/(r+\kappa)$. A version of this problem has been solved on Hopenhayn and Jarque (2010).

two constraints, the Lagrangian is

$$E\left[\int_0^\infty e^{-rt}\left((\nu_0+\zeta_0^t\lambda_0)u(c_t)-c_t\right)\,dt\right]-\nu_0 W_0-\lambda_0\Phi_0.$$

The first-order condition is

$$c_t=\arg\max_c\ \underbrace{(\nu_0+\zeta_0^t\lambda_0)}_{\nu_t}u(c)-c, \tag{35}$$

where ν_t is the multiplier on the agent's utility at time t. From (10), the laws of motion of the Lagrange multipliers can be expressed as

$$d\nu_t=\lambda_0\,d\zeta_0^t=\underbrace{e^{-\kappa t}\lambda_0}_{\lambda_t}\frac{r+\kappa}{\sigma}\frac{dX_t-\mu_t\,dt}{\sigma},\quad \text{and}\quad d\lambda_t=-\kappa\lambda_t\,dt. \tag{36}$$

This corresponds to the solution (33) after time τ.

Proposition 7 characterizes the functions $\underline{W}$, $\underline{\Phi}$ and $\underline{F}:\mathbb{R}\times[0,\infty)\to\mathbb{R}$, which describe the correspondence between the multipliers (ν_0,λ_0) and variables (W_0,Φ_0) and F_0 in problem (34).

Proposition 7 *Functions $\underline{W}$, $\underline{\Phi}$ and $\underline{F}$ solve the following system of parabolic equations*

$$r\underline{W}=u(c)-\kappa\lambda\,\underline{W}_\lambda+\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{\underline{W}_{\nu\nu}}{2}, \tag{37}$$

$$(r+\kappa)\underline{\Phi}=(r+\kappa)\lambda\frac{r+\kappa}{\sigma^2}\underline{W}_\nu-\kappa\lambda\,\underline{\Phi}_\lambda+\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{\underline{\Phi}_{\nu\nu}}{2} \tag{38}$$

$$\text{and}\quad r\underline{F}=-c-\kappa\lambda\,\underline{F}_\lambda+\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{\underline{F}_{\nu\nu}}{2}, \tag{39}$$

where c for any pair (ν,λ) maximizes $u(c)\nu-c$. All three solutions can also be derived from a single convex function $\underline{G}$ that solves

$$r\underline{G}(\nu,\lambda)=\max_c\ \nu u(c)-c-\kappa\lambda\,\underline{G}_\lambda(\nu,\lambda)+\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{\underline{G}_{\nu\nu}(\nu,\lambda)}{2}. \tag{40}$$

Then $\underline{W}=\underline{G}_\nu$, $\underline{\Phi}=\underline{G}_\lambda$ and $\underline{F}=\underline{G}-\nu\underline{G}_\nu-\lambda\underline{G}_\lambda$. Function $\underline{G}$ has the stochastic representation

$$\underline{G}(\nu_0,\lambda_0)=\max_{\{c_t\}}\ E\left[\int_0^\infty e^{-rt}\left(\nu_t u(c_t)-c_t\right)\,dt\right], \tag{41}$$

where (ν_t,λ_t) follow (36).

Proof. Equation (41) is a standard stochastic representation of the solution of the parabolic partial differential equation (40) (see Karatzas and Shreve (1991)). Since $\nu_t = \nu_0 + \zeta_0^t \lambda_0$, differentiating (41) with respect to ν_0 and using the Envelope theorem, we get

$$\underline{G}_\nu(\nu_0, \lambda_0) = E\left[\int_0^\infty e^{-rt} u(c_t)\, dt\right] = W_0. \tag{42}$$

Differentiating with respect to λ_0 we get

$$\underline{G}_\lambda(\nu_0, \lambda_0) = E\left[\int_0^\infty e^{-rt} \zeta_0^t\, u(c_t)\, dt\right] = \Phi_0. \tag{43}$$

Finally, from (41) directly, $\underline{G} = \underline{F} + \nu \underline{W} + \lambda \underline{\Phi}$. The equations for $\underline{W}$, $\underline{\Phi}$ and $\underline{F}$ can be obtained by differentiating (40) or by matching their stochastic representations with the corresponding equations. ■

The set of options available to the principal after termination can be found either by solving the system (37) through (39) or a single equation (40).[13] One needs to know these functions to determine the optimal termination time τ.

The relationship between $\underline{G}(\nu, \lambda)$ and the value function $\underline{F}(W, \Phi)$ over the space of pairs (W, Φ) attainable to the principal after termination is given by

$$\underline{G}(\nu, \lambda) = \max_{W, \Phi}\ \underline{F}(W, \Phi) + \nu W + \lambda \Phi.$$

Since the function $\underline{G}$ is strictly convex and C^2, it follows that the principal's value function $\underline{F}(W, \Phi)$ is concave and C^2.

The optimal contract before termination. The optimal termination time is a stopping time, at which the multipliers (ν_t, λ_t) reach the boundaries of the employment region $\mathcal{R} \subseteq [0, \infty) \times \mathbb{R}$. On $\mathcal{R}$, the maps from (ν_t, λ_t) to W_t, Φ_t and F_t are characterized by a system of equations

$$rW = u(c) + (a'(\Phi)(1 - \Phi\nu) - \kappa\lambda)\, W_\lambda + \lambda^2 \frac{(r+\kappa)^2}{\sigma^2} \frac{W_{\nu\nu}}{2}, \tag{44}$$

[13] Boundary conditions are not necessary: there is a single non-explosive solution to each of these equations because the process λ never reaches 0. However, for the purposes of numerical integration, it makes sense to impose $\underline{W}(\nu, \epsilon) = u(c)/r$, $\underline{\Phi}(\nu, \epsilon) = 0$ and $\underline{F}(\nu, \epsilon) = -c/r$, for ϵ close to 0, with c that maximizes $u(c)\nu - c$. Equation (40) or (39) through (39) can then be solved in the direction of increasing λ.

$$(r+\kappa)\Phi = (r+\kappa)\lambda\frac{r+\kappa}{\sigma^2}W_\nu + (a'(\Phi)(1-\Phi\nu)-\kappa\lambda)\,\Phi_\lambda + \lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{\Phi_{\nu\nu}}{2}, \quad (45)$$

$$\text{and} \quad rF = a(\Phi) - c + (a'(\Phi)(1-\Phi\nu)-\kappa\lambda)\,F_\lambda + \lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{F_{\nu\nu}}{2}, \quad (46)$$

where c for any pair (ν,λ) maximizes $u(c)\nu - c$. Alternatively, all three maps can be characterized by a single function $G : \mathcal{R} \to \mathbb{R}$ that solves

$$rG = a(G_\lambda) - c + \nu\,(u(c) - h(a(G_\lambda))) - \kappa\lambda G_\lambda + \lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{G_{\nu\nu}}{2}, \quad (47)$$

so that

$$W = G_\nu, \quad \Phi = G_\lambda \quad \text{and} \quad F = G - \nu G_\nu - \lambda G_\lambda. \quad (48)$$

The relevant smooth-pasting conditions on the boundary of $\mathcal{R}$ are

$$G(\nu,\lambda) = \underline{G}(\nu,\lambda) + L \quad \text{and} \quad \nabla G(\nu,\lambda) = \nabla\underline{G}(\nu,\lambda). \quad (49)$$

Proposition 8 *Suppose that function G solves equation (47) on $\mathcal{R} \subseteq [0,\infty)\times\mathbb{R}$ and satisfies the smooth-pasting conditions (49) on the boundary. Then, as long as the transversality conditions hold, W_t, Φ_t and F_t are given by (48) and functions $W(\nu,\lambda)$, $\Phi(\nu,\lambda)$ and $F(\nu,\lambda)$ solve equations (44) through (46) in the contract defined by (33).*

Sufficient conditions for the contract to solve the relaxed control problem (28) are as follows: on $\mathcal{R}$, $G(\nu,\lambda) \geq \underline{G}(\nu,\lambda)$ and the Hessian of G is positive definite, and outside $\mathcal{R}$,

$$r\underline{G} \geq \max_c \; a(\underline{G}_\lambda) - c + \nu\,(u(c) - h(a(\underline{G}_\lambda))) - \kappa\lambda\underline{G}_\lambda + \lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{\underline{G}_{\nu\nu}}{2}. \quad (50)$$

Proof. See Appendix. ■

The last paragraph of Proposition 8 gives sufficient conditions for the martingale verification argument to go through. They can be verified numerically or, for models near the special large-firm case, analytically (see Section 6). Strictly speaking, the stochastic maximum principle provides only necessary first-order conditions for the solution of a stochastic control problem, and a verification argument is required to ensure the optimality of the solution.[14]

