Moral Hazard and Long-Run Incentives*

Yuliy Sannikov

May 27, 2013

Abstract

This paper considers dynamic moral hazard settings, in which the agent's actions have consequences over a long horizon. Limited liability makes it difficult to tie the agent's compensation to long-run outcomes. 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.

*I am especially 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.

This paper builds a dynamic agency model, in which 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. I identify the variable Φ_t that summarizes the agent's incentives at any point of time.

Under the assumption that the impact of the agent's effort on future outcomes is exponentially decaying, the principal's problem, subject to only

[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.

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.

While a control problem with two state variables can be daunting, I find a clear and intuitive way to represent the solution in terms of the 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 rate κ that corresponds to the declining impact of the agent's past effort on output.

One component of the optimal contract, the map from λ_t to the volatility of the agent's compensation, is determined explicitly up front. To determine

[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 $x/\gamma\%$.

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.

Of course, it is adequate to derive the optimal contract using the agent's first-order incentive constraints only if these constraints are in fact sufficient. I derive a condition which guarantees that this is the case: 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 change sufficiently fast after he deviates to actually make it profitable deviate far away from the recommended strategy. In contrast, contracts in which Γ_t satisfies the bound are robust. I find that optimal contracts in settings with low signal-to-noise ratio typically satisfy the bound. Many applications, such as that of executive compensation, naturally have a low signal-to-noise ratio. When the contract obtained by imposing only the first-order incentive constraints violates the sufficient condition on Γ_t, then, assuming the agent does not deviate, one obtains only an upper bound on profit from the optimal contract. The optimal contract itself may be extremely hard to find. However, it is possible to impose the bound on Γ_t explicitly to get a *robust* contract that provides a lower bound on the principal's profit. In many cases, especially when the condition on Γ_t is violated only on particular portions of the state space, the two bounds may well end up extremely close, and the robust contract is also approximately optimal.

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 is exponentially decaying. It provides a sufficient condition, under which first-order incentive constraints guarantee the optimality of the agent's strategy. It also 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) and Garrett and Pavan (2012)). 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 total 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. The nonnegative function $f : [0, \infty) \to \mathbb{R}$ is assumed to satisfy

$$\int_0^\infty e^{-rt} f(t) \, dt = 0 \tag{2}$$

where r is the discount rate, to fix the rate at which effort adds to the present value of future output.

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 $h : \mathcal{A} \to [0, \infty)$ are C^2 functions that satisfy $u(0) = 0$, $u' > 0$, $u'' < 0$, $h(0) = 0$, $h'(0) = 0$ and $h'' > 0$, so the agent's preferred action is normalized to 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]

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

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

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

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

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

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

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.

[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.

To analyze 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*,

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

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 effort

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

determines the drift of current output.

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

There are two convenient expressions for the key process Φ_t that defines the agent's incentives on the margin. Besides (9), there is an alternative expression (shown to be equivalent in 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,*

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

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 $t \leq t' \leq s$. Then

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

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

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

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,

$$0 = E_t^a \left[\int_t^s e^{-r(t'-t)} f(t'-t) \Delta_{t'} \, dt' \right] + \underbrace{E_t^a \left[\int_t^s d\Phi_{t'}^t \right]}_{E_t^a[\Phi_s^t] - \Phi_t}.$$

Taking the limit $s \to \infty$, also assuming that the transversality condition $E_t^a [\Phi_s^t] \to 0$ 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. From the point of view of the principal, it is convenient to 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. This expression normalizes for the drift of X and takes into account only surprise innovations in output relative to what is expected. From the point of view of the principal, on the equilibrium path, once appropriate incentive conditions are imposed to fix the agent's strategy, Z defined by (17) is a Brownian motion. We only need to work with X, explicitly taking into account how the agent's effort affects the path of X, when analyzing the agent's incentives, i.e. in Proposition 1, as well as Proposition 6 that provides a sufficient incentive-compatibility condition.

4 The Large-Firm Case.

This section focuses on the *large-firm case*, 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.

Below I formalize the assumptions behind the large firm case and characterize the optimal contract. The solution is very tractable: the model implies a closed-form map from the performance signal to optimal compensation. Also, the optimal termination time is similar to the optimal exercise time of an American option.

The large-firm case is particularly transparent when one assumes quadratic effort cost,

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

Assuming that $\bar{a}$ is sufficiently large, the incentive constraint (12) reduces to

$$a_t = \hat{\Phi}_t \equiv E_t \left[\int_t^\infty e^{-r(s-t)} \hat{\zeta}_t^s \, u_s \, ds \right], \quad \text{where} \quad \hat{\zeta}_t^s \equiv \frac{1}{\psi} \int_t^s f(s'-t) dZ_{s'}. \qquad (18)$$

Then there are convenient expressions for the agent's effort and the principal's profit, in terms of the correlations between the normalized output histories and the agent's compensation utility.

[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.

Proposition 4 *If the time-0 expectations of the principal's profit and the agent's payoff are F_0 and W_0, then*

$$F_0 + \nu_0 W_0 = E\left[\int_0^\infty e^{-rt}\left(\nu_t u_t - c_t\right)\, dt + e^{-r\tau} L\right], \tag{19}$$

where

$$\nu_t = \nu_0 + \int_0^t \hat{\lambda}_s \, dZ_s \qquad \text{and} \quad \hat{\lambda}_t = \frac{1}{\psi}\int_0^{\min(t,\tau)} f(t-s)\, ds. \tag{20}$$

Proof. Substituting the incentive constraint (18) into the profit expression (3), we get

$$E^a\left[\int_0^\tau e^{-rt}\left(\int_t^\infty e^{-r(s-t)}\hat{\zeta}_t^s \, u_c \, ds\right) dt + e^{-r\tau} L - \int_0^\infty e^{-rt} c_t \, dt\right].$$

After changing the order of integration and adding $\nu_0 W_0$, we obtain (19). ■

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

In this limit, 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 will use expression (19) to characterize the optimal contract in closed form in this limit. In general, (19) can be used to derive both upper and lower bounds on the optimal contract profit, and these bounds get tighter as one approaches the limit.

In the limit, the agent's compensation before termination can be found as follows. Since the cost of effort is negligible

$$F_0 + \nu_0 W_0 = E\left[\int_0^\infty e^{-rt}\left(\nu_t u_t - c_t\right)\, dt + e^{-r\tau} L\right]. \tag{22}$$

Choosing the multiplier ν_0 on the agent's utility appropriately to match the constraint on W_0, we can maximize this expression by choosing consumption c_t directly after all histories $\{Z_s, s \in [0,t]\}$. Specifically,

$$c_t = \arg\max_c \nu_t u(c) - c, \tag{23}$$

where the multiplier on the agent's utility ν_t is given by (20).

