Learning in Dynamic Incentive Contracts

PETER M. DEMARZO AND YULIY SANNIKOV*

May 28, 2008

Abstract

We derive the optimal dynamic contract in a continuous-time principal-agent setting, in which both investors and the agent learn about the firm's profitability over time. Because investors learn about the firm's future profitability from current output, which also depends upon the agent's actions, deviations by the agent distort investors' beliefs. We characterize the optimal contract, and show that the performance sensitivity of the agent's payoff is set to account for both moral hazard and asymmetric information. We show that the optimal contract can be implemented by compensating the agent with equity and allowing him to manage the firm's cash reserves by setting the its payout policy. Under this contract, the firm accumulates cash until it reaches a target balance that depends on the agent's perceived productivity. Once this target balance is reached, the firm starts paying dividends equal to its expected future earnings, while any temporary shocks to earnings either add to or deplete the firm's cash reserves. The firm is liquidated if it exhausts its cash reserves. We also show that once the firm initiates dividends, these dividends are smooth relative to earnings, and that liquidation is first-best, despite the agency problem.

* Stanford University and UC Berkeley/NYU. This research is based on work supported in part by the NBER and the National Science Foundation under grant No. 0452686.

In many types of agency relationships output is informative about both the agent's current actions and the future quality of the project itself. Company earnings depend on both current managerial effort and unobservable exogenous factors that affect the firm's future profitability. Mortgage payments reflect the borrower's financial discipline and her ability to pay, which is influenced by such factors as income and nondiscretionary expenses. The performance of a worker may depend on both the worker's effort and skill. In all of these settings, the dynamic agency problem is complicated by the fact that both the agent and the principal are learning about the quality of the project over time, and that the agent's actions can distort the principal's inferences.

In a dynamic context, aligning incentives in such an environment is complicated by the fact that the agent's knowledge of his own actions will provide private information regarding the project's quality. Absent appropriate long-term contracts, incentives may even break down completely due to the "ratchet effect": When the principal has full bargaining power but cannot commit to a long-term contract, the principal will measure the agent's performance against a tougher benchmark if the agent reveals himself to be a more productive type. As shown by Laffont and Tirole (1988), in anticipation of the ratchet effect the agent will not want reveal his productivity when it is high, and so will offset it by reducing effort.[1]

In contrast, the literature on career concerns (e.g. Holmstrom (1999)) finds that if all bargaining power lies in the hands of the agent, whose wage is set competitively each period, then the agent's concern over his future wage will create incentives for effort even if employers only offer flat wage contracts. The reason is that workers with a better output history get higher wage offers because they are perceived to be more productive. Though incentives to both exert effort and reveal high ability exist in this setting, they may be too strong or too weak to achieve an optimal outcome depending on agent's ability to manipulate market beliefs through effort.

These incentive properties of short-term contracts certainly motivate the study of dynamic contracts under full commitment. Unfortunately, characterizing the optimal dynamic agency contract with learning has proven difficult because of technical

[1] See also Baron and Besanko (1987), and Freixas, Guesnerie, and Tirole (1985), for related results.

challenges. Standard dynamic agency models assume that the principal and the agent always have common knowledge about how the distribution of future output depends on the agent's effort.[2] This structure facilitates the application of recursive methods to solve for an optimal contract.[3] This is not the case when the principal is learning about the project's quality, as then the principal and agent may have asymmetric information about future output. Asymmetric information arises endogenously, as the agent's actions can distort the principal's inferences about future profitability. For example, suppose the agent lowers effort, causing a temporary drop in output. If the principal anticipates high effort by the agent, he attributes a drop in output to a negative exogenous shock, and so has lower expectations about future output than the agent. This effect provides an additional channel through which the agent's current actions may alter the terms of the contract going forward, and thereby affect the agent's incentives.

We construct a model of dynamic agency with learning that is both (1) tractable under full commitment and (2) adaptable to study a number of applications. The main ingredients of our model are as follows. Time is continuous. The agent is risk neutral but has limited liability. The agent's output is stochastic, with an expected flow on date t given by δ_t if the agent puts in full effort. The agent may also take actions that generate private benefits, but at the cost of reducing the expected output flow below δ_t. Importantly, the expected flow of benefits from the relationship, δ_t, is not directly observable and evolves over time stochastically. The source of the asymmetric information problem is that the principal and the agent learn about δ_t from realized output. Thus, the principal's belief about δ_t becomes distorted if the agent deviates. A contract provides the agent with incentives by specifying payments to the agent and conditions under which the relationship with the principal is terminated, both as a function of the history of observed output.

---

[2] See, e.g., Rogerson (1985), Thomas and Worrall (1990), Phelan and Townsend (1991) or more recently Albuquerque and Hopenhayn (2001), DeMarzo and Fishman (2007), DeMarzo and Sannikov (2006), Biais, Mariotti, Plantin and Rochet (2007) and Biais, Mariotti, Rochet and Villeneuve (2007).

[3] The use of recursive methods to solve for dynamic contracts was pioneered by Spear and Srivastava (1987). See also Stokey and Lucas (1989) and Ljungqvist and Sargent (2004).

For the sake of concreteness, our paper studies one application in which the agent is a manager of a firm, who can put effort in the firm and engage in outside activities to produce a private benefit. In this case output corresponds to the firm's profits and the hidden parameter δ_t reflects the firm's profitability. In general, however, the model can be applied to many other settings. For example, in the context of regulation and the procurement of public goods, "output" can be interpreted as the social benefit net of the contractor's realized costs, which depends on the contractor's effort, while δ_t represents the contractor's efficiency. In a labor contract setting (see e.g. Lazear (1986) and Holmstrom (1999)), we interpret δ_t as the agent's productivity in producing some measured output. Finally, we can apply or model in the context of mortgages as in Piskorski and Tchistyi (2007), where δ_t captures the borrower's expected ability to pay (i.e. income net of spending needs) and "output" corresponds to actual payments.

Our dynamic agency model delivers general conclusions about (1) the structure of incentive provision in the optimal contract and (2) inefficient termination that may be required to provide optimal incentives. Regarding incentives, the key idea is to understand how the agent's actions affect the principal's beliefs, which determine the benchmark against which the principal measures the agent's future performance. While the agent may be rewarded for current output, he may be reluctant to apply effort if doing so raises expectations regarding future output, and thereby makes it more difficult for the agent to earn future bonuses. To ensure the agent's incentive to exert effort, the optimal contract must compensate the agent for the effect of effort on current output as well as on investors' beliefs about future output. In that sense, the asymmetric information problem raises the optimal performance sensitivity of the contract.[4]

For our specific application, in which the agent is a manager who runs the firm, we show that a simple way to create appropriate incentives is by giving the agent a fraction of the firm's equity. The idea that equity creates incentives is not new, but the way that equity solves the incentive problem in our setting is new. The value of equity changes when the

---

[4] The performance sensitivity we are alluding to is with regard to the agent's utility. It may not be directly observed in the agent's immediate compensation or consumption. Indeed, in the optimal contract there are regions for which the agent's immediate consumption is insensitive to the firm's current performance, but an increase in utility still follows due to an increase in expected future bonuses. Thus, one must be careful relating these results to empirical measures of "pay for performance."

output realization, which measures the agent's performance, differs from expected output, which is a benchmark against which performance is measured. Why does the agent not want to ease the benchmark with lower effort, in order to create a better impression in the future? The reason is that by lowering the expectation about future output, the agent also lowers the value of his equity stake.

After solving for the optimal contract, we show that it can be implemented through a natural specification of the firm's payout policy and a capital structure. When the firm is young, it makes no payouts and accumulates cash until it reaches a target level of financial slack that is positively related to the manager's perceived productivity. Once this target balance is reached, the firm initiates dividend payments. From that point on, the firm pays dividends at a rate equal to its *expected* future earnings. Specifically, the firm smoothes its dividends by absorbing any temporary shocks to earnings through an increase or decrease of its cash reserves (or available credit). When the firm exhausts its available cash (or credit), the firm is liquidated and agency relationship is terminated.