In the proof of Proposition 8, I adapt the standard martingale verification argument to demonstrate the optimality of a candidate control policy described by (33). It is possible to carry out the martingale verification argument directly using the properties of the value function $F(W, \Phi)$ that corresponds to the HJB equation (32), which is implied by the solution $G(\nu, \lambda)$. It is also possible to construct an indirect verification argument using the function G itself. In the proof of Proposition 8, I demonstrate how to do the latter, but the argument can be translated into an argument about F. Under the conditions of Proposition 8, the relationship between functions F and G is given by

$$G(\nu, \lambda) = \max_{W,\Phi} \; F(W, \Phi) + \nu W + \lambda \Phi.$$

5.2 Numerical Examples.

This section provides several numerical examples. Before presenting them, I would like to summarize several key observations about the effects of κ, the key parameter that describes how the impact of the agent's effort is distributed is κ. First, the principal's profit, as a function of W_0, generally has very little sensitivity to κ.[15] This is in virtue of model specification, which assumes that κ affects only the horizon over which the agent's effort has impact, but not the present value created by effort. It ultimately matters how much value the agent's effort creates, not how this value is distributed over time. Second, while contract design - specifically, the rate at which the volatility of ν_t converges to its target level given by (52) - does depend on κ, the target volatility of ν_t itself has very little sensitivity to κ. Third, because of these observations, for any κ it is possible to design an approximately optimal contract by borrowing the target level of λ_t from the standard case of $\kappa = \infty$, and adjust for κ by letting λ_t converge to its target gradually (e.g. at rate κ) rather than instantaneously.

[14]This issue is separate from the sufficiency of the first-order conditions for the agent: here I am discussing the sufficiency of the first-order conditions for the principal.

[15]In contrast, the principal's profit is hugely sensitive to parameters θ and σ, especially in the large firm case. For example, if it is possible to reduce the volatility of the signal about the agent's performance by half, e.g. measuring firm's stock performance against an appropriate industry benchmark, the objective function for the principal's problem increases by a factor of about 4.

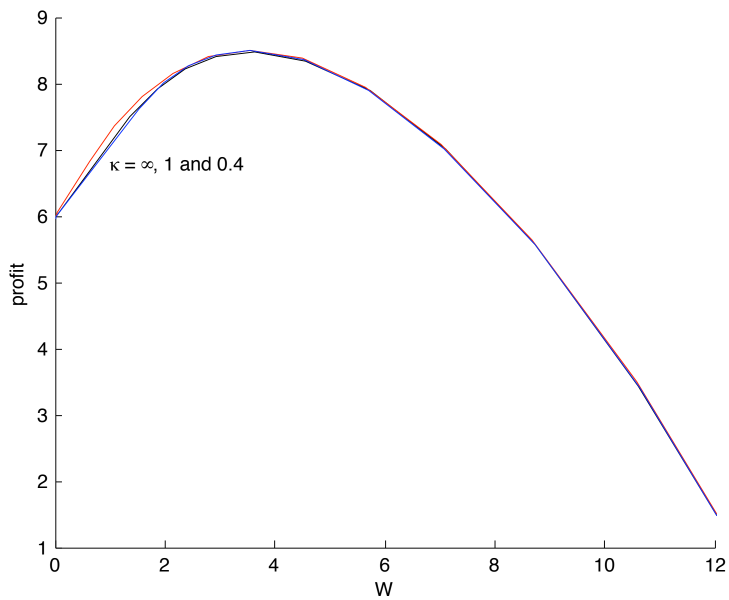

Figure 2: Principal's profit, as a function of W_0.

Let us take $u(c) = \sqrt{c}$, $h(a) = \theta a^2/2$ with $\theta = 0.5$, $r = 0.05$, $\sigma = 4$, and $L = 5$, and solve the relaxed control problem. Then Figure 2 compares the principal's profit in this example, for parameters $\kappa = 0.4$, 1 and ∞. The distinction between the three profit functions is minimal, and the difference in maximal profit across these examples is about 0.03.[16]

The sufficient incentive condition (31) is satisfied in these examples, and it holds more easily when κ is larger. Figure 3 shows the level of Γ at the target level of λ for $\nu = 0$. Γ tends to be the largest at the left termination boundary, and for large values of λ.

Figure 4 illustrates the dynamics of the state variables ν_t and λ_t, for $\kappa = 0.4$. For clarity, the vertical axis displays the volatility of ν_t, $\lambda_t(r+\kappa)/\sigma$, rather than λ_t itself. The target volatility of ν_t is given by (**??**); and at $\nu_t = 0$ it is $(\kappa + r)/(\kappa\theta)$. To illustrate the dynamics, the solid curves on

[16]While the reader may guess that higher κ leads to greater profit (since information about effort is revealed sooner), this is not always the case as there are forces that pull profit in the opposite direction. For example, this model the signal-to-noise ratio improves slightly as κ declines. The reason is that while the present value of output is invariant with κ, the level of output relative to volatility increases slightly as κ declines.

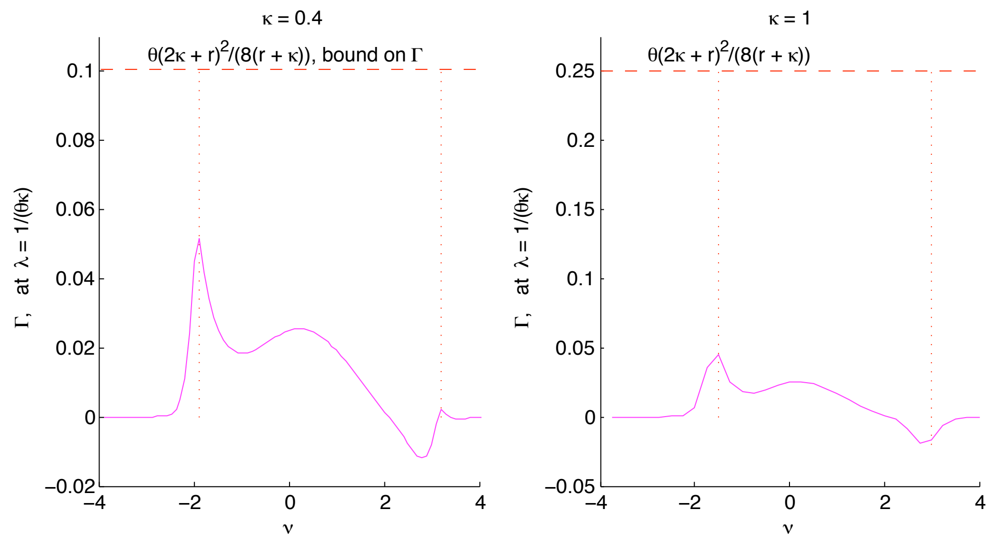

Figure 3: Γ is within bound (termination boundaries are dotted lines).

Figure 4 show the rate of change of the volatility of ν_t over one quarter, for several values of λ_t. Points where these curves intersect horizontal solid lines correspond to the target volatility of ν_t. The boundaries where termination occurs are indicated by dashed lines. Note that the volatility of ν_t under the optimal contract never goes above the level of 0.8, as the drift of the volatility (and the drift of λ_t) are uniformly negative at that level. Thus, the optimal contract does not use the full state space, but only a portion of it.

For comparison, Figure 4 also indicates the volatility of ν_t (multiplied by $(\kappa+r)/\kappa$ to account for the fact that the signal-to-noise ratio varies slightly with κ) for the standard benchmark of $\kappa = \infty$. The volatility of ν_t in the standard case matches remarkably well the target volatility for $\kappa = 0.4$. This close fit indicates that the benchmark with $\kappa = \infty$ is hugely informative about the structure of the optimal contract for other values of κ.[17]

The numerical relationship between the target volatility of ν_t for an arbitrary κ and for $\kappa = \infty$ suggests that the following procedure leads to an approximately optimal contract for arbitrary κ. First, solve for the optimal

[17]Note, however, that κ is significantly higher than r in all our examples. That is, information is revealed before the principal's ability to reward and punish the agent is eroded by discounting. This is a natural assumption for most applications in practice.