This is an incredibly simple characterization of the optimal contract. The specification of the optimal termination time τ aside, (23) together with (20) describe exactly how output paths map into the agent's compensation. Payments to the agent are continuously increasing in the multiplier ν_t on the agent's current utility, and may fall to zero (e.g. if $\nu_t \leq 0$). The multiplier ν_t follows a specific law of motion, and its sensitivity to normalized output $\hat{\lambda}_t$ is deterministic both before and after time τ. According to (20), $\hat{\lambda}_t$ increases from 0 towards a target level before termination, and decays to 0 after time τ. This characterization suggests a contract, which gradually raises the agent's deferred stock compensation towards a target level during employment, and lets it vest gradually after termination.

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. Denote by $G(\nu, t)$ and $\underline{G}(\nu_\tau, \tau)$ the value functions for the maximization problem (22) before termination and at termination.[6] Proposition 9 in the Appendix characterizes the function $\underline{G}$.

[6] 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 $\chi(\nu) \equiv \max_c\ \nu u(c) - c$.

Before termination, the value function $G(\nu, t)$ must solve the following simple parabolic equation

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

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 $\mathcal{R}$.[7] Since ν_t is the multiplier on the agent's utility, it follows immediately that the agent's continuation payoff is given by

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

4.1 Approximately Optimal Contracts.

When σ is large and θ is small, the contract defined by (23) and (20), with an appropriate termination time τ, is *approximately optimal* and it gives a *lower* bound on the principal's profit. The contract given by (23) optimizes with respect to the agent's compensation without taking into account that 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.[8]

One may wonder how closely profit from the approximately optimal contract approaches that from the optimal contract. It is possible to address

[7] 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.

[8] 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).

this issue directly, by finding the optimal contract, or indirectly be bounding the principal's profit from above. There are many approaches to getting an upper bound on the principal's profit. One natural starting point is to maximize (19) through an explicit choice of c_t and a_t inside the agent's utility flow $u_t = u(c_t) - 1_{t\leq\tau}h(a_t)$. This procedure disconnects the agent's cost of effort from expected output, and it provides a reasonable approximation when of optimal contract profit when the cost of effort is insignificant. A bound $\bar{a}$ on the set of effort levels needs to be imposed, so the optimal choice of a_t leads to $h(a_t) = h(\bar{a})$ when $\nu_t < 0$, and $h(a_t) = h(0)$ when $\nu_t \geq 0$.

To illustrate these bounds on the principal's profit on an example, consider an agent who manages a ten-billion dollar firm, whose volatility is 20%. Then $\sigma = 2000$ million dollars. Let $u(c) = \sqrt{c}$ and $r = 5\%$. The principal's profit depends on the effort parameter θ, the principal's outside option L, as well as the agent's payoff W_0. Profit is non-monotonic in W_0: the value of W_0 that maximizes profit is strictly positive in order to create enough room to give the agent incentives. In the examples below, I set L endogenously to be about 50% of the optimal contract profit.[9] Function f is chosen to be exponential with decay parameter $\kappa = 0.4$, although the specific value of κ has much less impact on the principal's profit than parameters such as σ and θ.

Figure 1 illustrates the principal's profit, as a function of W_0, for two examples: $\theta = .0003$ ($L = 50$ and $\bar{a} = 60$) and $\theta = .0002$ ($L = 140$ and $\bar{a} = 30$). The upper and lower bounds are illustrated by dashed curves. Solid curves between the two dashed bounds illustrate exact profit under the optimal contract (which is characterized in the next section). The bounds give reasonable estimates in the large-firm case, and the approximately optimal contract based on equations (19) approaches closely the optimal contract profit.

In these examples, the agent's effort varies over the contract: it generally grows with time and with ν (except for very high ν). The agent's effort is also highly sensitive to parameter θ (as well as σ). For $\theta = .0003$, effort creates value between \$10 and \$25 million a year. For $\theta = .0002$, effort

[9]The reader may be wondering why I take the relative risk-aversion coefficient to be $\gamma = 0.5$ in this example, instead of a more reasonable value of between 2 and 3. The reason is that CRRA utility with high risk aversion is unbounded from below, violating the assumption that $u(0) = 0$ made in Section 2. If the assumption is violated, there are unintended consequences as the principal can punish the agent arbitrarily harshly without firing him. I plan to address this issue in future research.

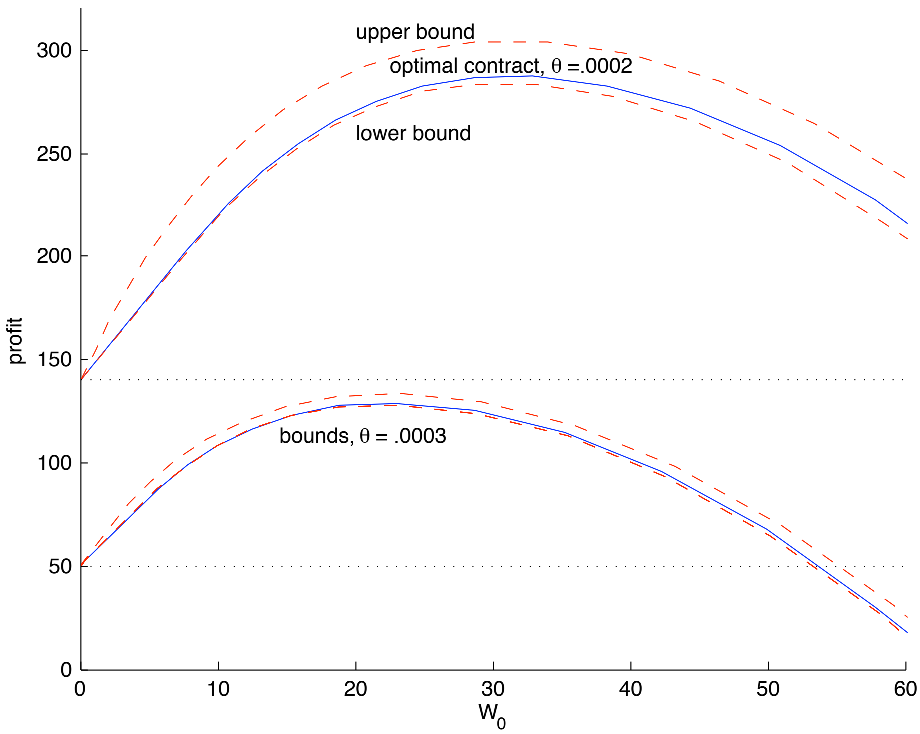

Figure 1: Bounds on the principal's profit.

creates between \$20 and \$50 million a year. This, the signal-to-noise ratio for effort detection in these examples is generally between 0.5% and 2.5%.