This payout policy captures well many of the stylized facts associated with observed payout policies. Immature firms do not pay dividends, but instead retain their earnings to invest, repay debt, and build cash reserves. For these firms the value of internal funds is high, as they risk running out of cash and being prematurely liquidated. But once the firm has sufficient financial slack, dividends are then paid at a level that appears to be a smooth estimate of the permanent component of earnings.[5] Because dividend changes reflect permanent changes to profitability, they are persistent and have substantial implications for firm value.

The payout policy we identify in our model is incentive compatible for the manager and is the unique implementation of the optimal mechanism that provides the fastest possible payout rate subject to the constraint that the firm will not need to raise external capital in the future.[6] Thus, in our model firms build up a target level of internal funds to ensure that they never need to liquidate inefficiently due to financial constraints. This constraint

---

[5] See Lintner (1956), Fama and Babiak (1968), as well as more recent studies by Allen and Michaely (2003) and Brav et. al. (2005).
[6] There are many alternative dividend policies that are optimal if there is no cost to raising equity capital in the event that the firm runs out of cash in the future. In practice, however, these costs are substantial.

alone cannot explain dividend smoothing, as once the target is reached we would expect all excess cash flows to be paid out as dividends. The key driver of dividend smoothing in our model is the fact that there is learning about the firm's profitability based on the current level of the firm's cash flows. When cash flows are high, the firm's perceived profitability increases. This raises the cost of liquidating the firm (we are liquidating a more profitable enterprise), and therefore raises the optimal level of financial slack. Thus, a portion of the firm's high current cash flow will optimally be used to increase its cash reserves, resulting in a smoothed dividend policy.[7]

In the next section of the paper, we describe the continuous-time principal-agent problem with learning about the firm's profitability. Then in Section 2, we present the solution to this problem when moral hazard is absent, which is based purely on option-value considerations. This solution is important to understand the optimal long-term contract with moral hazard, which we derive in Section 3. In Section 3, we also show how this optimal contract can be implemented in terms of the firm's payout policy. Section 4 presents a theoretical justification of the optimal contract, and Section 5 discusses several extensions of our basic model.

1. Model

Risk-neutral outside investors hire a risk-neutral manager to run a firm. Investors have unlimited wealth, whereas the agent has no initial wealth and must consume non-negatively.[8] Both the agent and investors discount the future at rate $r > 0$.

The profitability of the firm depends on the agent's managerial skill δ_t, which evolves stochastically with time t. The firm provides resources that allow the agent to use his skills more productively. That is, inside the firm the agent can use his skills to produce cash flows at an expected rate of δ_t, whereas outside the firm, his productivity is only $\lambda\delta_t$

[7] This payout policy is therefore also consistent with the evidence that firms' cash and leverage positions are strongly influenced by past profitability, even when firms are financially unconstrained (see, e.g., Fama and French (2002)).
[8] The assumption that the agent has no initial wealth is without loss of generality; equivalently, we can assume the agent has already invested any initial wealth in the firm. The agent's limited liability prevents a general solution to the moral hazard problem in which the firm is simply sold to the agent.

for some $\lambda \in (0, 1)$. If the agent is fired, his expected outside option equals the perpetuity value of his outside productivity, and so is given by $R(\delta_t) = \lambda\delta_t/r$.

When the agent works for the firm, there is moral hazard. The firm's cumulative cash flows X_t are

$$dX_t = (\delta_t - a_t)dt + \sigma dZ_t, \tag{1}$$

where σ is the volatility of cash flows, Z is a standard Brownian motion, and $a_t \geq 0$ is the extent to which the agent diverts his own effort and firm resources from the firm for private benefit. It is more efficient to use the agent's and firm's resources inside the firm to generate profit rather than for the agent's private benefit. If $a_t > 0$, the agent gets a private benefit at rate of λa_t.

The profitability of the firm δ_t, which depends on the agent's managerial skill, evolves over time.[9] The firm's cash flows convey information about future profitability. We model this by assuming that δ_t starts at $\delta_0 > 0$ and evolves according to

$$d\delta_t = \nu(dX_t - (\delta_t - a_t)dt) = \nu\sigma\, dZ_t, \tag{2}$$

as long as $\delta_t > 0$.

Equation (2) can be interpreted as the steady state of a filtering problem in which the firm's true profitability, δ_t^*, is unobservable, and the manager and investors attempt to learn the firm's profitability based on past cash flows. Specifically, suppose δ_0^* is initially normally distributed with standard deviation γ_0, and is subject to normally distributed shocks so that

$$d\delta_t^* = \sigma' dZ_t'$$

where Z' is a Brownian motion independent of Z. Then we can apply the Kalman-Bucy filter to compute the Bayesian posterior distribution for δ_t^* given the observed cash flows up to date t. In particular, if we define

$$\delta_t \equiv E_t[\ \delta_t^*\ ]$$

[9] Alternatively, we could interpret δ_t as the suitability of the agent's managerial style and skills to current market conditions.

then[10]

$$d\delta_t = \gamma_t/\sigma^2\,(dX_t - (\delta_t - a_t))\,dt \tag{3}$$

where γ_t is the standard deviation of the posterior distribution and evolves according to

$$d\gamma_t = ((\sigma')^2 - \gamma_t^2/\sigma^2)\,dt.$$

For simplicity, for most of the paper we focus the steady state of this problem, in which $\gamma_0 = \gamma_t = \sigma\sigma'$; in that case, (3) reduces to (2) with $\nu = \sigma'/\sigma$. In Section 5 we analyze an extension in which $\gamma_0 > \sigma\sigma'$ and γ_t is decreasing, so there is more uncertainty about firm's future profitability when it is young.

Since the agent's unobservable effort enters the learning equation (2), in our contractual environment the agent may have private information not only about his effort, but also about the firm's profitability. Indeed, if the principal expects the agent to choose effort a_t, the principal will update his belief $\hat{\delta}_t$ about firm profitability according to

$$d\hat{\delta}_t = \nu(dX_t - (\hat{\delta}_t - a_t)dt), \quad \hat{\delta}_0 = \delta_0. \tag{4}$$

Thus, if the agent chooses a different effort strategy $\hat{a}_t \neq a_t$, the principal's belief $\hat{\delta}_t$ will be incorrect. Unlike in a standard principal-agent settings, in our environment deviations by the agent generate asymmetric information between him and the principal.

The firm requires external capital of $K \geq 0$ to be started. The investors contribute this capital and in exchange receive the cash flows generated by the firm less any compensation paid to the agent. The agent's compensation is determined by a long-term contract. This contract specifies, based on the history of the firm's cash flows, non-negative compensation dC for the agent while he manages the firm, as well as a time τ when the agent is fired. Formally, a contract is a pair (C, τ), where C is a non-decreasing X-measurable process that represents the agent's cumulative compensation and τ is an X-measurable stopping time. When the agent leaves the firm, he receives his outside payoff of $\lambda\delta_t/r$, and the investors receive a payoff of

[10] Note that in this interpretation, the firm's cash flow process is given by $dX_t = (\delta_t^* - a_t)\,dt + \sigma\,dZ_t$, and so (2) represents the firm's expected cash flows after integrating over the posterior for δ_t^*.

$$L(\delta_\tau) = L_0 + \frac{\kappa\delta_\tau}{r}, \tag{5}$$

where $\lambda + \kappa < 1$.[11] We allow the value of the firm after the manager leaves to depend on δ_t to recognize that investors may be able to capture some fraction of the firm's profitably, e.g. by hiring a new manager.

A contract (C,τ) together with an X-measurable effort recommendation a is optimal given an expected payoff of W_0 for the agent if it maximizes the principal's profit

$$E\left[\int_0^\tau e^{-rt}((\delta_t - a_t)dt - dC_t) + e^{-r\tau}L(\delta_\tau)\right] \tag{6}$$

subject to

$$W_0 = E\left[\int_0^\tau e^{-rt}(\lambda a_t dt + dC_t) + e^{-r\tau}R(\delta_\tau)\ \right] \text{ given strategy } a \tag{7}$$

and

$$W_0 \geq E\left[\int_0^\tau e^{-rt}(\lambda \hat{a}_t dt + dC_t) + e^{-r\tau}R(\delta_\tau)\ \right] \text{ for any other strategy } \hat{a} \tag{8}$$

By varying $W_0 > R(\delta_0)$, we can use this solution to consider different divisions of bargaining power between the agent and the investors. For example, if the agent enjoys all the bargaining power due to competition between investors, then the agent will receive the maximal value of W_0 subject to the constraint that the investors' payoff be at least equal to their initial investment, K. We say that the effort recommendation a is *incentive-compatible* with respect to the contract (C,τ) if it satisfies (7) and (8) for some W_0.