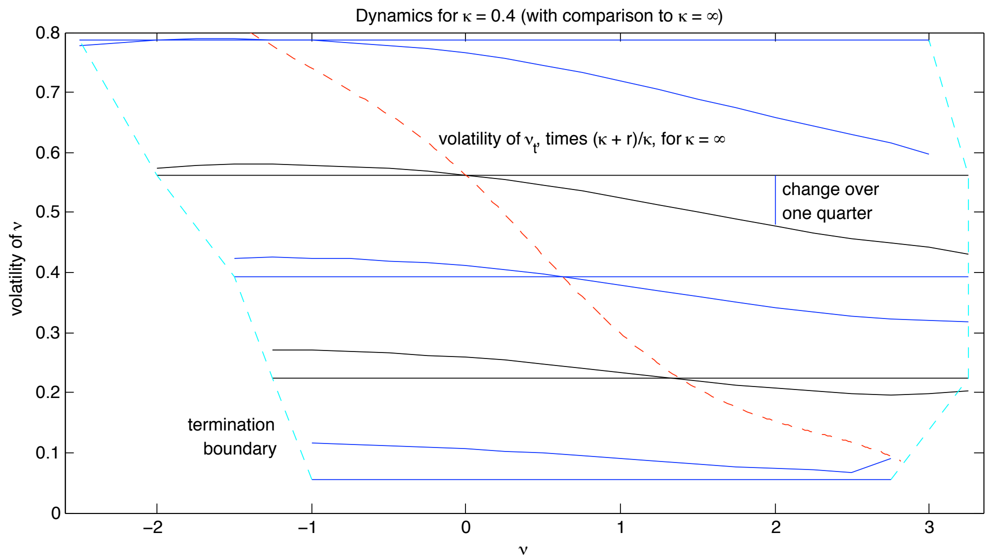

Figure 4: Dynamics of the volatility of ν_t, $\kappa = 0.4$.

contract for $\kappa = \infty$, and determine the agent's effort $\hat{a}(\nu)$. Then, let the state variables λ_t and ν_t for arbitrary κ evolve according to

$$d\nu_t = \lambda_t \frac{r+\kappa}{\sigma}, \quad d\lambda_t = 1_{t\leq\tau}\, a'(\Phi_t)|_{a(\Phi_t)=\hat{a}(\nu)}\,(1-\nu_t h'(\hat{a}(\nu)))\,dt - \kappa\lambda_t\,dt.$$

Determine the termination time τ optimally.

Figure 5 compares profit under the approximately optimal contract implied by this procedure to profit under the optimal contract. The approximately optimal contract does quite well - the distance between the two curves is only 0.1. For comparison, Figure 5 also presents profit from the contract that would be optimal in the large-firm case, in which λ_t follows $d\lambda_t = (1/\theta - \kappa\lambda_t)\,dt$ before termination (see Section 4.1). The contract designed for the large-firm case performs quite badly in this case.

6 Optimal Contracts near the Large-Firm Case.

This section adds transparency to the Lagrangian characterization of the optimal contract (33) by studying the form that it takes near the large-firm case. As in Section 4, assume that the cost of effort is quadratic of the form

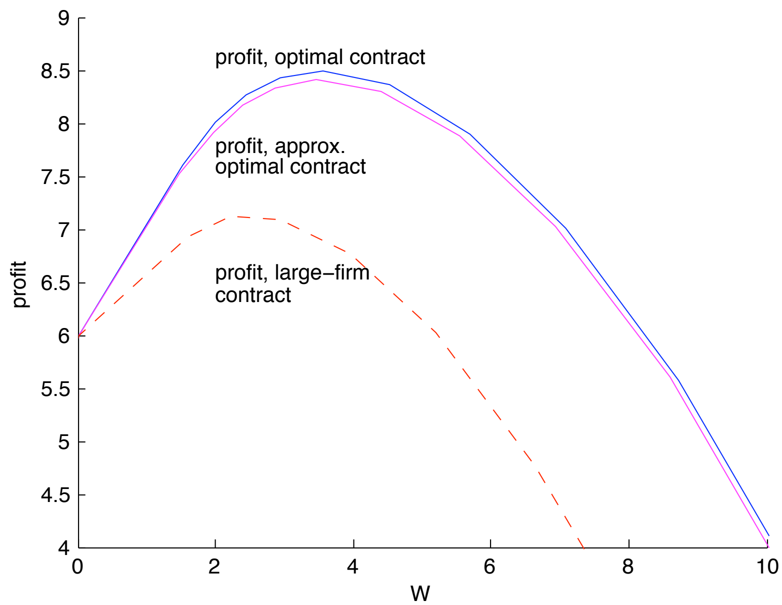

Figure 5: Approximating the optimal contract.

$h(a) = \theta a^2/2$. As θ rises away from 0 while keeping $\psi = \theta\sigma$ fixed, for small θ,

1. I describe how the optimal contract changes relative to the large-firm case

2. argue that the necessary first-order conditions on the Lagrange multipleirs lead to a full solution of the principal's control problem, i.e. that the function G satisfies the conditions of Proposition 8

3. argue that the solution to the principal's control problem leads to a fully incentive-compatible contract, i.e. the necessary first-order incentive conditions for the *agent* are also sufficient and

4. demonstrate that the optimal contract for the large-firm case is approximately optimal for small θ, i.e. the loss from efficiency is on the order of θ^2.

To perform the comparison, it is useful to replace the state variable Φ with $\hat{\Phi} = \Phi\sigma/(r + \kappa)$. Then the principal's control problem is characterized by equations

$$\begin{bmatrix} dW_t \\ d\hat{\Phi}_t \end{bmatrix} = \begin{bmatrix} rW_t - u(c_t) + \theta a_t^2/2 \\ (r+\kappa)\hat{\Phi}_t - \hat{\Delta}_t \end{bmatrix} dt + \begin{bmatrix} \hat{\Delta}_t \\ \hat{\Gamma}_t \end{bmatrix} dZ_t,$$

where $\hat{\Delta}_t = \Delta_t \sigma$ and $\hat{\Gamma}_t = \Gamma_t \sigma^2/(r+\kappa)$. The multiplier on $\hat{\Phi}$ is given by $\lambda_t(r+\kappa)/\sigma$, and the law of motion of the multipliers on W and $\hat{\Phi}_t$ is captured by

$$d\nu_t = \hat{\lambda}_t \, dZ_t \quad \text{and} \quad \frac{d\hat{\lambda}_t}{dt} = 1_{t \le \tau} \frac{r+\kappa}{\psi} (1 - \nu_t \theta a_t) - \kappa \hat{\lambda}_t. \tag{51}$$

In this characterization, $\hat{\lambda}_t$ directly plays the role of the volatility of ν_t. Also, note that at $\nu_t = 0$ the drift of $\hat{\lambda}_t$ does not depend on θ.

First, we can obtain the form that the optimal contract takes near $\theta = 0$ by differentiating (51) with respect to θ. We find that

$$\frac{d\hat{\lambda}_t}{dt} = 1_{t \le \tau} \frac{r+\kappa}{\psi} \Big(1 - \nu_t \theta a^0(\nu_t, \hat{\lambda}_t) + o(\theta)\Big) - \kappa \hat{\lambda}_t,$$

where $a^0(\nu, \hat{\lambda})$ is the agent's effort in the optimal contract of the large-firm case. During employment, the volatility $\hat{\lambda}_t$ of ν_t tends to its target level of $\hat{\lambda}$, at which the drift of $\hat{\lambda}$ equals zero, satisfies the equation

$$\bar{\lambda}(\nu) = \frac{r+\kappa}{\kappa\psi} \left(1 - \nu \theta a^0(\nu, \bar{\lambda}(\nu)) + o(\theta)\right). \tag{52}$$

The target level of $\hat{\lambda}$ is $(r+\kappa)/(\kappa\psi)$ when $\nu = 0$, and it is higher when $\nu < 0$ and lower when $\nu > 0$. Intuitively, when $\nu < 0$, the agent's punishment involves greater risk exposure, which also forces the agent to expand a greater effort cost. When $\nu > 0$, the agent is rewarded by lower risk exposure, which allows him to work less.

Figure ... compares the target level of $\hat{\lambda}$ given by the approximate formula with the actual target for several values of θ near 0.

Next, it is useful to evaluate how the principal's profit changes with θ as $\psi = \sigma\theta$ is kept fixed. The following proposition provides such an estimate (which is valid for all $\theta > 0$).

Proposition 9 *The derivative of the principal's profit $F(W, \hat{\Phi})$ with respect to θ is given by*

$$F^\theta(W, \hat{\Phi}) \equiv -E\left[\int_0^\tau e^{-rt} \nu_t \frac{(r+\kappa)^2 \hat{\Phi}(\nu_t, \hat{\lambda}_t)^2}{2\psi^2} \, dt \mid \nu_0, \hat{\lambda}_0\right], \tag{53}$$

where $(\nu_0, \hat{\lambda}_0)$ *are chosen so that* $W = W(\nu_0, \hat{\lambda}_0)$ *and* $\hat{\Phi} = \hat{\Phi}(\nu_0, \hat{\lambda}_0)$. *The The derivative of the function* $G(\nu, \hat{\lambda})$ *with respect to* θ *at* $\theta = 0$ *is given by* $F^\theta(W(\nu, \hat{\lambda}), \hat{\Phi}(\nu, \hat{\lambda}))$.