The agent is exposed to significant risk. For example, when $\theta = .0003$, if the principal chooses $W_0 \approx 22$ to maximize profit, the certainty equivalent value of the agent's contract is a bit over \$1 million a year. The agent starts by consuming 0, and he receives positive payments only if realized performance X exceeds expectation. When $\hat{\lambda}_t$ reaches its target level of $\frac{1}{\theta\sigma}\frac{r+\kappa}{\kappa}$, the standard deviation of ν_t is 1.875 per year. The agent's pay increases to \$1 million per year after 1 standard deviation rise in ν_t, \$3.5 million after two standard deviations, and \$8 million after three standard deviations. The agent is fired if ν_t drops by about five standard deviations. When the agent is fired, the certainty equivalent value of his deferred compensation is a perpetuity of about \$100 thousand a year.

5 Contract Design via Optimal Control.

In general, it is possible to investigate the principal's problem using methods from optimal stochastic control. To ensure that the principal's problem has a recursive structure, this section assumes 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 process Φ_t, which characterizes the agent's incentives on the margin, has a recursive representation characterized by the following proposition.

Proposition 5 *Fix a contract (c, τ) and a strategy a, for which Φ_0 is finite. Then Φ_t is the process 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

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

can be framed as an optimal stochastic control problem. In this problem, the full set of incentive-compatibility constraints (5) is replaced with just the first-order condition (12). From Propositions 2 and 5, the control problem is defined by 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. I will refer to this problem as the *relaxed control problem.* The method of solving the relaxed problem, hoping to find the optimal contract, is called the *first-order approach.*

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*[10]

$$\Gamma_t \leq \frac{\theta(2\kappa + r)^2}{8(r + \kappa)}. \tag{31}$$

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

[10]Condition (31) can be weakened to $\Gamma_t \leq \frac{\theta(2\kappa+r)^2}{4(r+\kappa)}$ for $t \geq \tau$.

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.

The HJB equation for the relaxed control problem takes the form

$$rF(W,\Phi) = \max_{c,\Delta,\Gamma} \ a(\Phi) - c + [F_W \ F_\Phi] \left[\begin{array}{c} rW - u(c) + h(a(\Phi)) \\ (r+\kappa)(\Phi - \Delta) \end{array} \right] + \frac{\sigma^2}{2} \frac{\partial^2}{\partial \epsilon^2} F(W + \Delta\epsilon, \Phi + \Gamma\epsilon). \tag{32}$$

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 the endogenous direction. I will show that the solution to the control problem becomes significantly more tractable (and intuitive) in the domain of Lagrange multipliers on the variables W_t and Φ_t. The solution is analogous that given by (20) and (23) in the large-firm case, but taking into account the agent's nonnegligible disutility of effort.

Let me contrast these control problems to that, which arises in the standard case when $\kappa = \infty$ and the agent's effort has immediate impact on output. In that case, the principal solves a control problem that has the same objective (28) but a single state variable W_t, which follows

$$dW_t = (rW_t - u(c_t) + 1_{t\le\tau} h(a(\Phi_t)))\, dt + \Phi_t\, \sigma\, dZ_t. \tag{33}$$

The principal directly chooses Φ_t as a control, as well as c_t and the stopping time τ. The solution to this problem always leads to the optimal contract, since the incentive constraint $a_t = a(\Phi_t)$ is not only necessary but also sufficient in this case.[11] Sannikov (2008) characterizes the optimal contract via the HJB equation

$$rF(W) = \max_{c,\Phi}\ a(\Phi) - c + (rW - u(c) + h(a(\Phi)))F'(W) + \frac{1}{2}\Phi^2\sigma^2 F''(W). \tag{34}$$

In the optimal contract, the agent is employed while $W \in [0, u(\bar{c})/r]$, and termination time τ is reached when W_t reaches one of the endpoints of this interval. Function F must satisfy the conditions $F(0) = L$, $F(u(\bar{c})) = L - \bar{c}/r$ and $F'(u(\bar{c}))u'(\bar{c}) = -1/r$ at the boundaries.

5.1 The Lagrangian Approach: the Case of $\kappa = \infty$.

The Lagrangian approach often provides a much more tractable characterization of the optimization problem than the direct approach. We already saw its power in Section 4, where we were able to characterize the exact map from the histories of output to the agent's compensation in the large-firm case. This section demonstrates how the Lagrangian approach works in the standard case when $\kappa = \infty$.

One useful way to derive the laws of motion of Lagrange multipliers is by differentiating the HJB equation with respect to state variables (and another way relies on the Hamiltonian - see for example Yong and Zhou (1999)). Differentiating (34) with respect to W and using Envelope theorem, we get

$$0 = (rW - u(c) + h(a(\Phi)))F''(W) + \frac{1}{2}\Phi^2\sigma^2 F'''(W).$$

From Ito's lemma and the law of motion (33) of W_t, the right hand side of this equation is the drift of $F'(W)$, hence, the drift of the multiplier ν_t is 0. The (absolute) volatility of ν_t is $-\Phi\sigma F''(W)$. Since

$$a'(\Phi)(1 - \nu_t h'(a(\Phi_t))) + \Phi\sigma^2 F''(W) = 0$$

[11]Note also that as $\kappa \to \infty$, the right hand side of (31) becomes infinite.

by the first-order condition with respect to a, it follows that

$$d\nu_t = \frac{a'(\Phi_t)(1 - \nu_t h'(a(\Phi_t)))}{\sigma}\, dZ_t, \qquad c_t = \arg\max_c \nu_t u(c) - c. \tag{35}$$

Equations (35) do not completely characterize the optimal contract, as they include an unknown function Φ of ν_t, but they provide important intuition behind the optimal contract. Let me list my interpretations, some of them loose, in the bullet points below:

- The multiplier ν_t, the inverse of the agent's marginal utility, is the marginal cost in money of an extra util. When the agent is risk averse, ν_t maps one-to-one into the agent's compensation. To interpret (35) empirically, it is useful to note that for the relative risk aversion coefficient γ, the *relative* volatility of the agent's consumption is $1/\gamma$ times the *relative* volatility of ν_t. The relative volatility of ν_t is

$$\frac{a'(\Phi_t)(1 - \nu_t h'(a(\Phi_t)))}{\nu_t \sigma}. \tag{36}$$

- $\nu_t h'(a(\Phi_t))$ is the monetary marginal disutility of effort required to produce one dollar of output. The most direct empirical interpretation of this quantity is the pay-performance sensitivity. If noise σ is small, then the optimal contract basically sells the firm to the agent, and so pay-performance sensitivity is nearly 1. The gap between $\nu_t h'(a(\Phi_t))$ and 1 is determined so that the volatility of the present value of the agent's compensation is approximately equal to the volatility of output (and so the agent effectively absorbs all of project risk). As σ gets larger, pay-performance sensitivity has to decline. In particular, for applications such as CEO compensation, this term is essentially 0.