REMARKS. For simplicity, we specify the contract assuming that the agent's compensation and the termination time τ are determined by the cash flow process, ruling out public randomization. This assumption is without loss of generality, as we will later verify that public randomization would not improve the contract.

[11] In principle, the value of the firm after the agent leaves may depend both on both the agent's and investors' beliefs about firm profitability. However, because those beliefs coincide on the equilibrium path, without loss of generality we may specify L as a function of a single variable δ_t.

2. The First-Best Solution.

Before solving for the optimal contract, we derive the first-best solution as a benchmark. In the first-best, the principal can control the agent's effort, and so we can ignore the incentive constraints (8). Then it is optimal to let the agent take action $a_t = 0$ until liquidation, since it is cheaper to provide the agent with a flow of utility by paying him than by letting him divert attention to private activities. Then the total cost of providing the agent with a payoff of W_0 is

$$E\left[\int_0^\tau e^{-rt} dC_t\right] = W_0 - E\left[e^{-r\tau} R(\delta_\tau)\right],$$

and the principal's payoff is

$$E\left[\int_0^\tau e^{-rt}\delta_t\, dt + e^{-r\tau}(L(\delta_\tau) + R(\delta_\tau))\right] - W_0 .$$

Thus, without moral hazard the principal chooses a stopping time τ that solves

$$\bar{b}(\delta_0) = \max_\tau \; E\left[\int_0^\tau e^{-rt}\delta_t\, dt + e^{-r\tau}(L(\delta_\tau) + R(\delta_\tau))\right]. \tag{9}$$

This is a standard real-option problem that can be solved by the methods of Dixit and Pindyck (1994). Because liquidation is irreversible, it is optimal to trigger liquidation when the expected profitability δ_t reaches a critical level of $\underline{\delta}$ that is below the level δ^* such that

$$R(\delta^*) + L(\delta^*) = \delta^*/r$$

See Figure 1.

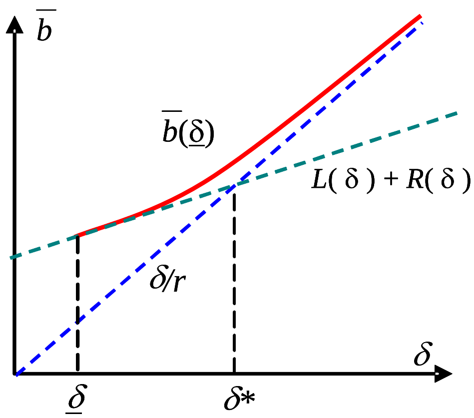

Figure 1: First-best Liquidation Threshold and Value Function

We have the following explicit solution for the first-best liquidation threshold and value of the firm:

Proposition 1. *Under the first-best contract, the firm is liquidated if $\delta \le \underline{\delta}$ where*

$$\underline{\delta} = \delta^* - \frac{\nu\sigma}{\sqrt{2r}}..$$

The principal's payoff is $\bar{b}(\delta_0) - W_0$, *where*

$$\bar{b}(\delta) = \frac{\delta}{r} + \left(L(\underline{\delta}) + R(\underline{\delta}) - \frac{\underline{\delta}}{r} \right) \exp\left(-\frac{\sqrt{2r}}{\nu\sigma} (\delta - \underline{\delta}) \right)$$

if $\delta \ge \underline{\delta}$, *and* $\bar{b}(\delta) = L(\delta) + R(\delta)$ *otherwise.*

Proof: Note that $\bar{b}(\delta)$ is the solution on $[\underline{\delta}, \infty)$ to the ordinary differential equation

$$r\bar{b}(\delta) = \delta + \frac{1}{2}\nu^2\sigma^2\bar{b}''(\delta) \tag{10}$$

with boundary conditions

(a) $\bar{b}(\underline{\delta}) = L(\underline{\delta}) + R(\underline{\delta})$,

(b) $\bar{b}'(\underline{\delta}) = \dfrac{\lambda + \kappa}{r}$ (smooth-pasting),

(c) and $\bar{b}(\delta) - \delta / r \to 0$ as $\delta \to \infty$.

Let us show that $\bar{b}(\delta)$ gives the maximal profit attainable by the principal. For an arbitrary contract (C, τ), consider the process

$$G_t = e^{-rt}\bar{b}(\delta_t) + \int_0^t e^{-rs}\delta_s ds.$$

Let us show that G_t is a submartingale. Using Ito's lemma, the drift of G_t is

$$-re^{-rt}\bar{b}(\delta_t) + e^{-rt}\frac{1}{2}\nu^2\sigma^2\bar{b}''(\delta_t) + e^{-rt}\delta_t,$$

which is equal to 0 when $\delta_t > \underline{\delta}$ and $-re^{-rt}(L(\delta_t) + R(\delta_t)) + e^{-rt}\delta_t < 0$ when $\delta_t < \underline{\delta}$. Therefore, the principal's expected profit at time 0 is

$$\begin{aligned} E\left[e^{-r\tau}L(\delta_\tau) + \int_0^\tau e^{-rt}(\delta_t dt - dC_t) \right] &= E\left[e^{-r\tau}(L(\delta_\tau) + R(\delta_\tau)) + \int_0^\tau e^{-rt}\delta_t dt - W_0 \right] \\ &\le E\left[G_\tau\right] - W_0 \le G_0 - W_0 \\ &= \bar{b}(\delta_0) - W_0. \end{aligned}$$

The inequalities above become equalities if and only if $\delta_\tau \le \underline{\delta}$ and $\delta_t > \underline{\delta}$ before time τ. ■

Our characterization of the first-best contract can be interpreted in terms of the firm's capacity to sustain operating losses. At any moment of time, the firm must be able to withstand a productivity shock of up to $d\delta = -(\delta - \underline{\delta})$. From (2), this corresponds to a cash flow shock equal to

$$dX_t - \delta_t dt = \sigma dZ_t = \frac{d\delta_t}{\nu} = -\left(\frac{\delta - \underline{\delta}}{\nu}\right). \tag{11}$$

We can view Equation (11) as specify the minimal level of "financial slack" the firm will require in order to avoid inefficient liquidation. This result will play an important role in our implementation of the optimal contract, which we consider next.

3. An Implementation

Having characterized the first-best outcome, we now consider the problem of finding the optimal dynamic contract in our setting with both moral hazard and asymmetric information. The task of finding the optimal contract is complex due to the huge space of fully contingent history-dependent contracts to consider. A contract (C,τ) must specify how the agent's consumption and the liquidation time depend on the entire history of cash flows. In classic settings with uncertainty only about the agent's effort but not the firm's productivity, there are standard recursive methods to deal with such complexity. These methods rely on dynamic programming using the agent's future expected payoff (a.k.a. continuation value) as a state variable.[12]

But with additional uncertainty and the potential for asymmetric information about the firm's productivity, these standard methods do not apply directly to our model. Thus we will take a different approach. We begin instead by conjecturing a simple and intuitive implementation for the contract. In our setting with moral hazard, if the agent had deep pockets the first-best liquidation policy could be attained by letting the agent own the firm. If the agent's wealth is limited, however, negative cash flow shocks can lead to inefficient liquidation. In order to minimize this inefficiency, it is natural to expect that in an optimal contract, the firm will build up cash reserves until it has an optimal level of financial slack. In this section we consider an implementation based on this intuition, and then show this implementation is incentive compatible. Though our implementation is rather simple, we will then verify the optimality of this contract in the space of all possible contracts in the following section.