Proof. See Appendix. ■

We can use Proposition 9 to check the sufficiency of first-order conditions for the principal and the agent near $\theta = 0$. For the principal, by Proposition 12 in the Appendix, function G is strictly convex for the large-firm case. The stochastic expression (53) implies that G changes in a differentiable manner, and must therefore remain convex for small θ.[18]

For the agent, note that the sufficient condition of Proposition 6 takes the form

$$\hat{\Gamma}_t \leq \frac{\sigma\psi\kappa}{2(r+\kappa)}. \tag{54}$$

This condition is trivially satisfied when $\sigma = \infty$. Moreover, since

$$\hat{\Gamma} = d\hat{\Phi}/dZ = \hat{\Phi}_\nu(\nu, \hat{\lambda})\hat{\lambda} = G_{\nu\hat{\lambda}}(\nu, \hat{\lambda})\hat{\lambda},$$

it follows that $d/d\theta\, \hat{\Gamma} = G_{\nu\hat{\lambda}}(\nu, \hat{\lambda})\hat{\lambda}$. Thus, from Proposition 9, it follows that $\hat{\Gamma}$, as a function of $(\nu, \hat{\lambda})$, changes continuously in θ, and so the sufficient condition (54) must hold for all θ close enough to 0.

Finally, let us argue that the optimal contract for the large-firm case remains approximately optimal in general when θ is close to 0. In this contract, the joint law of motion of $(\nu_t, \hat{\lambda}_t)$ is given by (55), i.e.

$$d\nu_t = \hat{\lambda}_t \frac{dX_t - a_t\, dt}{\sigma} \quad \text{and} \quad \frac{d\hat{\lambda}_t}{dt} = 1_{t\leq\tau}\frac{r+\kappa}{\psi} - \kappa\hat{\lambda}_t, \tag{55}$$

where a_t is the agent's optimal effort. This contract gives a *lower* bound on the principal's profit. It optimizes with respect to the agent's compensation without taking into account how the agent's cost of effort affects his utility, and therefore incentives. This approximation is valid when the cost of effort is, in fact, insignificant relative to the utility of consumption.[19]

[18] The remaining conditions of Proposition 8 follow from the optimal choice of the termination time τ.

[19] The agent's continuation value $W(\nu, \hat{\lambda})$, his incentives $\Phi(\nu, \hat{\lambda})$ and the principal's value function $F(\nu, \hat{\lambda})$ can be computed by solving a system of parabolic equations, analogous to (24).

We would like to argue that the derivative of the principal's value function with respect to θ under this contract is still given by $F^{\theta}(W, \hat{\Phi})$, as in Proposition (9). That is, loss of efficiency of contract (55) relative to the optimal contract is of $o(\theta)$. The following proposition evaluates how the maps from the pair $(\nu, \hat{\lambda})$ to W and $\hat{\Phi}$ change with θ.

Proposition 10 *Under the contract (55),*

$$\frac{d}{d\theta} W(\nu_0, \hat{\lambda}_0) = E\left[r \int_0^{\tau} e^{-rt} \frac{a_t^2}{2}\, dt \right] \quad \text{and}$$

$$\frac{d}{d\theta} \hat{\Psi}(\nu_0, \hat{\lambda}_0) = E\left[\int_0^{\tau} e^{-(r+\kappa)t} \frac{d}{d\theta} W_{\nu}(\nu_t, \hat{\lambda}_t) \hat{\lambda}_t\, dt \right]$$

where $a_t = (r+\kappa)\hat{\Psi}(\nu_t, \hat{\lambda}_t)/\psi$.

Proof. The expression for $d/d\theta W$ follows from the definition of W_t, and the expression for $d/d\theta \hat{\Psi}$ follows from the facts that

$$\hat{\Psi}_0 = E\left[\int_0^{\tau} e^{-(r+\kappa)t} \hat{\Delta}_t\, dt \right],$$

and that $\hat{\Delta} = W_{\nu} \hat{\lambda}$. ■

Corollary 1 *As θ rises above 0, principal's value function under the contract given by (55) changes at the rate $d/d\theta F(W, \hat{\Phi})$ given by Proposition 9. That is, the contract given by (55) remains approximately optimal for θ close to 0.*

Proof. Proposition implies, in particular, that the controls $(c, \hat{\Delta}, \hat{\Gamma})$, as functions of W and $\hat{\Phi}$, change in a differentiable manner. Since the first-order conditions of the principal's problem have to hold, it follows that this change in controls has only a second-order effect on the principal's value function. ■

To illustrate on a numerical example how closely the contract given by (55) approaches profit under the optimal contract, consider an example, consider an agent who manages a ten-billion dollar firm, whose volatility is 20%. Then $\sigma = 2000$ million dollars. Let $u(c) = \sqrt{c}$, $r = 5\%$, and $\kappa = 0.4$, i.e. the effect of the agent's effort decays by 40% per year. Figure 6 illustrates the principal's profit, as a function of W_0, for two examples: $\theta = .0003$, $L = 50$ and $\theta = .0002$, $L = 140$. Red dashed curves illustrate profit under

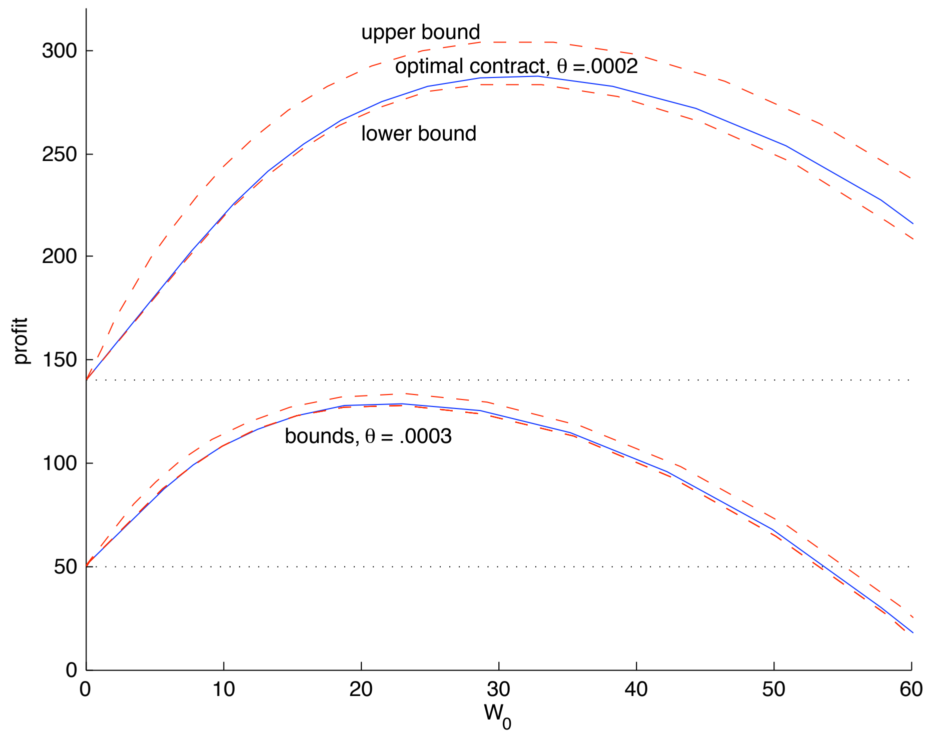

Figure 6: Bounds on the principal's profit.

the approximately optimal contract, and solid blue ones, under the optimal contract.[20] The large-firm contract does well for these values of θ.

7 Conclusions.

This paper aims to enhance our understanding of environments where the agent's actions can have delayed consequences. If a contract is thought of as a derivative on project value, which pays in the units of utility to the agent, and *Delta* is the sensitivity of derivative value to the performance signal, then the agent's incentives on the margin are captured by a discounted expectation of future contract *Deltas*. Contracts based on first-order incentive constraints are fully incentive-compatible if the discounted expectation of future contract

[20]The principal's profit (on the vertical axis) can be interpreted as the value that the agent adds to the firm, in millions of dollars.

Gammas is bounder by an appropriate constant. The first-order incentive constraints alone allow us to frame the problem of finding an optimal contract as an optimal stochastic control problem. I characterize a solution to this problem using the method of Lagrange multipliers.

The optimal contract becomes particularly tractable under the assumption that the signal about the agent's performance is noisy. In this case the first-order approach holds automatically. In these settings, the map from the signal about the agent's performance to the agent's compensation is determined in closed form. The problem of determining the optimal termination time is a real options problem.

The intuitive optimal contract is based on two variables: the multiplier on the agent's utility ν_t, which fully determines the agent's compensation flow, and the multiplier on incentives λ_t, which determines the sensitivity of ν_t to the performance signal. Generally, λ_t rises towards a target level during the agent's tenure and falls to 0 after the agent is fired. The agent is paid only when $\nu_t > 0$. When $\nu_t \leq 0$, the contract is "out of the money:" variable ν_t still adjusts to performance, but the agent is no longer paid. When the contract is sufficiently far out of money, the agent is fired, but he may still get paid after termination if favorable performance signals are realized after termination.

Appendix A: Proofs.