- Φ_t is the sensitivity of *utility* to an extra dollar of firm value. Therefore, a loose interpretation for $a'(\Phi_t)/\nu_t$ is the marginal value of extra pay-performance sensitivity. Consequently, $a'(\Phi_t)/(\nu_t\sigma)$ is the marginal value of exposing the agent's wealth to an extra dollar of risk.

While these interpretations are loose, they are a useful starting point for thinking about model calibration. For example, consider a ten-billion dollar firm, with annual volatility 20%. If the CEO's target annual pay is $5 million, then what level of risk exposure is optimal?

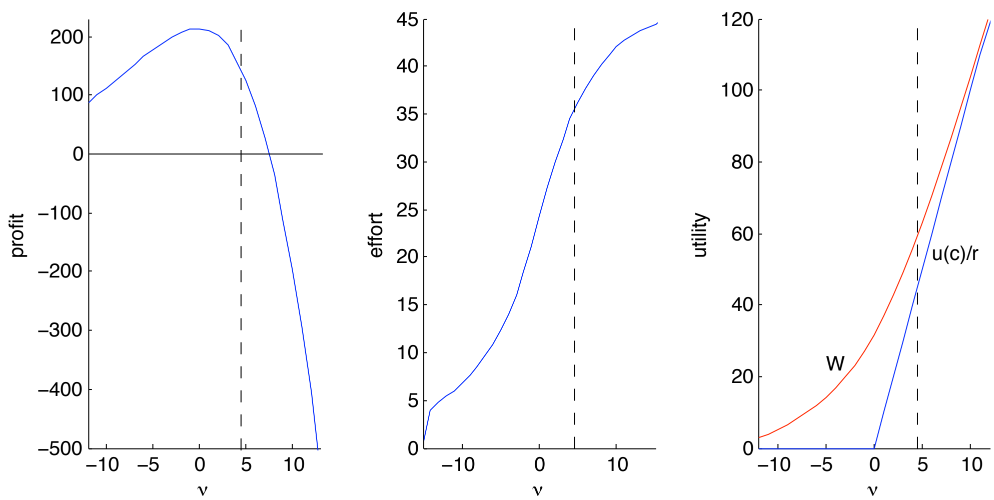

Figure 2: Profit, the agent's effort and utility in the optimal contract.

According to the interpretations above, the relevant question is: what is the marginal benefit of exposing the CEO to risk? Suppose that 20% volatility of the agent's pay (effectively the present value of the agent's future compensation is invested in stock) creates \$10 million in value a year. With $r = 5\%$, and the agent's wealth estimated at $\$5/r = \$100$ million, this corresponds to the risk exposure of \$20 million a year, so $a'(\Phi_t)/(\nu_t\sigma) \approx 0.5$. If the agent has relative risk aversion $\gamma = 2$, then optimal volatility of the agent's pay (or wealth) is $0.5/\gamma = 25\%$.

Let us see how this informal logic pans out on a computed example. Since $a'(\Phi_t)/(\nu_t\sigma) = 1/(\theta\nu_t\sigma) = 0.5$, we have $\theta = 1/(2\sqrt{5} \cdot 2000 \cdot 0.5) = .00022$. In this case, Figure 2 shows the principal's profit and the agent's effort, and compares the agent's continuation value and the utility from consuming c_t in perpetuity for $u(c) = \sqrt{c}$. The agent's effort creates output of \$32 million a year at the pay level of \$5 million ($\nu_t \approx 4.5$). This is less, but on the same order, than the amount predicted by informal logic (which leads to the volatility of ν of 50%, the volatility of the agent's pay of 100%, and output of \$50 million a year). The main reason for the difference is that, due to the agent's limited liability, the contract delivers a lower pay-performance sensitivity than the informal logic implies, see the last panel of Figure 2.

5.2 The Optimal Contract for an Arbitrary κ.

In general, the characterization of the optimal contract in terms of the multipliers ν_t and λ_t on W_t and Φ_t is clearer and more directly interpretable than the characterization in terms of the state variables W_t and Φ_t. Specifically, the solution to the relaxed control problem (which generates the optimal contract when the sufficient condition (31) on Γ_t is satisfied) is characterized by

$$d\nu_t = \lambda_t \frac{r+\kappa}{\sigma}\, dZ_t \qquad \text{and} \tag{37}$$

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

and the agent's compensation is determined by ν_t directly through

$$c_t = \arg\max_c \; \nu_t u(c) - c. \tag{38}$$

The advantages of the Lagrangian characterization in this case are as follows:

- The agent's compensation is determined directly from ν_t by (38), rather than indirectly through a function on the space of W and Φ
- 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,[12] 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: variable Φ_t. This variable determines the agent's effort, the drift of λ_t and expected output μ_t (needed to calculated 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.

It is useful to point out that both characterizations (20) in the large-firm case and (35) in the standard case are special cases of (37). Indeed, (20) is obtained as $\nu_t h'(a(\Phi_t))$ becomes negligible in the limit as $\theta \to 0$. Variable $\hat{\lambda}_t$ is a normalization of λ_t to ensure that the target level of this variable

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

remains invariant in the limit. Ignoring the term $\nu_t h'(a(\Phi_t))$, the target level of λ_t is $a'(\Phi_t)/\kappa = 1/(\theta\kappa)$ converges to 0 as $\theta \to 0$, but the target level of $\hat{\lambda}_t = \lambda_t(r+\kappa)/\sigma$ stays at $(r+\kappa)/(\kappa\psi)$.

In general, we can think of the *target* level of λ_t as the level at which

$$\bar{\lambda}(\nu_t) = \frac{a'(\Phi_t)}{\kappa}(1 - h'(a(\Phi_t))\nu_t), \tag{39}$$

where the drift of λ_t is 0. The rate of convergence of λ_t to its target level depends on κ (as well as on how Φ_t varies with λ_t). As $\kappa \to \infty$, λ_t converges to its target level instantaneously. The volatility of ν_t at the target level of λ_t is

$$\frac{a'(\Phi_t)}{\sigma}(1 - h'(a(\Phi_t))\nu_t)\frac{\kappa + r}{\kappa}. \tag{40}$$

As $\kappa \to \infty$, this leads to the volatility of ν_t given by (35) for any level of ν_t

Before deriving this form of the optimal contract, and justifying it, let me provide several numerical examples.

5.3 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.[13] 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 (39) - 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.

[13] 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.