[12] For example, see Spear and Srivastava (1987), Abreu, Pearce and Stacchetti (1990) (in discrete time) and DeMarzo and Sannikov (2006) and Sannikov (2007a) (in continuous time) for the development of these methods, and Piskorski and Tchistyi (2006) and Philippon and Sannikov (2007) for their applications.

3.1. Cash Reserves and Payout Policy

Consider an all-equity financed firm that uses cash reserves to provide financial slack.[13] Denote the level of its cash reserves by $M_t \geq 0$. Because the firm earns interest at rate r on these balances, its earnings at date t are given by

$$dE_t \equiv r\, M_t\, dt + dX_t = (r\, M_t + \delta_t)\, dt + \sigma\, dZ_t\,, \tag{12}$$

where we have assumed the agent's action $a_t = 0$.

If the firm uses these earnings to pay dividends dD_t, then its cash reserves will grow by

$$dM_t \equiv dE_t - dD_t = (r\, M_t + \delta_t)\, dt - dD_t + \sigma\, dZ_t\,. \tag{13}$$

Consider a contract in which the firm is forced to liquidate if it depletes its reserves and $M_t = 0$.[14] In order to avoid inefficient liquidation, we know from (11) that the firm must have reserves M_t of at least

$$M^1(\delta_t) \equiv \left(\delta_t - \underline{\delta}\right)/\nu\,. \tag{14}$$

Therefore, it is natural to suppose that if $M_t < M^1(\delta_t)$, the firm will retain 100% of its earnings in order to increase its reserves and reduce the risk of inefficient liquidation. In that case, dividends are equal to zero:

$$dD_t = 0 \text{ if } M_t < M^1(\delta_t). \tag{15}$$

Suppose the firm achieves the efficient level of reserves, so that $M_t = M^1(\delta_t)$. In order to maintain its reserves at the efficient level, using (14) and (2) we must have

$$dM_t = dM^1(\delta_t) = \frac{d\delta_t}{\nu} = \sigma dZ_t\,. \tag{16}$$

[13] We note that our proposed implementation is not unique, nor is it clear that it is optimal. The firm could also maintain financial slack through alternative means, such as a credit line or loan commitment. The analysis would be similar; for convenience we focus on the simplest implementation in terms of cash reserves. See Biais et al (2006) for a similar implementation based on cash reserves in a moral hazard setting without learning.

[14] Over any finite time period, the firm will experience operating losses with probability one; therefore, absent cash or credit, the firm must shut down. However, it is not yet clear whether it is optimal to deny the firm funds and force liquidation if $M_t = 0$. We will address optimality in the following section.

That is, to maintain the efficient level of reserves, the firm should adjust its cash balances by the "surprise" component of its earnings. Then from (13), dividends are equal to the firm's expected earnings:

$$dD_t = E\left[dE_t\right] = (rM_t + \delta_t)dt \quad \text{if } M_t = M^1(\delta_t). \tag{17}$$

The following result demonstrates that with this payout policy, the liquidation policy is first-best:

Proposition 2. *If $M_t = M^1(\delta_t)$ and if the firm follows the payout policy (17) after time t, then $M_\tau = 0$ if and only if $\delta_\tau = \underline{\delta}$.*

Proof: Given the payout policy (17), the firm's cash balance evolves according to (16). Therefore, $M_\tau = 0$ implies

$$\delta_\tau = \delta_t + \int_t^\tau d\delta_s = \delta_t + \nu\int_t^\tau dM_s = \delta_t + \nu\left(-M_t\right) = \delta_t - \nu M^1(\delta_t) = \underline{\delta}\,.$$

■

Finally, because there is no benefit from maintaining reserves in excess of the amount needed to avoid inefficient liquidation, we assume the firm pays out any excess cash immediately. Thus, we can summarize the firm's payout policy as follows:

$$dD_t = \begin{cases} 0 & \text{if } M_t < M^1(\delta_t) \\ (rM_t + \delta_t)dt & \text{if } M_t = M^1(\delta_t) \\ M_t - M^1(\delta_t) & \text{if } M_t > M^1(\delta_t) \end{cases}. \tag{18}$$

Under the payout policy described by (18), the firm accumulates cash as quickly as possible until it either runs out of cash and is inefficiently liquidated, or its reserves reach the efficient level. Once the efficient level of reserves is attained, the firm begins paying dividends at a rate equal to its expected earnings. It will continue to operate in this fashion unless δ_t falls to $\underline{\delta}$, in which case $M_t = 0$ and the firm is liquidated as in the first-best.

Figure 2 presents contract dynamics for an example. Until time 1.5 the firm has cash balances below the efficient level, and it stands the risk of being liquidated inefficiently. However, in this example inefficient liquidation does not happen. At time 1.5 the firm's cash level reaches the efficient target, and the firm initiates dividends. Dividends continue until the firm's profitability falls sufficiently and it is liquidated at date 5. The right panel of Figure 2 illustrates that the total quarterly dividends are significantly smoother than earnings.

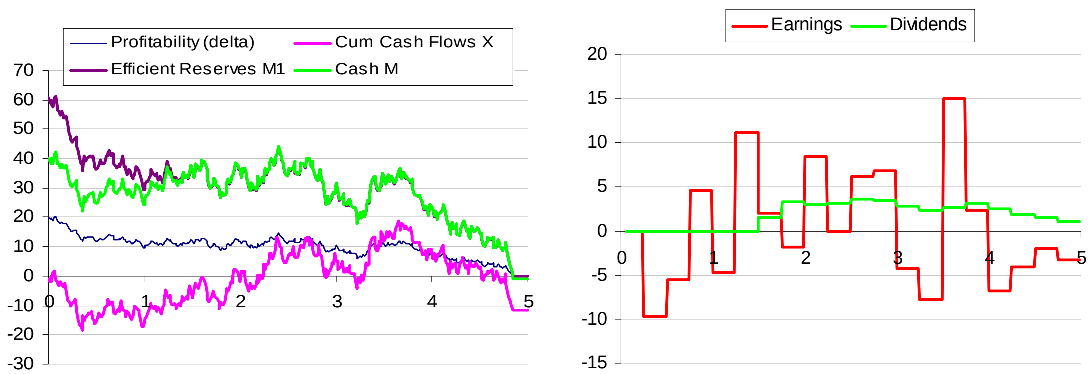

Figure 2: Contract dynamics when r = 5%, σ = 15, ν = 33%. The liquidation threshold is $\underline{\delta} = 0$.

3.2. Compensation and Incentive Compatibility

The requirement of cash reserves combined with the payout policy described above determines the liquidation time, τ, of the contract. To complete this implementation, we need specify the agent's compensation, C, and then assess whether the contract provides appropriate incentives.

Note that if the agent had unlimited wealth, we could provide the agent with appropriate incentives for effort by paying him a fraction λ of the firm's cash flows. This solution is not possible, however, since the firm's cash flows may be negative and the agent has limited liability. Given the implementation above, a natural alternative to consider is to pay the agent the fraction λ of the firm's *dividends* (rather than its cash flows), which are always non-negative.