Proof of Proposition 5. Consider the process Φ is defined by (9), and let us show that there exists a process Γ in L^2 such that (27) holds. Note that the process

$$\bar{\Phi}_t \equiv E_t\left[\int_0^\infty (r+\kappa)e^{-(r+\kappa)s}\Delta_s\,ds\right] = \int_0^t (r+\kappa)e^{-(r+\kappa)s}\Delta_s\,ds + e^{-(r+\kappa)t}\Phi_t$$

Since $\bar{\Phi}_t$ is a martingale, by the Martingale Representation Theorem, there exists a process Γ in L^2 such that

$$d\bar{\Phi}_t = e^{-(r+\kappa)t}\,\Gamma_t\sigma\,dZ_t. \tag{56}$$

Differentiating $\bar{\Phi}_t$ with respect to t, we get

$$d\bar{\Phi}_t = (r+\kappa)e^{-(r+\kappa)t}\Delta_t\,dt - (r+\kappa)e^{-(r+\kappa)t}\Phi_t\,dt + e^{-(r+\kappa)t}\,d\Phi_t.$$

Combining with (56), we get

$$d\Phi_t = (r+\kappa)(\Phi_t - \Delta_t)\,dt + \Gamma_t \underbrace{(dX_t - \mu_t\,dt)}_{\sigma dZ_t},$$

as required. The transversality condition holds because Φ_0 is finite.

Conversely, suppose that Φ is simply a process that satisfies (27) and the transversality condition. Then $\bar{\Phi}^t_{t'}$ defined as

$$E_{t'}\left[\int_t^\infty (r+\kappa)e^{-(r+\kappa)(s-t)}\Delta_s\,ds\right] = \int_t^{t'} (r+\kappa)e^{-(r+\kappa)s}\Delta_s\,ds + e^{-(r+\kappa)(t'-t)}\Phi_{t'}$$

is a martingale. Therefore,

$$\Phi_t = \bar{\Phi}^t_t = \lim_{t'\to\infty} E_t[\bar{\Phi}^t_{t'}] =$$

$$\lim_{t'\to\infty} E_t\left[\int_t^{t'} (r+\kappa)e^{-(r+\kappa)s}\Delta_s\,ds\right] + \lim_{t'\to\infty} E_t[e^{-(r+\kappa)(t'-t)}\Phi_{t'}].$$

Then the transversality condition implies that Φ_t satisfies (9). ■

Proposition 11 *Function $\underline{G}(\nu,\tau)$ of Section 4 can be characterized as follows. Let*

$$\hat{\lambda}(\tau,t) = \int_0^\tau f(t-s)\,ds,$$

and let $\underline{F}^\tau(\nu,t)$ be the solution to the following parabolic partial differential equation

$$r\underline{F}^\tau(\nu,t) = \max_c\,(\nu u(c) - c) + \underline{F}^\tau_2(\nu,t) + \frac{\hat{\lambda}(\tau,t)^2}{2}\underline{F}^\tau_{11}(\nu,t). \tag{57}$$

Then $\underline{G}(\nu,\tau) = \underline{F}^\tau(\nu,\tau)$.

Proof. Equation (57) represents the expectation

$$E_t\left[\int_t^\infty e^{-r(s-t)}\chi(\nu_s)\,ds\right], \quad \text{when } d\nu_t = \hat{\lambda}(\tau,s)\,dZ_s,$$

and $\chi(\nu) \equiv \max_c \ \nu u(c) - c$, see Karatzas and Shreve (1991). Therefore, $\underline{F}^\tau(\nu, \tau)$ correctly represents the boundary condition for the optimization of the objective (18) ■

Proof of Proposition 6. Let us evaluate the agent's payoff under an alternative strategy $\hat{a}$. Denote by μ the level of fundamentals under the original strategy, and by $\hat{\mu}$, under $\hat{a}$. Consider the process

$$\hat{V}_t = \int_0^t e^{-rs}(u(c_s) - h(\hat{a}_s))\, ds + e^{-rt}\left(W_t + \frac{\Phi_t}{r+\kappa}(\hat{\mu}_t - \mu_t)\right).$$

Differentiating $\hat{V}_t$ with respect to t, we find that

$$e^{rt}\frac{d\hat{V}_t}{dt} = (u(c_t) - h(\hat{a}_t))\, dt - r\left(W_t + \frac{\Phi_t}{r+\kappa}(\hat{\mu}_t - \mu_t)\right)\, dt$$

$$+ \underbrace{(rW_t - u(c_t) + h(a_t))\, dt + \Delta_t\,(dX_t - \mu_t\, dt)}_{dW_t} +$$

$$\frac{1}{r+\kappa}\overbrace{((r+\kappa)(\Phi_t - \Delta_t)\, dt + \Gamma_t(dX_t - \mu_t\, dt))}^{d\Phi_t}(\hat{\mu}_t - \mu_t)$$

$$+\frac{\Phi_t}{r+\kappa}\underbrace{((r+\kappa)(\hat{a}_t - a_t) - \kappa(\hat{\mu}_t - \mu_t))\, dt}_{d(\hat{\mu}_t - \mu_t)}.$$

Using the fact that $dX_t = \hat{\mu}_t\, dt + \sigma\, dZ_t$, under the strategy $\hat{a}$ the drift of $\hat{V}_t$ is e^{-rt} times

$$\overbrace{h(a_t) + h'(a_t)(\hat{a}_t - a_t) - h(\hat{a}_t)}^{-\frac{\theta}{2}(\hat{a}_t - a_t)^2} + \frac{\Gamma_t}{r+\kappa}(\hat{\mu}_t - \mu_t)^2.$$

Thus, the agent's utility under the alternative strategy is

$$E^{\hat{a}}\left[\int_0^\infty e^{-rs}(u(c_s) - h(\hat{a}_s))\, ds\right] = E^{\hat{a}}[\hat{V}_\infty] = \hat{V}_0 + E^{\hat{a}}\left[\int_0^\infty d\hat{V}_t\right] =$$

$$W_0 + E^{\hat{a}}\left[\int_0^\infty e^{-rt}\left(\frac{\Gamma_t}{r+\kappa}(\hat{\mu}_t - \mu_t)^2 - \frac{\theta}{2}(\hat{a}_t - a_t)^2\right)\, dt\right].$$

This expression has a clear interpretation. An investment $\hat{a}_t - a_t$ in the divergence of fundamentals $\hat{\mu}_t - \mu_t$ carries a cost at time t, but creates benefit down the line if $\Gamma_t > 0$. Thus, the agent can get utility higher than W_0 only if Γ_t are large enough.

Let us show that if $dm_t = ((r+\kappa)\alpha_t - \kappa m_t)\, dt$ and $|\alpha_t| \leq \bar{\alpha}$ then

$$E^{\hat{a}}\left[\int_0^\infty e^{-rt}\left(\frac{\Gamma_t}{r+\kappa}m_t^2 - \frac{\theta}{2}\alpha^2\right)\, dt\right] \leq 0.$$

Consider the strategy α that maximizes the expression above, and denote by t the moment of time when the agent us choosing the maximal in absolute value action α_t. Without loss of generality, assume that $\alpha_t = \bar{\alpha}$. Then

$$\mu_{t+s} \leq \frac{r+\kappa}{\kappa}\bar{\alpha}$$

for all $s \geq 0$, so the derivative of payoff with respect to α_t is less than or equal to

$$-\theta\bar{\alpha} + E\left[\int_0^\infty e^{-rs}\frac{\Gamma_s}{r+\kappa}\frac{r+\kappa}{\kappa}\bar{\alpha}\cdot 2(r+\kappa)e^{-\kappa(s-t)}\, ds\right].$$

This expression is always negative if $\Gamma_s \leq \theta\kappa/2$, so the agent benefits by reducing deviation, the strategy is not optimal, a contradiction.

■

Alternative Proof of Proposition 6 (with a stronger sufficient condition). Denote by μ the level of fundamentals under the original strategy, and by $\hat{\mu}$, under a possible deviation strategy $\hat{a}$. We claim that after the agent deviated from time 0 until time t, his future expected payoff is bounded from above by

$$\hat{W}_t(\hat{\mu}_t) = W_t + \frac{\Phi_t}{r+\kappa}(\hat{\mu}_t - \mu_t) + L(\hat{\mu}_t - \mu_t)^2, \qquad (58)$$

where the constant L will be specified below. Then it follows immediately that when $\hat{\mu}_t = \mu_t$, the agent's continuation payoff is bounded from above by W_t, which is also the payoff he receives by following the strategy a. Thus, if the bound (58) is valid, then the the full set of incentive-compatiblity constraints (5) holds.