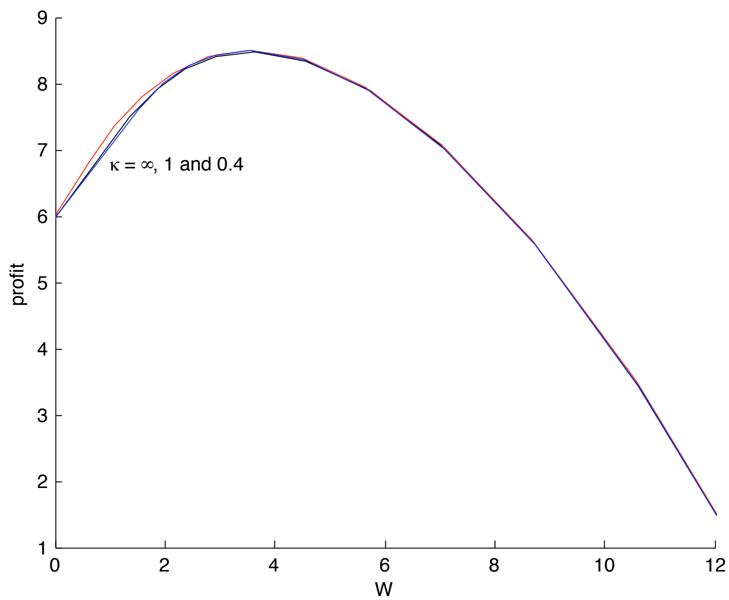

Figure 3: Principal's profit, as a function of W_0.

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 3 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.[14]

The sufficient incentive condition (31) is satisfied in these examples, and it holds more easily when κ is larger. Figure 4 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 5 illustrates the dynamics of the state variables ν_t and λ_t, for $\kappa = 0.4$. For clarity, and to facilitate comparison across different values of κ, 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 (40); and at $\nu_t = 0$ it is $(\kappa + r)/(\kappa\theta)$.

[14]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.

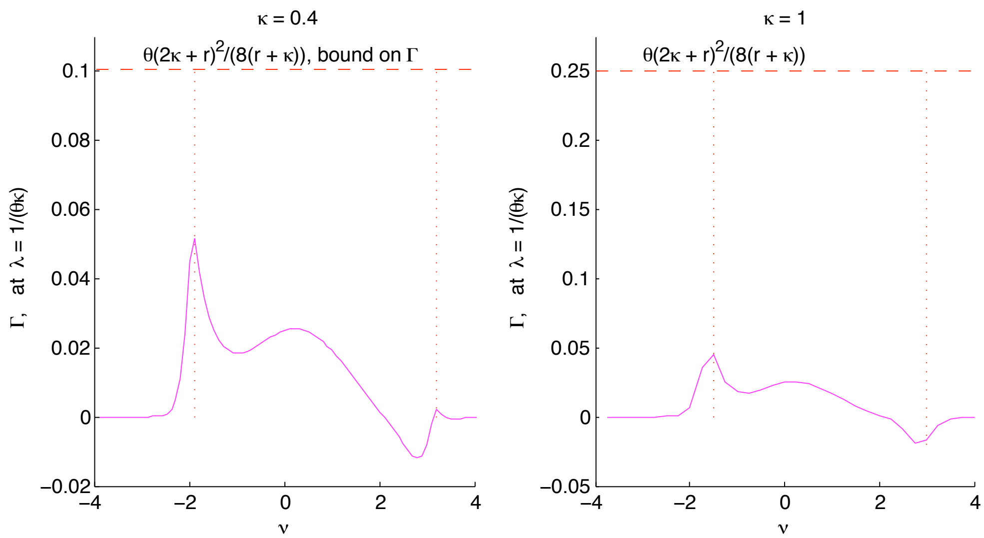

Figure 4: Γ is within bound (termination boundaries are dotted lines).

To describe the dynamics, for several values of λ_t, Figure 5 shows the rate of change of the volatility of ν_t over one quarter. Points where the curves that illustrate the change in the volatility of ν_t intersect the solid horizontal lines indicate the target volatility of ν_t. The boundaries where termination occurs are indicated by dashed lines.

For comparison, Figure 5 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.[15]

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

[15]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.

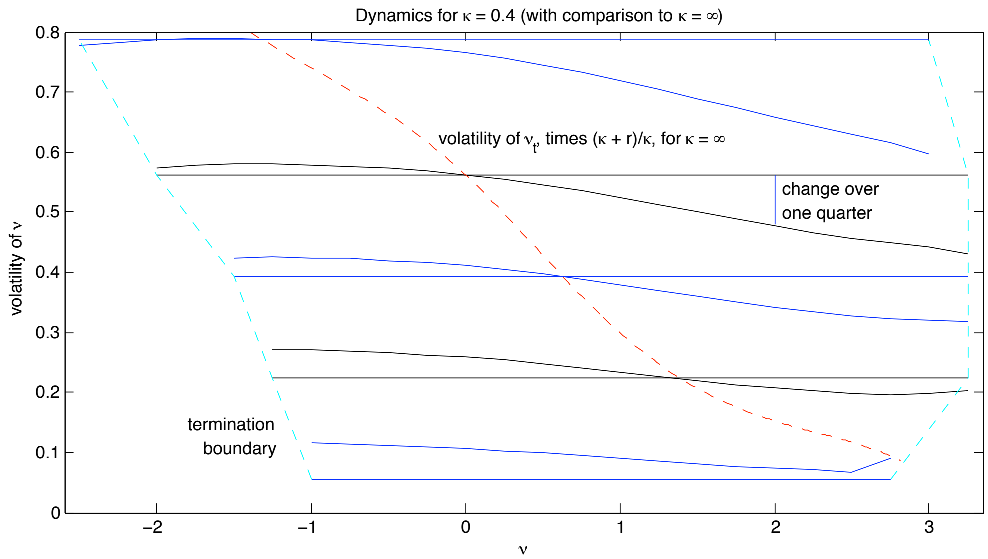

Figure 5: Dynamics of the volatility of ν_t, $\kappa = 0.4$.

to

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

Determine the termination time τ optimally.

Figure 6 compares profit under the approximately optimal contract given 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 6 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.

5.4 Justification.

I justify the characterization (38)-(37), and determine the optimal termination time τ, by backward induction in two steps. First, I derive the form of the optimal contract after time τ. Second, I characterize the optimal contract in the region of employment $\mathcal{R} \subseteq \mathbb{R} \times [0, \infty)$ and solve the optimal stopping problem to determine τ.