This compensation can be interpreted as providing the agent a fraction λ of the firm's equity, with the proviso that the agent not receive any proceeds from a liquidation should it occur. This outcome could be implemented, for example, by giving outside investors preferred stock with complete priority in the event of liquidation. (Alternatively, the agent may receive a zero interest loan to purchase the shares, which becomes due in the event of liquidation.) We refer to the agent's compensation as equity that is *rescindable in the event of liquidation.*

Now we are ready to consider the agent's incentives. To verify incentive compatibility, we must determine the agent's payoff given different effort choices and payout policies. Consider the case in which the firm is already paying dividends; that is, $M_t = M^1(\delta_t)$. Suppose the agent follows the proposed implementation. Then from (17), the agent's expected payoff W_t is given by

$$\begin{aligned} E\left[\int_t^\tau e^{-r(s-t)}\lambda dD_s + e^{-r(\tau-t)}R(\underline{\delta})\right] &= E\left[\int_t^\tau e^{-r(s-t)}\lambda(rM_s+\delta_s)ds + e^{-r(\tau-t)}R(\underline{\delta})\right] \\ &= E\left[\int_t^\tau e^{-r(s-t)}r(\lambda M_s + R(\delta_s))ds + e^{-r(\tau-t)}(\lambda M_\tau + R(\delta_\tau))\right] \\ &= \lambda M_t + R(\delta_t) \end{aligned}$$

where the last equation follows from the fact that M_s and δ_s are martingales for $s > t$.[15]

Now consider a deviation in which the agent "cashes out" by immediately paying out all cash as a dividend, $dD_t = M_t$, and then defaulting. Under this strategy, the agent again receives a payoff of

$$W_t = \lambda\, M_t + R(\delta_t). \tag{19}$$

Thus there is no incentive for the agent to deviate in this way. We can similarly show that this implementation is robust to other types of deviations for the agent. For example, because (19) implies that the agent's payoff increases by λ for each dollar of additional cash held by the firm, there is no incentive for the agent to shirk and engage in outside

[15] Note that if X_s is a martingale, then $X_t = X_t\left(r\int_t^\tau e^{-r(s-t)}ds + e^{-r(\tau-t)}\right) = E\left[r\int_t^\tau e^{-r(s-t)}X_s ds + e^{-r(\tau-t)}X_\tau\right]$.

activities. The following result establishes the incentive compatibility, with regard to both the agent's effort choice and payout policy, of our proposed implementation:

Proposition 3. *Suppose the agent holds a fraction* λ *of the firm's equity, rescindable in the event of liquidation, and that liquidation occurs if the firm's cash balance falls to zero. Then for any effort strategy and payout policy, the agent's expected payoff is given by (19). Thus, it is optimal for the agent to choose actions* $a_t = 0$ *and to adopt the payout policy in (18).*

Proof: Consider an arbitrary payout policy D and effort strategy a. Define

$$V_t = \int_0^t e^{-rs}(\lambda dD_s + \lambda a_s)ds + e^{-rt}W_t$$

Then using the fact that δ is a martingale and that

$$E\left[dM_t\right] = rM_t dt + E\left[dX_t\right] - dD_t = rM_t dt + (\delta_t - a_t)dt - dD_t\,,$$

the drift of V_t is

$$\begin{aligned} E\left[dV_t\right] &= e^{-rt}\left(\lambda dD_t + \lambda a_t ds - rW_t + E\left[dW_t\right]\right) = \\ &= e^{-rt}\left(\lambda dD_t + \lambda a_t ds - r(\lambda M_t + \lambda\,\delta_t/r) + \lambda E\left[dM_t\right]\right) \\ &= 0 \end{aligned}$$

and so V_t is a martingale. Thus, because $M_\tau = 0$ so that $W_\tau = R(\delta_\tau)$, the agent's expected payoff from this arbitrary strategy is

$$\begin{aligned} E\left[\int_0^\tau e^{-rt}(\lambda dD_t + \lambda a_t dt) + e^{-r\tau}R(\delta_\tau)\ \right] &= E\left[V_\tau - e^{-r\tau}W_\tau + e^{-r\tau}R(\delta_\tau)\ \right] \\ &= E\left[V_\tau\ \right] = V_0 = W_0 \end{aligned}$$

and therefore the implementation is incentive compatible. ■

4. Justification of the Optimal Contract

Because standard methods do not apply directly to our model, in this section we develop a new approach to justify the optimality of our conjectured contract. While the specific solution is unique to our problem, we propose a three-step method to solve similar problems:

(1) Isolate the necessary incentive constraints, which are most important in limiting the attainable expected profit.

(2) Show that the conjectured contract solves the principal's optimization problem subject to just the necessary incentive constraints.

(3) Verify that the conjectured contract is fully incentive-compatible.

We conjectured a contract in the previous section, and verified its full incentive-compatibility in Proposition 3. We need to execute steps 1 and 2 of the verification argument.

Before we proceed, we note that the agent must take action $a_t = 0$ at all times in the optimal contract. The reason is that because $\lambda < 1$, it is cheaper to pay the agent directly rather than let him take actions for private benefit.

Lemma 1 (High Effort). In the optimal contract $a_t = 0$ until time τ.

Proof. Consider any contract in which sometimes $a_t > 0$, and let us show that there exists a better contract. Let us change it, by giving the agent an option to ask the principal for extra payments dC_t'. If the agent exercises this option at least once, then the agent's wages and termination time τ are determined as if the true path of output were

$$\hat{X}_t = X_t - \int_0^t \frac{dC_s'}{\lambda}.$$

Then the agent is indifferent between raising a_t above 0 and simply asking the principal for extra money. If he asks for money now whenever he was lowering effort previously, then the agent's strategy is incentive-compatible, and the principal's profit is strictly higher. Therefore, the original contract cannot be optimal. ■

From now on, we restrict attention to contracts with recommended effort $a_t = 0$.

4.1. Necessary Incentive Constraints

The necessary incentive-compatibility constraints are formulated using appropriately chosen state variables.[16] For our problem, we must include as state variables at least the principal's current belief about the agent's skill $\hat{\delta}_t$, which evolves according to

$$d\hat{\delta}_t = \nu(dX_t - \hat{\delta}_t dt), \quad \hat{\delta}_0 = \delta_0,$$

and the agent's continuation value when the agent follows the recommended strategy *(a_s)* after time *t*, and the principal has a correct belief about the agent's skill

$$W_t = E_t\left[\int_t^\tau e^{-r(s-t)} dC_s + e^{-r(\tau-t)} R(\delta_\tau) \mid \delta_t = \hat{\delta}_t\right] \text{ given strategy } \{a_t = 0\}.$$

The variables $\hat{\delta}_t$ and W_t are well-defined for any contract (C, τ), after any history of cash flows $\{X_s,\ s \in [0, t]\}$. However, they do not fully summarize the agent's incentives, which depend on the agent's *deviation payoffs*, the payoff that the agent would obtain if $\hat{\delta}_t \neq \delta_t$ due to the agent's past deviations. Therefore, we can formulate only necessary conditions for incentive compatibility using the variables $\hat{\delta}_t$ and W_t.

Lemma 2A which is standard in continuous-time contracting, provides a stochastic representation for the dependence of W_t on the cash flows $\{X_t\}$ in a given contract (C,τ). The connection between W_t and X_t matters for the agent's incentives.

Lemma 2A (Representation). *There exists a process $\{\beta_t, t \geq 0\}$ in L^* such that*

$$dW_t = rW_t dt - dC_t + \beta_t(dX_t - \hat{\delta}_t dt). \tag{20}$$

Proof. See Appendix A.

The process β_t determines the agent's exposure to the firm's cash flows shocks and therefore the strength of the agent's incentives under the contract. It is therefore natural

[16] For example, in continuous time Sannikov (2007b) solves an agency problem with adverse selection using the continuation values of the two types of agents as state variables, Williams (2007) solves an example with hidden savings using the agent's continuation value and his marginal utility as state variables.

to expect that β_t must be sufficiently large for the contract to be incentive compatible. For example, consider a deviation $\hat{a}_t = 1$. The agent can gain in two ways from this deviation. First, the agent earns the payoff λdt outside the firm. Second, from (4), the lower output of the firm reduces principal's estimate of the firm's productivity. If this deviation is the agent's first, then

$$d\hat{\delta}_t = \nu(dX_t - \delta_t dt) = d\delta_t - \nu dt\,.$$

Given these lowered expectations, the agent can continue to shirk and reduce effort by ν from that point onward and still generate cash flows consistent with the principal's expectations, for an additional expected perpetual gain of $\lambda\nu dt$. Because the deviation reduces the agent's contractual payoff by $\beta_t\, dt$, this deviation is profitable if

$$\beta_t < \lambda + \lambda\,\nu/r = \lambda\left(1 + \nu/r\right)$$

Lemma 2B below formalizes this intuition and establishes a necessary condition on β_t for a contract (C, τ) to be incentive compatible.