Consider the process

$$\hat{V}_t = \int_0^t e^{-rs}(u(c_s) - h(\hat{a}_s))\, ds + e^{-rt}\hat{W}_t(\hat{\mu}_t)$$

under the deviation strategy $\hat{a}$, so that

$$d\hat{\mu}_t = (r+\kappa)\hat{a}_t\, dt - \kappa\hat{\mu}_t\, dt, \quad \hat{\mu}_0 = 0.$$

To prove that the bound (58) is valid, it is enough to show that $\hat{V}$ is a supermartingale. Indeed, then

$$\hat{V}_t \geq E_t[\hat{V}_\infty] \;\; \Rightarrow \;\; \hat{W}_t(\hat{\mu}_t) \geq E_t\left[\int_t^\infty e^{-r(s-t)} u_s\, ds\right].$$

Differentiating $\hat{V}_t$ with respect to t, we find that

$$e^{rt}\frac{d\hat{V}_t}{dt} = (u(c_t) - h(\hat{a}_t))\, dt - r\overbrace{\left(W_t + \frac{\Phi_t}{r+\kappa}(\hat{\mu}_t - \mu_t) + L(\hat{\mu}_t - \mu_t)^2\right)}^{\hat{W}_t(\hat{\mu}_t)} dt$$

$$+\underbrace{(rW_t - u(c_t) + h(a_t))\, dt + \Delta_t\, (dX_t - \mu_t\, dt)}_{dW_t} +$$

$$\frac{1}{r+\kappa}\overbrace{((r+\kappa)(\Phi_t - \Delta_t)\, dt + \Gamma_t(dX_t - \mu_t\, dt))}^{d\Phi_t}(\hat{\mu}_t - \mu_t)$$

$$+\left(\frac{\Phi_t}{r+\kappa} + 2L(\hat{\mu}_t - \mu_t)\right)\underbrace{((r+\kappa)(\hat{a}_t - a_t) - \kappa(\hat{\mu}_t - \mu_t))\, dt}_{d(\hat{\mu}_t - \mu_t)}.$$

Using the fact that $dX_t = \hat{\mu}_t\, dt + \sigma\, dZ_t$ the drift of $\hat{V}_t$ is e^{-rt} times

$$\overbrace{h(a_t) + h'(a_t)(\hat{a}_t - a_t) - h(\hat{a}_t)}^{-\frac{\theta}{2}(\hat{a}_t - a_t)^2} +$$

$$\left(\frac{\Gamma_t}{r+\kappa} - (r+2\kappa)L\right)(\hat{\mu}_t - \mu_t)^2 + 2L(r+\kappa)(\hat{\mu}_t - \mu_t)(\hat{a}_t - a_t),$$

where we used $\Phi_t = h'(a_t)$ (and we have to set $a_t = \hat{a}_t$ if $t > \tau$.)

Now, in order to guarantee that $\hat{V}_t$ is a supermartingale for $t < \tau$, we need the matrix

$$\begin{bmatrix} -\theta/2 & L(r+\kappa) \\ L(r+\kappa) & \frac{\Gamma_t}{r+\kappa} - (r+2\kappa)L \end{bmatrix} \tag{59}$$

to be negative semidefinite. Note that

$$\max_L \frac{\theta(r+2\kappa)}{2}L - \frac{\theta\Gamma_t}{2(r+\kappa)} - L^2(r+\kappa)^2 = \frac{\theta^2(r+2\kappa)^2}{16(r+\kappa)^2} - \frac{\theta\Gamma_t}{2(r+\kappa)}.$$

If

$$\Gamma_t \leq \frac{\theta(r+2\kappa)^2}{8(r+\kappa)},$$

the determinant of the matrix (59) is guaranteed to be non-negative, and the matrix itself is negative semidefinite. If so, then the process $\hat{V}_t$ is a supermartingale if we set

$$L = \frac{\theta(r+2\kappa)}{4(r+\kappa)^2}.$$

For $t \geq \tau$, $a_t = \hat{a}_t = 0$, and so $\hat{V}_t$ is a supermartingale under a weaker condition

$$\Gamma_t \leq (r+2\kappa)(r+\kappa)L = \frac{\theta(r+2\kappa)^2}{4(r+\kappa)}.$$

■

Lemma 1 provides a partial converse to Proposition 6: it provides a necessary second-order condition for strategy a, which already satisfies the first-order condition $h'(a_t) = \Phi_t$, to be optimal.

Lemma 1 *A necessary second-order condition for the strategy a to be optimal is that the backward stochastic differential equation (BSDE)*

$$dL_t = \left(L_t(r+2\kappa) - 1_{t \leq \tau} \frac{2L_t^2(r+\kappa)^2}{\theta} - \frac{\Gamma_t}{r+\kappa} \right) dt + \sigma_t^L \, dZ_t, \tag{60}$$

must have a solution that remains finite until time τ, and satisfies the transversality condition $E[e^{-(r+2\kappa)t}L_t] \to 0$ after time τ.

Lemma 1 provides a condition on Γ_t, which guarantees that the agent cannot improve upon his continuation value through local deviations. Basically, condition (60) it requires that Γ_t cannot remain large for an extended period of time. Indeed, if $\Gamma_t \geq 0$ and L_t ever becomes negative, then it must violate the transversality condition (and may blow up in finite time to $-\infty$ before time τ, due to the quadratic term). Even if L_t starts near ∞, in finite time it falls towards the "stable" interval $(0, \theta(r+2\kappa)/2(r+\kappa^2))$, in which the terms excluding Γ_t are positive. However,

$$\Gamma_t > \frac{\theta(r+2\kappa)^2}{4(r+\kappa)}$$

ensures that the drift of L_t is negative even in the stable interval. So, if Γ_t remains large with a sufficiently large probability for long enough, L_t is guaranteed to "blow up" or at least violate the transversality condition.

Proof. Consider the second-order approximation of the agent's value function near $\hat{\mu}_t = \mu_t$,

$$\hat{W}_t(\hat{\mu}_t) = W_t + \frac{\Phi_t}{r+\kappa}(\hat{\mu}_t - \mu_t) + L_t(\hat{\mu}_t - \mu_t)^2.$$

Then under the optimal strategy

$$\hat{V}_t \equiv \int_0^t e^{-rs}(u(c_s) - h(\hat{a}_s))\, ds + e^{-rt}\hat{W}_t(\hat{\mu}_t)$$

must be a martingale, up to terms of order higher than $(\hat{\mu}_t - \mu_t)^2$. Using a calculation similar to that in the proof of Proposition 6, the drift of $\hat{V}_t$ is e^{-rt} times

$$-\frac{\theta}{2}(\hat{a}_t - a_t)^2 + 2(r+\kappa)L_t(\hat{\mu}_t - \mu_t)(\hat{a}_t - a_t) +$$
$$\left(\frac{\Gamma_t}{r+\kappa} - (r+2\kappa)L_t + \mu_t^L\right)(\hat{\mu}_t - \mu_t)^2 + O((\hat{\mu}_t - \mu_t)^3),$$

where μ_t^L is the drift of L_t. The drift of $\hat{V}_t$ at the optimal level of effort $\hat{a}_t = 1_{t\leq\tau} 2L_t(r+\kappa)/\theta$ must be 0, which implies that

$$\mu_t^L = L_t(r+2\kappa) - 1_{t\leq\tau}\frac{2L_t^2(r+\kappa)^2}{\theta} - \frac{\Gamma_t}{r+\kappa}.$$

This implies that L_t must satisfy the BSDE (60). It cannot "blow up" in finite time. Moreover, the transversality condition $E[e^{-(r+2\kappa)t}L_t] \to 0$ follows from the requirement that $E[e^{-rt}\hat{W}_t(\hat{\mu}_t)]$, since $(\hat{\mu}_t - \mu_t)^2$ shrinks at rate 2κ after time τ. ■

Proof of Proposition 8. First, let us show that if G solves (47) on $\mathcal{R} \subseteq [0,\infty) \times \mathbb{R}$ and satisfies the smooth-pasting conditions (49) on the boundary, then $W_t = G_\nu(\nu_t, \lambda_t)$, $\Phi_t = G_\lambda(\nu_t, \lambda_t)$ and the principal's continuation payoff is $G(\nu_t, \lambda_t) - \nu_t W_t - \lambda_t \Phi_t$ in the contract defined by (33).

Differentiating (47) with respect to ν and using the Envelope Theorem, we get

$$rG_\nu - u(c) + h(a(G_\lambda)) =$$

$$(a'(G_\lambda)\,(1 - \nu h'(a(G_\lambda))) - \kappa\lambda)\, G_{\nu\lambda} + \lambda^2 \frac{(r+\kappa)^2}{\sigma^2} \frac{G_{\nu\nu\nu}}{2}. \tag{61}$$

The right hand side represents the drift of the process $G_\nu(\nu_t, \lambda_t)$ when (ν_t, λ_t) follow (33). Also,

$$G_\nu(\nu_\tau, \lambda_\tau) = \underline{G}_\nu(\nu_\tau, \lambda_\tau) = W_\tau,$$

by Proposition 7. Therefore, as long as the transversality condition holds, Proposition 2 implies that $G_\nu(\nu_t, \lambda_t)$ is the agent's continuation value W_t under the effort strategy $\{a(G_\lambda(\nu_t, \lambda_t))\}$.