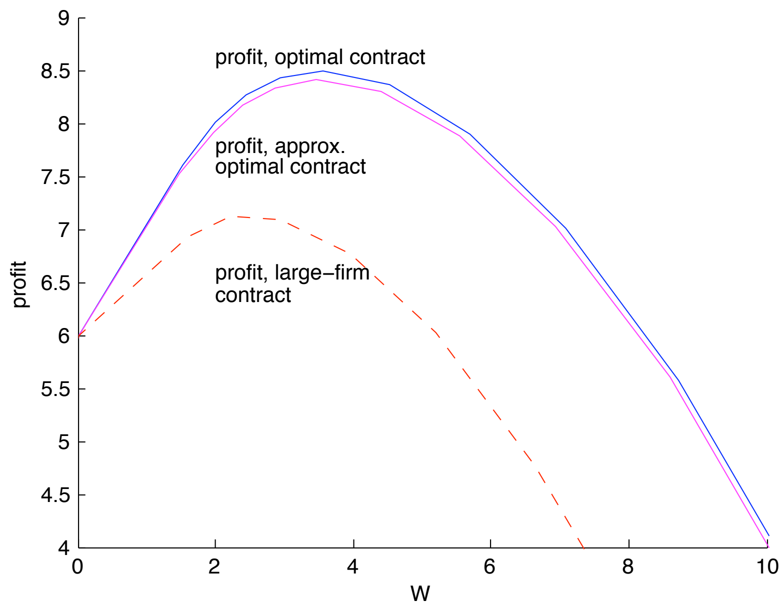

Figure 6: Approximating the optimal contract.

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 marginal benefit of effort 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 we set τ to 0 to simplify notation):[16]

$$\max_c \; E\left[-\int_0^\infty e^{-rt} c_t \, dt\right] \tag{41}$$

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

Problem (41) is easy to solve. Letting ν_0 and λ_0 be the multipliers on the 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

$$c_t = \arg \max_c \; \underbrace{(\nu_0 + \zeta_0^t \lambda_0)}_{\nu_t} u(c) - c, \tag{42}$$

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

$$d\nu_t = \lambda_0 \, d\zeta_0^t = \underbrace{e^{-\kappa t} \lambda_0}_{\lambda_t} \frac{r + \kappa}{\sigma} \frac{dX_t - \mu_t \, dt}{\sigma}, \quad \text{and} \quad d\lambda_t = -\kappa \lambda_t \, dt. \tag{43}$$

This corresponds to the solution (37) after time τ.

Proposition 7 characterizes the principal's profit, as well as the correspondence between the multipliers (ν_0, λ_0) and variables (W_0, Φ_0), in problem (41). These relationships are conveniently represented through a single function $\underline{G}(\nu_0, \lambda_0)$, which solves a tractable parabolic partial differential equation (45).

[16]Interestingly, problem (41) 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

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

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).

Proposition 7 *Define*

$$\underline{G}(\nu_0, \lambda_0) = \max_{\{c_t\}} E\left[\int_0^\infty e^{-rt}\left(\nu_t u(c_t) - c_t\right) dt\right], \tag{44}$$

where (ν_t, λ_t) *follow (43). Then* $\underline{G}$ *solves equation*

$$r\underline{G}(\nu, \lambda) = \max_c \nu u(c) - c - \kappa\lambda\, \underline{G}_\lambda(\nu, \lambda) + \lambda^2 \frac{(r+\kappa)^2}{\sigma^2} \frac{\underline{G}_{\nu\nu}(\nu, \lambda)}{2}. \tag{45}$$

Also, $W_0 = \underline{G}_\nu(\nu_0, \lambda_0)$, $\Phi_0 = \underline{G}_\lambda(\nu_0, \lambda_0)$ *and*

$$E\left[-\int_0^\infty e^{-rt} c_t\, dt\right] = \underline{G}(\nu_0, \lambda_0) - \nu_0 W_0 - \lambda_0 \Phi_0. \tag{46}$$

Proof. Equation (44) is a standard stochastic representation of the solution of the parabolic partial differential equation (45) (see Karatzas and Shreve (1991)). Since $\nu_t = \nu_0 + \zeta_0^t \lambda_0$, differentiating (44) with respect to ν_0 and using the Envelope theorem, we get

$$\underline{G}_\nu(\nu_0, \lambda_0) = E\left[\int_0^\infty e^{-rt} u(c_t)\, dt\right] = W_0. \tag{47}$$

Differentiating with respect to λ_0 we get

$$\underline{G}_\lambda(\nu_0, \lambda_0) = E\left[\int_0^\infty e^{-rt} \zeta_0^t\, u(c_t)\, dt\right] = \Phi_0. \tag{48}$$

Finally, from the stochastic representation (44) itself, we have

$$\underline{G}(\nu_0, \lambda_0) = \nu_0 W_0 + \lambda_0 \Phi_0 - E\left[\int_0^\infty e^{-rt} c_t\, dt\right]. \tag{49}$$

Equation (49) implies (46). ■

Function $\underline{G}$ provides information about the maximal attainable profit for any pair (W_τ, Φ_τ) at the time of termination. It can be used to determine the optimal termination time τ.

The optimal contract before termination. I will show that the equation analogous to (45), which characterizes the optimal contract before termination, is

$$rG = \max_c \ a(G_\lambda) - c + \nu\,(u(c) - h(a(G_\lambda))) - \kappa\lambda G_\lambda + \lambda^2 \frac{(r+\kappa)^2}{\sigma^2}\frac{G_{\nu\nu}}{2}. \quad (50)$$

The following proposition shows that as long as equation (50) has an appropriate solution on a subset $\mathcal{R} \subseteq [0,\infty) \times \mathbb{R}$ of the state space, with smooth-pasting conditions

$$G(\nu,\lambda) = \underline{G}(\nu,\lambda) + L \quad \text{and} \quad \nabla G(\nu,\lambda) = \nabla \underline{G}(\nu,\lambda). \quad (51)$$

on the boundary of $\mathcal{R}$, then the optimal contract is characterized by the laws of motion of the state variables (37). In this contract,

$$W_t = G_\nu(\nu_t,\lambda_t), \quad \Phi_t = G_\lambda(\nu_t,\lambda_t), \quad (52)$$

and τ is the time when the pair (λ_t,ν_t) reaches the boundary of the region $\mathcal{R}$ for the first time.

Proposition 8 *Suppose that function G solves equation (50) on $\mathcal{R} \subseteq [0,\infty)\times \mathbb{R}$ and satisfies the smooth-pasting conditions (51) on the boundary. Then, as long as the transversality conditions hold, W_t and Φ_t are given by (52), and the principal's future profit is $G(\nu_t,\lambda_t) - \nu_t W_t - \lambda_t\Phi_t$, in the contract defined by (37).*

In addition, if on $\mathcal{R}$, $G(\nu,\lambda) \geq \underline{G}(\nu,\lambda)$ and the Hessian of G is positive definite, and outside $\mathcal{R}$,

$$r\underline{G} \geq \max_c \ a(\underline{G}_\lambda) - c + \nu\,(u(c) - h(a(\underline{G}_\lambda))) - \kappa\lambda \underline{G}_\lambda + \lambda^2 \frac{(r+\kappa)^2}{\sigma^2}\frac{\underline{G}_{\nu\nu}}{2} \quad (53)$$

and the Hessian of $\underline{G}$ is positive definite, then the contract is optimal.