Lemma 2B (Incentive Compatibility). *Consider a contract (C,τ), for which the agent's continuation value evolves according to* (20). *A necessary condition for $\{a_t = 0\}$ to be incentive-compatible with respect to (C,τ) is that $\beta_t \geq \lambda(1 + \nu/r)$.*

Proof. See Appendix A.

4.2. Verification of Optimality

In this section we verify that our conjectured implementation is indeed an optimal contract. Recall from Section 3 that the agent's payoff in this contract is defined by

$$W_t = \lambda\, M_t + R(\delta_t),$$

where

$$dM_t = (r\, M_t + \delta_t)\, dt + \sigma\, dZ_t \text{ and } dD_t = 0 \text{ while } M_t < M^1(\delta_t) = \left(\delta_t - \underline{\delta}\right)/\nu \text{ and}$$

$$dM_t = \sigma\, dZ_t,\; M_t = M^1(\delta_t) \text{ and } dD_t = (rM_t + \delta_t)dt \text{ thereafter.}$$

It follows that

$dW_t = rW_t\,dt + \lambda(1+\nu/r)\sigma\,dZ_t$ and $dC_t = 0$ until W_t reaches $W^1(\delta_t)$, and

$$W_t = W^1(\delta_t) \text{ and } dC_t = rW_t\,dt \text{ thereafter,} \tag{21}$$

where

$$W^1(\delta) = R(\underline{\delta}) + \lambda(\tfrac{1}{\nu} + \tfrac{1}{r})(\delta - \underline{\delta}).$$

This evolution happens until W_t reaches $R(\delta_t)$, triggering liquidation.

Let us show that this contract attains the highest expected profit among all contracts that deliver value W_0 to an agent of skill level δ_0 and satisfy the necessary incentive-compatibility condition of Lemma 2B. The set of such contracts includes all fully incentive-compatible contracts. Since the conjectured contract is incentive-compatible, as shown in Proposition 3, it follows that it is also optimal.

Let us present a roadmap of our verification argument. First, we define a function $b(W_0,\delta_0)$, which gives the expected profit that a contract of Section 3 attains for any pair (W_0, δ_0) with $W_0 \geq R(\underline{\delta})$ and $\delta_0 \geq \underline{\delta}$. Proposition 4 verifies that this definition is indeed the expected payoff of outside equity holders in our implementation. After that, Proposition 5 shows that the principal's profit in any alternative contract that satisfies the necessary incentive-compatibility condition of Lemma 2B is at most $b(W_0, \delta_0)$ for any pair (W_0, δ_0) with $W_0 \geq R(\underline{\delta})$ and $\delta_0 \geq \underline{\delta}$. It follows that the conjectured contract of Section 3 is optimal.

For $W \geq R(\delta)$ and $\delta \geq \underline{\delta}$, define a function $b(W,\delta)$ as follows.

(i) For $W > W^1(\delta)$, let $b(W,\delta) = \bar{b}(\delta) - W$.

(ii) For $W = R(\delta)$, let $b(W,\delta) = L(\delta)$.

Otherwise, for $\delta > \underline{\delta}$ and $W \in (R(\delta), W^1(\delta))$, let $b(W,\delta)$ solve the equation

$$\begin{aligned} rb(W,\delta) &= \delta + rWb_W(W,\delta) + \\ &\tfrac{1}{2}\lambda^2(1+\tfrac{\nu}{r})^2\sigma^2 b_{WW}(W,\delta) + \tfrac{1}{2}\nu^2\sigma^2 b_{\delta\delta}(W,\delta) + \lambda(1+\tfrac{\nu}{r})\nu\sigma^2 b_{W\delta}(W,\delta) \end{aligned} \tag{22}$$

with boundary conditions given by (i) and (ii).

For an arbitrary contract (C,τ) with an incentive-compatible effort recommendation a, in which the agent's continuation value evolves according to (21), define the process

$$G_t = \int_0^t e^{-rs}(\delta_s ds - dC_s) + e^{-rt} b(W_t, \delta_t).$$

Note that on the equilibrium path we always have $\delta_t = \hat{\hat{\delta}}_t$.

Lemma 3 helps us prove both Propositions 4 and 5.

Lemma 3. *When $\delta_t \geq \underline{\delta}$ and C_t is continuous at t, then*

$$dG_t = e^{-rt}(\nu b_\delta(W_t,\delta_t) + \beta_t b_W(W_t,\delta_t))\sigma dZ_t - e^{-rt}(b_W(W_t,\delta_t)+1)dC_t +$$
$$e^{-rt}\sigma^2\left(\tfrac{1}{2}(\beta_t - \lambda(1+\tfrac{\nu}{r}))^2 b_{WW}(W_t,\delta_t) + (\beta_t - \lambda(1+\tfrac{\nu}{r}))(\lambda(1+\tfrac{\nu}{r})b_{WW}(W_t,\delta_t) + \nu b_{W\delta}(W_t,\delta_t))\right)dt$$

Proof. See Appendix A.

Proposition 4. *The conjectured optimal contract of subsection 3.1 attains profit* $b(W_0,\delta_0)$.

Proof. Under that contract, the process G_t is a martingale. Indeed, for all $t > 0$, the continuous process C_t increases only when $W_t = W^1(\delta_t)$ (where $b_W(W^1(\delta_t),\delta_t) = -1$) and $\beta_t = \lambda(1+\frac{\nu}{r})$, so G_t is a martingale by Lemma 3. At time 0, the agent consumes positively only in order for W_0 to drop to $W^1(\delta_0)$, and $b_W(W,\delta_0) = -1$ for $W \geq W^1(\delta_0)$, so G_t is a martingale there as well. Therefore, the principal attains the profit of

$$E\left[e^{-r\tau}b(W_\tau,\delta_\tau) + \int_0^\tau e^{-rt}(\delta_t dt - dC_t)\right] = E\left[G_\tau\right] = G_0 = b(W_0,\delta_0).$$

QED.

Proposition 5. *In any alternative incentive-compatible contract (C,τ) the principal's profit is bounded from above by $b(W_0,\delta_0)$.*

Proof. Let us argue that G_t is a supermartingale for any alternative incentive-compatible contract (C, τ) while $\delta_t \geq \underline{\delta}$.

First, whenever C_t has an upward jump of ΔC_t, G_t has a jump of $e^{-rt} (b(W_t+\Delta C_t, \delta_t) - b(W_t, \delta_t) - \Delta C_t) \leq 0$, since $b_W(W,\delta) \geq -1$ for all pairs (W,δ) (see Appendix B, which shows that b is concave in W).

Second, whenever C_t is continuous, then $\beta_t \geq \lambda(1+\nu/r)$ by Lemma 2B. By Lemma 3, the drift of G_t is

$$-e^{-rt}(b_W(W_t,\delta_t)+1)dC_t + \\ e^{-rt}\sigma^2\left(\tfrac{1}{2}(\beta_t-\lambda(1+\tfrac{\nu}{r}))^2 b_{WW}(W_t,\delta_t)+(\beta_t-\lambda(1+\tfrac{\nu}{r}))(\lambda(1+\tfrac{\nu}{r})b_{WW}(W_t,\delta_t)+\nu b_{W\delta}(W_t,\delta_t))\right)dt<0$$

since $b_W(W,\delta) \geq -1$ and, as shown in Appendix B,

$$b_{WW}(W,\delta) \leq 0 \text{ and } \lambda(1+\frac{\nu}{r})b_{WW}(W,\delta)+\nu b_{W\delta}(W,\delta) \leq 0 \tag{23}$$

for all pairs (W,δ).

Now, let $\underline{\tau}$ be the earlier of the liquidation time or the time when δ_t reaches the level $\underline{\delta}$. Then Proposition 1 implies that the principal's profit at time $\underline{\tau}$ is bounded from above by $b(W_{\underline{\tau}}, \delta_{\underline{\tau}})$. It follows that the principal's total expected profit is bounded from above by

$$E\left[e^{-r\tau}b(W_{\underline{\tau}},\delta_{\underline{\tau}})+\int_0^{\underline{\tau}} e^{-rt}(\delta_t dt - dC_t)\right] = E\left[G_{\underline{\tau}}\right] \leq G_0 = b(W_0,\delta_0).$$

QED

We conclude that Section 3 presents the optimal incentive-compatible contract for any pair (W_0, δ_0) such that $W_0 \geq R(\underline{\delta})$ and $\delta_0 \geq \underline{\delta}$. If $W_0 \geq W^1(\delta_0)$, then this contract attains the first-best profit, and liquidation always occurs at the efficient level of profitability of $\delta_\tau = \underline{\delta}$. If $W_0 < W^1(\delta_0)$, then liquidation happens inefficiently with positive probability.