Similarly, differentiating (47) with respect to λ and using the Envelope Theorem, we get

$$(r+\kappa)G_\lambda - \underbrace{\lambda \frac{(r+\kappa)^2}{\sigma^2} G_{\nu\nu}}_{(r+\kappa)\Delta_t} =$$

$$\underbrace{(a'(G_\lambda)(1 - \nu h'(a(G_\lambda))) - \kappa\lambda)G_{\lambda\lambda} + \lambda^2 \frac{(r+\kappa)^2}{\sigma^2} \frac{G_{\lambda\nu\nu}}{2}}_{\text{drift of } G_\lambda}. \tag{62}$$

Since also $G_\lambda(\nu_\tau, \lambda_\tau) = \underline{G}_\lambda(\nu_\tau, \lambda_\tau) = \Phi_\tau$ by Proposition 7, Proposition 3 implies that $\Phi_t = G_\lambda(\nu_t, \lambda_t)$ under the effort strategy $\{a(G_\lambda(\nu_t, \lambda_t))\}$ (as long as the transversality condition holds).

Finally, subtracting ν times (61) and λ times (62) from (47), we get

$$r(G - \nu G_\nu - \lambda G_\lambda) = a(G_\lambda) - c + (a'(G_\lambda)\,(1 - \nu h'(a(G_\lambda))) - \kappa\lambda)\underbrace{(-\nu G_{\nu\lambda} - \lambda G_{\lambda\lambda})}_{\frac{\partial(G - \nu G_\nu - \lambda G_\lambda)}{\partial \lambda}}$$

$$+\frac{1}{2}\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\underbrace{(-G_{\nu\nu}-\nu G_{\nu\nu\nu}-\lambda G_{\lambda\nu\nu})}_{\frac{\partial^2(G-\nu G_\nu-\lambda G_\lambda)}{\partial\nu^2}}. \quad (63)$$

Hence, the process

$$\bar{F}_t=\int_0^t e^{-rs}\left(a_s-c_s\right)\,ds+e^{-rt}(G(\nu_t,\lambda_t)-\nu_t W_t-\lambda_t\Phi_t).$$

is a martingale. Since

$$\bar{F}_t=E_t[\bar{F}_\tau]=\int_0^t e^{-rs}\left(a_s-c_s\right)\,ds+$$

$$e^{-rt}E_t\left[\int_t^\tau e^{-r(s-t)}\left(a_s-c_s\right)\,ds+e^{-r(\tau-t)}(\underline{G}(\nu_\tau,\lambda_\tau)-\nu_\tau W_\tau-\lambda_\tau\Phi_\tau)\right],$$

where $\underline{G}(\nu_\tau,\lambda_\tau)-\nu_\tau W_\tau-\lambda_\tau\Phi_\tau$ is the principal's continuation payoff at time τ by Proposition 7, it follows that $G(\nu_t,\lambda_t)-\nu_t W_t-\lambda_t\Phi_t$ is the principal's continuation payoff in the contract defined by (33).

Next, we will show that under any alternative contract, for which $W_0=G_\nu(\nu_0,\lambda_0)$ and $\Phi_0=G_\lambda(\nu_0,\lambda_0)$, the principal's profit is bounded from above by $G(\nu_0,\lambda_0)-\nu_0W_0-\lambda_0\Phi_0$. The key step in the argument is showing that the process $\bar{F}_t$ is a supermartingale for appropriate processes (ν_t,λ_t) chosen to match the law of motion of (W_t,Φ_t) under the alternative contract.

Lemma 2 *Consider an alternative contract, characterized by controls* (c,Δ,Γ) *and termination time* τ*, and denote by* W *and* Φ *the state variables under those controls. Define* $G(\nu,\lambda)=\underline{G}(\nu,\lambda)$ *outside* $\mathcal{R}$*. If the Hessian of* G *is positive definite, then there exist processes*

$$d\nu_t=\mu_t^\nu\,dt+\sigma_t^\nu\,dZ_t \quad \text{and} \quad d\lambda_t=\mu_t^\lambda\,dt+\sigma_t^\lambda\,dZ_t \quad (64)$$

such that $W_t=G_\nu(\nu_t,\lambda_t)$ *and* $\Phi_t=G_\lambda(\nu_t,\lambda_t)$ *for* $t\le\tau$*.*

Proof. We would like to make sure that there are processes σ_t^ν, σ_t^λ, μ_t^ν and μ_t^λ such that the laws of motion of $G_\nu(\nu_t,\lambda_t)$ and $G_\lambda(\nu_t,\lambda_t)$ are identical to those of W_t and Φ_t. To match volatilities, Ito's lemma requires that σ_t^ν and σ_t^λ be determined by equations

$$\underbrace{\begin{bmatrix} G_{\nu\nu} & G_{\nu\lambda} \\ G_{\lambda\nu} & G_{\lambda\lambda}\end{bmatrix}}_{H(G)}\begin{bmatrix}\sigma_t^\nu \\ \sigma_t^\lambda\end{bmatrix}=\begin{bmatrix}\Delta_t\sigma \\ \Gamma_t\sigma\end{bmatrix}. \quad (65)$$

There is a unique solution because $H(G)$, the Hessian of G, is invertible.

Similarly, to match drifts, let μ_t^ν and μ_t^λ be determined from equations

$$H(G)\begin{bmatrix}\mu_t^\nu\\ \mu_t^\lambda\end{bmatrix}+\ldots=\begin{bmatrix}rW_t-u(c_t)+h(a(\Phi_t))\\ (r+\kappa)(\Phi_t-\Delta_t)\end{bmatrix},$$

where "$\ldots$" stand for terms that depend on the volatilities of ν_t and λ_t and not the drifts. Again, the solution exists because the Hessian of G is invertible. ■

In order to prove that the alternative contract cannot be superior to the contract defined in Proposition 8, we will first show that the drift of the process $\bar{F}_t$ defined above is non-positive when ν_t and λ_t follow (64).

Using Ito's lemma and the laws of motion of W_t and Φ_t, the drift of $G(\lambda_t,\nu_t)-\nu_tW_t-\lambda_t\Phi_t$ is

$$G_\nu\mu_t^\nu+G_\lambda\mu_t^\lambda+\frac{1}{2}[\sigma_t^\nu\ \sigma_t^\lambda]H(G)\begin{bmatrix}\sigma_t^\nu\\ \sigma_t^\lambda\end{bmatrix}-\mu_t^\nu W_t-\mu_t^\lambda\Phi_t-[\sigma_t^\nu\ \sigma_t^\lambda]\begin{bmatrix}\Delta_t\sigma\\ \Gamma_t\sigma\end{bmatrix}$$

$$-\nu_t(rW_t-u(c_t)+h(a_t))-\lambda_t(r+\kappa)(\Phi_t-\Delta_t)=$$

$$-\frac{1}{2}[\sigma_t^\nu\ \sigma_t^\lambda]H(G)\begin{bmatrix}\sigma_t^\nu\\ \sigma_t^\lambda\end{bmatrix}-\nu_t(rW_t-u(c_t)+h(a_t))-\lambda_t(r+\kappa)(\Phi_t-\Delta_t),$$

where we used (65). Without loss of generality, we can assume that $c_t=\chi(\nu_t)$, which maximizes the drift of $\bar{F}_t$.

For comparison, when λ_t and ν_t follow (33) then the drift of $G(\nu_t,\lambda_t)-\nu_tW_t-\lambda_t\Phi_t$ is

$$-(r+\kappa)^2\frac{\lambda_t^2}{\sigma^2}\frac{G_{\nu\nu}}{2}-\nu_t(rW_t-u(c_t)+h(a_t))-\lambda_t(r+\kappa)\left(\Phi_t-\lambda_t\frac{r+\kappa}{\sigma^2}G_{\nu\nu}\right),$$

which, according to (63), leads to a drift of $\bar{F}_t$ of zero in $\mathcal{R}$ and negative outside $\mathcal{R}$, by (50).

Now, when λ_t and ν_t follow (64) instead of (33) the drift of $\bar{F}_t$ changes by e^{-rt} times

$$-\frac{1}{2}[\sigma_t^\nu\ \sigma_t^\lambda]\,H(G)\begin{bmatrix}\sigma_t^\nu\\ \sigma_t^\lambda\end{bmatrix}+\lambda_t(r+\kappa)\underbrace{\frac{\sigma_t^\nu G_{\nu\nu}+\sigma_t^\lambda G_{\nu\lambda}}{\sigma}}_{\Delta_t}-\frac{1}{2}(r+\kappa)^2\frac{\lambda_t^2}{\sigma^2}G_{\nu\nu}=$$

$$-\frac{1}{2}[\sigma_t^\nu-(r+\kappa)\lambda_t/\sigma,\ \ \sigma_t^\lambda]\,H(G)\begin{bmatrix}\sigma_t^\nu-(r+\kappa)\lambda_t/\sigma\\ \sigma_t^\lambda\end{bmatrix}\le 0,$$

since the matrix $H(G)$ is positive definite. Hence, the drift of $\bar{F}_t$ under the alternative contract cannot be greater than that under the contract, in which λ_t and ν_t follow (33), so it must be negative. In other words, $\bar{F}_t$ is a supermartingale.