Proof. See Appendix. ■

The method of solving stochastic control problems using Lagrange multipliers on state variables is called the *stochastic maximum principle.* Yong and Zhou (1999) describe this approach for a general class of mathematical problems that have finite time horizon. I apply a similar approach to an infinite-horizon problem. However, in general, the laws of motion of the Lagrange multipliers, such as (37), are only necessary first-order conditions for a candidate control policy to be optimal.[17]

I adapt the standard martingale verification argument to demonstrate the optimality of a candidate control policy described by (37). Function $G(\nu, \lambda)$ implies a value function F on the state space given by (W_t, Φ_t), and the properties of F required for the martingale verification argument to go through correspond to the properties of G that are assumed in the second paragraph of Proposition 8. The approach of carrying out a verification argument through the properties of G instead of F is new methodologically, and offers promise for solving other economic applications via the stochastic maximum principle.

6 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 *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

[17]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.

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.

Proposition 9 *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

$$r\underline{F}^\tau(\nu, t) = \max_c\, (\nu u(c) - c) + \underline{F}_2^\tau(\nu, t) + \frac{\hat{\lambda}(\tau, t)^2}{2} \underline{F}_{11}^\tau(\nu, t). \tag{54}$$

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

Proof. Equation (54) 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 problem (22). ■

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

$$d\bar{\Phi}_t = e^{-(r+\kappa)t}\,\Gamma_t \sigma\, dZ_t. \tag{55}$$

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 (55), 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). ■

Proof of Proposition 6. 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

$$\hat{W}_t(\hat{\mu}_t) = W_t + \frac{\Phi_t}{r+\kappa}(\hat{\mu}_t - \mu_t) + L(\hat{\mu}_t - \mu_t)^2, \tag{56}$$

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 (56) 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 (56) 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

$$\begin{bmatrix} -\theta/2 & L(r+\kappa) \\ L(r+\kappa) & \frac{\Gamma_t}{r+\kappa} - (r+2\kappa)L \end{bmatrix} \tag{57}$$

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 (57) 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

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

■

The following Lemma shows that if the condition of Proposition 6 is violated for all t, then there is in fact a strategy that attains infinite utility for the agent.

Lemma 1 *Suppose that the agent's cost of effort is quadratic of the form $\theta a^2/2$, $\theta \in \mathbb{R}$. Consider any contract in which $\tau = \infty$, $\Gamma_t \geq \frac{\theta(2\kappa+r)^2}{8(r+\kappa)} + \epsilon_\Gamma$ for all t, and the laws of motions of state variables satisfy the growth conditions*

$$rW_t - u_t \leq (r - \epsilon_W)W_t + K_W, \quad \Delta_t \leq L_\Delta W_t^{1/2} + K_\Delta,$$

$$(r+\kappa)(\Phi_t - \Delta_t) \leq \frac{r}{2}\Phi_t + K_\Phi, \quad \Gamma_t \leq K_\Gamma$$

for some positive constants ϵ_W, K_W, L_Δ, K_Δ, K_Φ and K_Γ. Then, the agent has a strategy that attains infinite utility.

Proof. Suppose that $\hat{\mu}_t \neq \mu_t$, and consider the strategy that sets $(r+\kappa)(\hat{a}_t - a_t) = (\kappa + r/2 - \epsilon)(\hat{\mu}_t - \mu_t)$, where $\hat{\mu}_t$ will be specified later. Then, under this strategy

$$d(\hat{\mu}_t - \mu_t) = (r/2 - \epsilon)(\hat{\mu}_t - \mu_t)$$

and the drift of $\hat{V}_t$ is e^{-rt} times

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

Setting $L = \frac{\theta(r+2\kappa)}{4(r+\kappa)^2}$, this expression becomes greater than or equal to $(\hat{\mu}_t - \mu_t)^2$ times

$$-\frac{\theta}{2}\frac{(\kappa + r/2 - \epsilon)^2}{(r+\kappa)^2} + \frac{\epsilon_\Gamma}{r+\kappa} - \frac{\theta(r+2\kappa)^2}{8(r+\kappa)^2} + \frac{\theta(r+2\kappa)(\kappa + r/2 - \epsilon)}{2(r+\kappa)^2}.$$

If $\epsilon = 0$ then this is $\frac{\epsilon_\Gamma}{r+\kappa}$. Let us choose ϵ slightly above 0, so that this expression is still greater than $\frac{\epsilon_\Gamma}{2(r+\kappa)}$, and the drift of $\hat{V}_t$ is greater than or equal to

$$e^{-rt}\frac{\epsilon_\Gamma}{2(r+\kappa)}e^{(r-2\epsilon)t}(\hat{\mu}_0 - \mu_0) = \frac{\epsilon_\Gamma}{2(r+\kappa)}e^{-2\epsilon t}(\hat{\mu}_0 - \mu_0).$$

Then

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

$$\hat{V}_0 + (\hat{\mu}_0 - \mu_0)\int_0^t \frac{\epsilon_\Gamma}{2(r+\kappa)}e^{-2\epsilon s}\,dt \approx \hat{V}_0 + (\hat{\mu}_0 - \mu_0)\frac{\epsilon_\Gamma}{4\epsilon(r+\kappa)}$$

when t is large. To show that the agent's utility becomes infinite as $\epsilon \to 0$, we need to prove the transversality condition that $E[e^{-rt}\hat{W}_t(\hat{\mu}_t)] \to 0$ as $t \to \infty$.

Note that $e^{-rt}L(\hat{\mu}_t - \mu_t)^2$ converges to 0, because $(\hat{\mu}_t - \mu_t)^2$ grows at rate $r - 2\epsilon$, less than r. Furthermore, Φ_t grows at rate less than $r/2 + \epsilon'$ for any arbitrarily small ϵ', since

$$(r+\kappa)(\Phi_t - \Delta_t) \leq \frac{r}{2}\Phi_t + K_\Phi \quad \text{and} \quad \Gamma_t(\hat{\mu}_t - \mu_t) = O(e^{r/2-\epsilon}) = o(e^{r/2+\epsilon'}).$$

Thus, $\Phi_t(\hat{\mu}_t - \mu_t)$ grows at rate less than $r - \epsilon/2$. Finally, assume without loss of generality that $\epsilon_W/2 < \epsilon$. Then W_t grows at rate less than $r - \epsilon_W/2$, because

$rW_t - u_t \leq (r-\epsilon_W)W_t + K_W$ and then $\Delta_t(\hat{\mu}_t - \mu_t)$ is $O(e^{r-\epsilon_W/4-\epsilon}) = o(e^{r-\epsilon_W/2})$. It follows that the transversality conditions hold. ■

It is easy to see that there is a nonempty set of contracts that satisfy the conditions of Lemma 1. Indeed, let

$$dx_t = \gamma(dX_t - \mu_t\, dt), \quad W_t = x_t^2 + K,$$

where $K \geq \gamma^2\sigma^2/r$. Then the drift of W_t is $\gamma^2\sigma^2 = rW_t - u_t$, so $u_t = rW_t - \gamma^2\sigma^2 \geq 0$, i.e. c_t is well-defined. The agent's problem is linear-quadratic. We have $\Phi_t = \Delta_t = 2\gamma x_t$ and $\Gamma_t = 2\gamma^2$, and the growth conditions of Lemma 1 are satisfied. If

$$2\gamma^2 > \frac{\theta(2\kappa + r)^2}{8(r+\kappa)},$$

then by Lemma 1 the agent can get negative utility.