5. Extensions of the Basic Model.

In this section we extend the basic model in a number of directions of practical interest. Specifically, we relax the assumptions that (1) the firm's cash flows are a sufficient statistic about future profitability and that (2) the agent's effort affects only the current cash flow and not future profitability. The principal may have other sources to learn information about the firm's future profitability, e.g. performance of comparable firms, and the agent may also have additional private information about future profitability, e.g. from observing firm's operations first-hand. More generally, assume that Brownian motion Z_t captures the firm's idiosyncratic risk, Z_t^M captures aggregate market (or industry) risk observable by both the principal and the agent, and Z_t^O captures the information that the agent observes privately. Without loss of generality, these Brownian motions can be taken to be independent. Let the cash flows to follow

$$dX_t = (\delta_t - a_t)dt + \sigma dZ_t + \sigma^M dZ_t^M,$$

and expected profitability evolve according to

$$d\delta_t = \nu(dX_t - (\delta_t - a_t)dt + \eta^M dZ_t^M) + \chi(\delta_t - a_t)dt + \eta^O dZ_t^O.$$

In this expression, parameters η^M and χ (which are typically positive) measure how news about other companies in the industry and the agent's current effort, respectively, affect the firm's future profitability.

To analyze this setting, assume at first that $\eta^O = 0$. Then, analogously to our basic model, the agent's continuation value can be represented as

$$dW_t = rW_t\,dt - dC_t + \beta_t\,(dX_t - \hat{\delta}_t\,dt) + \beta_t^M dZ_t^M.$$

A necessary condition for the optimality of effort $a_t = 1$ is

$$\beta_t \geq \lambda\,(1 + (\nu - \chi)/r). \qquad (*)$$

The logic is the following: If the agent lowers the firm's cash flow by 1 dollar, he derives an immediate private benefit of λ. At the same time, the principal's belief about δ_t goes down by ν, while the true value of δ_t goes down by χ relative to what δ_t would have been had the agent not deviated. As a result, the deviation fools the principal into thinking that the value of δ_t is lower by $\nu - \chi$ than its true value. The agent can "cash in" on this difference in beliefs to derive a value of at least $\lambda(\nu - \chi)/r$ by diverting cash from the firm at rate $\nu - \chi$ in perpetuity. We see that when the agent's effort adds to the firm's future profitability, he needs less motivation to put effort.

The necessary condition is in fact sufficient if it holds with equality at all times. Moreover, if $\eta^O \neq 0$ and the principal adjusts the agent's continuation value according to

$$dW_t = rW_t\,dt - dC_t + \lambda(1 + (\nu - \chi)/r)\,(dX_t - \hat{\delta}_t\,dt) + \beta_t^M dZ_t^M + \lambda(\nu - \chi)/r\;\eta^O d\hat{Z}_t^O,$$

where $\hat{Z}_t^O$ is the agent's report of Z_t^O, then the agent also has incentives to report privately observed innovations to firm profitability truthfully.

What about the value of β_t^M? One natural conjecture is that the value of β_t^M completely insures the agent against aggregate risk, that is, $\beta_t^M = -\lambda(1+(\nu-\chi)/r)\,\sigma^M$. This conjecture is incorrect, which we demonstrate by characterizing the optimal first best contract for the case when $\eta^M = 0$ and $\eta^O = 0$. (Note that if $\eta^M \neq 0$ or $\eta^O \neq 0$ then first best cannot be attained for any finite level of continuation value of the agent. Indeed, a sufficiently large drop in Z_t necessarily results in liquidation, which is always inefficient if there is a compensating rise in Z_t^M or Z_t^O that keeps δ_t fixed).

In order to attain first-best when $\eta^M = 0$ and $\eta^O = 0$, the firm must be able to absorb a negative cash flow shock of size $(\delta_t - \underline{\delta})/\nu$. If firm maintains cash balances of this amount, and if the agent holds a fraction $\lambda(1 - \chi/r)$ of firm equity, then the agent's continuation value is

$$\lambda\underline{\delta}/r + \lambda(1 + (r-\chi)/r)\,(\delta_t - \underline{\delta})/\nu,$$

and first-best is attained. Note that this is the minimal value for the agent required to attain first best, since a drop of Z_t, which lowers δ_t to $\underline{\delta}$, pulls the agent's continuation value down by at least $\lambda(1 + (r-\chi)/r)\,(\delta_t - \underline{\delta})/\nu$ to the agent's outside option of $\lambda\underline{\delta}/r$.

Note that the agent receives the same exposure to idiosyncratic and aggregate cash flow risk in this contract. The contract gives the agent more value when the firm becomes more profitable to prevent inefficient liquidation even for the case when the positive shock to profitability is aggregate.

6. Appendix A.

Proof of Lemma 2A. Note that

$$V_t = \int_0^t e^{-rs} dC_s + e^{-rt} W_t$$

is a martingale when the agent follows the recommended strategy (a_s). By the Martingale Representation Theorem there exists a process $\{\beta_t, t \geq 0\}$ in L^* such that

$$V_t = V_0 + \int_0^t e^{-rs} \beta_s (dX_s - \hat{\delta}_s ds),$$

since $dX_s - \hat{\delta}_s ds = \sigma dZ_s$ under the strategy $(a_s = 0)$. Differentiating with respect to t, we find that

$$dV_t = e^{-rt} dC_t + e^{-rt} dW_t - re^{-rt} W_t dt = e^{-rt} \beta_t (dX_t - \hat{\delta}_t dt) \Rightarrow$$

$$dW_t = rW_t dt - dC_t + \beta_t (dX_t - \hat{\delta}_t dt).$$

This expression shows how W_t determined by X_t (since $(\hat{\delta}_t)$ itself is determined by X_t), and therefore it is valid even if the agent followed an alterative strategy in the past. In this case W_t is interpreted as the continuation value that the agent would have gotten after a history of cash flows $\{X_s, s \leq t\}$ if his estimate of the firm's profitability coincided with the principal's, and he planned to follow strategy $(a=0)$ after time t. QED.

Proof of Lemma 2B. Suppose that $\beta_t < \lambda(1 + \nu/r)$ while $\hat{\delta}_t > 0$ on a set of positive measure. Let us show that the agent has a strategy $(\hat{a})$ that attains an expected payoff greater than W_0. Let $\hat{a}_t = \delta_t - \hat{\delta}_t$ when $\beta_t \geq \lambda(1 + \nu/r)$ and $\hat{a}_t = 1 + \delta_t - \hat{\delta}_t$ when $\beta_t < \lambda(1 + \nu/r)$ before the time τ when the agent is fired. Define the process

$$V_t = e^{-rt}\left(W_t + (\delta_t - \hat{\delta}_t)\frac{\lambda}{r}\right) + \int_0^t e^{-rs}(dC_s + \lambda \hat{a}_s ds).$$

If the agent follows the strategy described above, then before time $\hat{\tau}$,[17]

$$d\delta_t - d\hat{\delta}_t = \nu(dX_t - (\delta_t - \hat{a}_t)dt - (dX_t - \hat{\delta}_t dt)) = \nu(\hat{\delta}_t + \hat{a}_t - \delta_t) = 0 \text{ or } 1,$$

$$dW_t = rW_t dt - dC_t + \beta_t((\delta_t - \hat{a}_t)dt + \sigma dZ_t - \hat{\delta}_t dt),$$

and

$$\frac{dV_t}{e^{-rt}} = -r\left(W_t + (\delta_t - \hat{\delta}_t)\frac{\lambda}{r}\right)dt +$$
$$rW_t dt - dC_t + \beta_t((\delta_t - \hat{a}_t)dt + \sigma dZ_t - \hat{\delta}_t dt) + \lambda\nu(\hat{\delta}_t + \hat{a}_t - \delta_t) +$$
$$dC_t + \lambda \hat{a}_t dt = (\beta_t - \lambda - \nu\frac{\lambda}{r})\underbrace{(\delta_t - \hat{a}_t - \hat{\delta}_t)}_{0 \text{ or } -1}dt + \sigma dZ_t.$$