Hence,

$$\bar{F}_0 = G(\nu_0, \lambda_0) - \nu_0 W_0 - \lambda_0 \Phi_0 \geq E[\bar{F}_\tau] =$$

$$E\left[\int_0^\tau e^{-rs}(a_s - c_s)\, ds + e^{-r\tau}(\underline{G}(\nu_\tau, \lambda_\tau) - \nu_\tau W_\tau - \lambda_\tau \Phi_\tau)\right]$$

$$\geq E\left[\int_0^\tau e^{-rs} a_s\, ds - \int_0^\infty e^{-rs} c_s\, ds\right],$$

where we used Proposition 7 for the last inequality.[21] Therefore, the contract, in which λ_t and ν_t follow (33), is optimal. ■

Proposition 12 *In the large-firm case, the function G is strictly convex.*

Proof. Let us show that G is a strictly convex function. Consider an initial condition $(\nu_0'', \hat{\lambda}_0'') = \omega(\nu_0, \hat{\lambda}_0) + (1-\omega)(\nu_0', \hat{\lambda}_0')$, for $\omega \in (0,1)$, which is a convex combination of two other initial conditions. Let us show that

$$G(\nu_0'', \hat{\lambda}_0'') < \omega G(\nu_0, \hat{\lambda}_0) + (1-\omega) G(\nu_0', \hat{\lambda}_0'). \tag{66}$$

If τ'' is the optimal stopping time under the contract that achieves the pair $(\nu_0'', \hat{\lambda}_0'')$, then certainly

$$G(\nu_0, \hat{\lambda}_0) \geq E\left[\int_0^{\tau''} e^{-rt} \chi(\nu_t)\, dt + e^{-r\tau''} L\right],$$

when $d\nu_t = \hat{\lambda}_t\, dZ_t$ and

$$\hat{\lambda}_t = \frac{1_{t \leq \tau''}(r+\kappa)}{\theta\sigma} - \kappa \hat{\lambda}_t\, dt.$$

The same is true about $G(\nu_0', \hat{\lambda}_0')$ when $d\nu_t' = \hat{\lambda}_t'\, dZ_t$ and

$$\hat{\lambda}_t' = \frac{1_{t \leq \tau''}(r+\kappa)}{\theta\sigma} - \kappa \hat{\lambda}_t'\, dt.$$

[21] The transversality condition $\liminf E[1_{t<\tau} e^{-rt}(G(\lambda_t, \nu_t) - \lambda_t W_t - \nu_t \Phi_t)] \geq 0$ needs to hold in order to extend the supermartingale $\bar{F}$ to time τ.

Note also that, under these definitions,

$$(\nu_t'', \hat{\lambda}_t'') = \omega(\nu_t, \hat{\lambda}_t) + (1-\omega)(\nu_t', \hat{\lambda}_t')$$

for all histories $\{Z_s, s \in [0,t]\}$. Therefore,

$$\chi(\nu_t'') \leq \omega\chi(\nu_t) + (1-\omega)\chi(\nu_t'),$$

with strict inequality almost everywhere. This implies (66). ■

Proof of Proposition 9. Using the Envelope Theorem, we can evaluate the derivative of the principal's profit with respect to θ can be computed while keeping the policy (c, Δ, Γ), as a function of $(W, \hat{\Phi})$, fixed. This policy changes the laws of motion of W and $\hat{\Phi}$ slightly:

$$dW_t = (rW_t - u(c_t) + \theta a_t^2/2)\, dt + \hat{\Delta}_t\, dZ_t, \quad d\hat{\Phi}_t = ((r+\kappa)\hat{\Phi}_t - \hat{\Delta}_t)\, dt + \hat{\Gamma}_t\, dZ_t.$$

Since only the drift of W changes, and the derivative of the drift with respect to θ is $a_t^2/2$, it follows that the effect of θ on the principal's value function can be evaluated as

$$F^\theta(W, \hat{\Phi}) = E\left[\int_0^\tau e^{-rt} F_W(W, \hat{\Phi}) \frac{a_t^2}{2}\, dt\right],$$

which is equivalent to (53), since $\nu_t = -F_W(W, \hat{\Phi})$.

Now, by the Envelope theorem, we can evaluate the derivative of $G(\nu, \hat{\lambda}) = \max_{W,\hat{\Phi}} F(W, \hat{\Phi}) + \nu W + \hat{\lambda}\hat{\Phi}$ with respect to θ while keeping W and $\hat{\Phi}$ fixed. Therefore, $G^\theta(\nu, \hat{\lambda}) = F^\theta(W, \hat{\Phi})$. ■

Bibliography.

Abreu, D., Pearce, D., and Stacchetti, E. (1990) "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." Econometrica, Vol. 58, pp. 1041-1063.

Ang, J., R. Cole and J. Lin (2000) "Agency Costs and Ownership Structure," *Journal of Finance*, 55(1), 81-106.

Biais, B., T. Mariotti, G. Plantin and J.-C. Rochet (2007): "Optimal Design of Dynamic Pricing of Securities," *Review of Economic Studies*.

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

DeMarzo, P. and Y. Sannikov (2011) "Learning in Dynamic Incentive Contracts," working paper, Princeton University.

Edmans, Gabaix, Sadzik and Sannikov (2012) "Dynamic CEO Compensation," forthcoming in Journal of Finance.

Eso, Peter and Balasz Szentes (2013) "Dynamic Contracting with Adverse Selection: An Irrelevance Result," LSE working paper.

Fama E. and K. French (2000) "Forecasting Profitability and Earnings," *Journal of Business* 73, 161-175.

Gabaix, X. and A. Landier (2008) "Why Has CEO Pay Increased So Much?," *Quarterly Journal of Economics*, 123(1), 49-100

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

He, Z., B. Wei and J. Yu (2012) "Optimal Long-Term Contracting with Learning," working paper

Holmstrom, B. and P. Milgrom (1987) "Aggregation and Linearity in the Provision of Intertemporal Incentives." Econometrica, Vol 55, pp. 303-328

Hopenhayn, H. and A. Jarque (2010) "Unobservable Persistent Productivity and Long Term Contracts," Review of Economic Dynamics, 13(2), 333-349.

Jarque, A. (2011) "Repeated Moral Hazard and Effort Persistence," forthcoming in the *Journal of Economic Theory.*

Jensen, Michael and Kevin J. Murphy (1990) "Performance Pay and Top-Management Incentives." *Journal of Political Economy*, 98, 225-64.

Jovanovic, Boyan and Julien Prat (2013) "Dynamic Contracts when Agent's Quality is Unknown," forthcoming in *Theoretical Economics.*

Kapicka, M. (2011) "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," forthcoming in *Review of Economic Studies.*

Kocherlakota, N. (2004) "Figuring out the Impact of Hidden Savings on Optimal Unemployment Insurance," *Review of Economic Dynamics*, 7(3), 541-554.

Lazear, E.P. (1979) "Why Is There Mandatory Retirement?" *Journal of Political Economy*, Vol. 87, no. 6, 1261-1284.

Morck, R., A. Shleifer and R. W. Vishny (1988), "Management Ownership and Market Valuation: An Empirical Analysis," *Journal of Financial Economics,* 20, 293-315.

Murphy, K. (1999) "Executive Compensation," *Handbook of Labor Economics*, 3(B), 2485-2563.

Phelan, C. and Townsend, R. (1991) "Computing Multi-Period, Information-Constrained Optima." *Review of Economic Studies*, Vol. 58, pp. 853-881

Radner, R. (1985) "Repeated Principal-Agent Games with Discounting." Econometrica, Vol. 53, pp. 1173-1198

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

Sannikov, Y. (2013) "Contracts: The Theory of Dynamic Principal-Agent Relationships and the Continuous-Time Approach," in Daron Acemoglu, Manuel Arellano and Eddie Dekel (eds.), *Advances in Economics and Econometrics, 10th World Congress of the Econometric Society*, Cambridge University Press.

Spear, S. E. and S. Srivastava (1987) "On Repeated Moral Hazard with Discounting." Review of Economic Studies, Vol. 54, pp. 599-617

Williams, N. (2011) “Persistent Private Information,” Econometrica, 79(4), 1233-1274.

Yong, J. and X. Y. Zhou (1999) *Stochastic Controls: Hamiltonian Systems and HJB Equations*, Springer.