Proof of Proposition 8. First, let us show that if G solves (50) on $\mathcal{R} \subseteq [0,\infty)\times\mathbb{R}$ and satisfies the smooth-pasting conditions (51) 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 (37).

Differentiating (50) with respect to ν and using the Envelope Theorem, we get

$$rG_\nu - u(c) + h(a(G_\lambda)) =$$
$$(a'(G_\lambda)\,(1-\nu h'(a(G_\lambda))) - \kappa\lambda)\, G_{\nu\lambda} + \lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{G_{\nu\nu\nu}}{2}. \tag{58}$$

The right hand side represents the drift of the process $G_\nu(\nu_t,\lambda_t)$ when (ν_t,λ_t) follow (37). 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 (50) with respect to λ and using the Envelope Theorem, we get

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

$$\underbrace{(a'(G_\lambda)(1-\nu h'(a(G_\lambda)))-\kappa\lambda)G_{\lambda\lambda}+\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\frac{G_{\lambda\nu\nu}}{2}}_{\text{drift of } G_\lambda}. \tag{59}$$

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 (58) and λ times (59) from (50), we get

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

$$+\frac{1}{2}\lambda^2\frac{(r+\kappa)^2}{\sigma^2}\underbrace{(-G_{\nu\nu}-\nu G_{\nu\nu\nu}-\lambda G_{\lambda\nu\nu})}_{\frac{\partial^2(G-\nu G_\nu-\lambda G_\lambda)}{\partial\nu^2}}. \tag{60}$$

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 (37).

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 (see Theorem ??). Define $G(\nu,\lambda) = \underline{G}(\nu,\lambda)$ outside $\mathcal{R}$. If the Hessian of G is positive definite, then there exist processes*

$$d\nu_t = \mu_t^\nu\, dt + \sigma_t^\nu\, dZ_t \quad \text{and} \quad d\lambda_t = \mu_t^\lambda\, dt + \sigma_t^\lambda\, dZ_t \tag{61}$$

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

$$\underbrace{\begin{bmatrix} G_{\nu\nu} & G_{\nu\lambda} \\ G_{\lambda\nu} & G_{\lambda\lambda} \end{bmatrix}}_{H(G)} \begin{bmatrix} \sigma_t^\nu \\ \sigma_t^\lambda \end{bmatrix} = \begin{bmatrix} \Delta_t\sigma \\ \Gamma_t\sigma \end{bmatrix}. \tag{62}$$

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 "…" 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 (61).

Using Ito's lemma and the laws of motion of W_t and Φ_t, the drift of $G(\lambda_t,\nu_t) - \nu_t W_t - \lambda_t \Phi_t$ 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 (62). 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 (37) then the drift of $G(\nu_t, \lambda_t) - \nu_t W_t - \lambda_t \Phi_t$ 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 (60), leads to a drift of $\bar{F}_t$ of zero in $\mathcal{R}$ and negative outside $\mathcal{R}$, by (53).

Now, when λ_t and ν_t follow (61) instead of (37) the drift of $\bar{F}_t$ changes by e^{-rt} times

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

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

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 (37), 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.[18] Therefore, the contract, in which λ_t and ν_t follow (37), is optimal. ■

Appendix B: Numerical Procedures.

To generate numerical examples, I assume that the agent is able to work only until time T (so that $\tau \leq T$) and let T go to infinity. Then it is possible to find the optimal contract by solving for relevant functions of T.

[18] 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.

To obtain the boundary conditions at $T = 0$ and at termination, I solve the following system of simple parabolic equations, which rely on the law of motion of λ given by $d\lambda_t = -\kappa\lambda_t\, dt$:

$$r\underline{F}(\nu,\lambda) = -c - \kappa\lambda\underline{F}_\lambda(\nu,\lambda) + \frac{1}{2}\frac{\lambda^2(r+\kappa)^2}{\sigma^2}\underline{F}_{\nu\nu}(\nu,\lambda),$$

$$r\underline{W}(\nu,\lambda) = u(c) - \kappa\lambda\underline{W}_\lambda(\nu,\lambda) + \frac{1}{2}\frac{\lambda^2(r+\kappa)^2}{\sigma^2}\underline{W}_{\nu\nu}(\nu,\lambda) \quad \text{and}$$

$$(r+\kappa)\underline{\Phi}(\nu,\lambda) = \underbrace{\lambda\frac{(r+\kappa)^2}{\sigma^2}\underline{W}_\nu(\lambda,\nu)}_{(r+\kappa)\Delta} - \kappa\lambda\underline{\Phi}_\lambda(\nu,\lambda) + \frac{1}{2}\frac{\lambda^2(r+\kappa)^2}{\sigma^2}\underline{\Phi}_{\nu\nu}(\nu,\lambda).$$

For $T > 0$, I solve

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

$$W_T + rW = u(c) - h(a(\Phi)) + (a'(\Phi)(1-\nu\Phi) - \kappa\lambda)W_\lambda + \frac{1}{2}\frac{\lambda^2(r+\kappa)^2}{\sigma^2}W_{\nu\nu} \quad \text{and}$$

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

using $F(\nu,\lambda,0) = \underline{F}(\nu,\lambda) + L$, $W(\nu,\lambda,0) = \underline{W}(\nu,\lambda)$ and $\Phi(\nu,\lambda,0) = \underline{\Phi}(\nu,\lambda)$ as initial conditions. Moreover, termination is triggered, and functions $F(\nu,\lambda,T)$, W and Φ are replaced with $\underline{F}(\nu,\lambda)$, $\underline{W}$ and $\underline{\Phi}$, whenever

$$F(\nu,\lambda,T) + W(\nu,\lambda,T)\nu + \Phi(\nu,\lambda,T)\lambda \ \leq \ \underline{F}(\nu,\lambda) + \underline{W}(\nu,\lambda)\nu + \underline{\Phi}(\nu,\lambda)\lambda.$$

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.

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.

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.