The drift of V_t is 0 if $\beta_t \geq \lambda(1 + \nu/r)$, and it equals $\lambda(1 + \nu/r) - \beta_t > 0$ when $\beta_t < \lambda(1 + \nu/r)$. At time τ the agent gets the payoff of $W_\tau + (\delta_\tau - \hat{\delta}_\tau)\frac{\lambda}{r} = R(\hat{\delta}_\tau) + (\delta_\tau - \hat{\delta}_\tau)\frac{\lambda}{r} = R(\delta_\tau)$. Therefore, the agent's total payoff from the strategy $(\hat{a})$ is

$$E\left[e^{-r\tau}\left(W_\tau + (\delta_\tau - \hat{\delta}_\tau)\frac{\lambda}{r}\right) + \int_0^\tau e^{-rs}(dC_s + \lambda \hat{a}_s ds)\right] = E[V_\tau] > V_0 = W_0.$$

We conclude that $\beta_t \geq \lambda(1 + \nu/r)$ when $\hat{\delta}_t > 0$ is a necessary condition for the incentive compatibility of the agent's strategy. ■

Proof of Lemma 3. Note that for $\delta \geq \underline{\delta}$, the function b satisfies partial differential equation (22) even if $W > W^1(\delta)$. Indeed, since $b(W,\delta) = \bar{b}(\delta) - W$ and $b_W = -1$ in that region, the equation reduces to

$$r(\bar{b}(\delta) - W) = \delta - rW + \tfrac{1}{2}\nu^2\sigma^2\bar{b}''(\delta).$$

This equation holds by the definition of $\bar{b}$.

[17] Note that $d\delta_t - d\hat{\delta}_t = \nu(\hat{\delta}_t + \hat{a}_t - \delta_t) = 0$ or ν implies that $\delta_t \geq \hat{\delta}_t$ for all t.

When C_t is continuous at t, then using Ito's lemma,

$$db(W_t,\delta_t) = (rW_t dt - dC_t)b_W(W_t,\delta_t) + \sigma^2\left(\tfrac{1}{2}\beta_t^2 b_{WW}(W_t,\delta_t) + \tfrac{1}{2}\nu^2 b_{\delta\delta}(W_t,\delta_t) + \beta_t \nu b_{W\delta}(W_t,\delta_t)\right)dt$$
$$+(\nu b_\delta(W_t,\delta_t) + \beta_t b_W(W_t,\delta_t))\sigma dZ_t = rb(W_t,\delta_t)dt - \delta_t dt - b_W(W_t,\delta_t)dC_t +$$
$$\sigma^2\left(\tfrac{1}{2}(\beta_t - \lambda(1+\tfrac{\nu}{r}))^2 b_{WW}(W_t,\delta_t)dt + (\beta_t - \lambda(1+\tfrac{\nu}{r}))(\lambda(1+\tfrac{\nu}{r})b_{WW}(W_t,\delta_t) + \nu b_{W\delta}(W_t,\delta_t))\right)dt,$$

where the second equality follows from (22). From the definition of G_t, it follows that Lemma 3 correctly specifies how G_t evolves. QED

7. Appendix B.

We must show that for all pairs (δ, W), the function $b(\delta, W)$ satisfies

$$b_{WW}(W,\delta) \le 0 \text{ and } \lambda(1+\frac{\nu}{r})b_{WW}(W,\delta) + \nu b_{W\delta}(W,\delta) \le 0.$$

It is useful to understand the dynamics of the pair (δ_t, W_t) under a conjectured optimal contract first. From (2) and (21), the pair (δ_t, W_t) follows

$$d\delta_t = \nu\sigma dZ_t \text{ and } dW_t = rW_t dt + \lambda(1+\nu/r)\sigma dZ_t \text{ until } W_t \text{ reaches } W^1(\delta_t),$$

$$\text{and } W_t = W^1(\delta_t) \text{ thereafter.} \tag{24}$$

When W_t reaches the level $R(\delta_t)$, termination results. The lines parallel to $W^1(\delta)$ are the paths of the joint volatilities of (W_t, δ_t). Due to the positive drift of W_t, the pair (W_t, δ_t) moves across these lines in the direction of increasing W_t. See the figure below for reference.

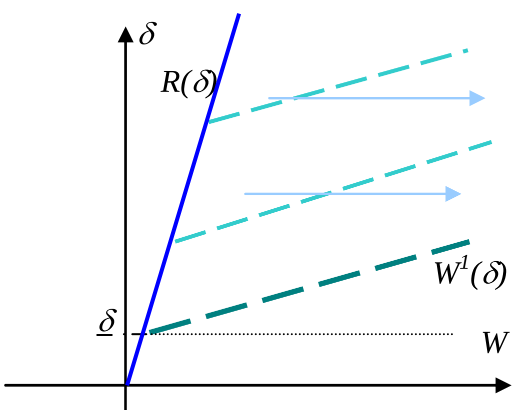

The phase diagram of (W_t, δ_t) provides two important directions: the direction of joint volatilities, in which $dW/d\delta = \lambda(1+\frac{\nu}{r})/\nu$, and the direction of drifts, in which W increases but δ stays the same. We need to prove that $b_W(\delta,W)$ weakly decreases in both of these directions.

To study how $b_W(W,\delta)$ depends on (W,δ), it is useful to know that $b_W(W_t, \delta_t)$ is a martingale (Lemma 4) and that $b_W(R(\delta),\delta)$ increases in δ (Lemma 5).

Lemma 4. *When the evolution of (W_t, δ_t) is given by* (24)*, then $b_W(W_t,\delta_t)$ is a martingale.*

Proof. Differentiating the partial differential equation for $b(W,\delta)$ with respect to W, we obtain

$$0 = rWb_{WW}(W,\delta) + \tfrac{1}{2}\lambda^2(1+\frac{\nu}{r})^2\sigma^2 b_{WWW}(W,\delta) + \tfrac{1}{2}\nu^2\sigma^2 b_{W\delta\delta}(W,\delta) + \lambda(1+\frac{\nu}{r})\nu\sigma^2 b_{WW\delta}(W,\delta).$$

The right hand side of this equation is the drift of $b_W(W,\delta)$ when $W_t < W^1(\delta_t)$ by Ito's lemma. When $W_t = W^1(\delta_t)$, then $b_W(W_t, \delta_t) = -1$ at all times. Therefore, $b_W(W_t, \delta_t)$ is always a martingale. QED

Lemma 5. *$b_W(R(\delta), \delta)$ weakly increases in δ.*

Proof. Note that

$$b(W_0,\delta_0) = \bar{b}(\delta_0) - W_0 - E\left[e^{-r\tau}(\bar{b}(\delta_\tau) - L(\delta_\tau) - R(\delta_\tau)) \mid \delta_0, W_0\right]. \tag{25}$$

That is, the principal's profit equals first-best minus the loss of payoff due to early inefficient liquidation. Let us show that for all $\varepsilon > 0$, $b(R(\delta_0)+\varepsilon,\delta_0) - b(R(\delta_0),\delta_0)$ increases in δ_0. Consider the processes (W_t^i,δ_t^i) (i = 1,2) that follow (24) starting from values δ_0^1 and $\delta_0^2 = \delta_0^1 + \Delta$ and $W_0^i = R(\delta_0^1) + \varepsilon$. Then for any path of Z, the process for i = 1 ends up in liquidation at a sooner time τ_i and at a higher value of $\delta^1_{\tau_1}$. Indeed, from the law of motion (24), it is easy to see that while the difference δ_t^2 -δ_t^1 stays constant at all times, $W_t^2 - W_t^1$ becomes larger than $\lambda/r\Delta$ after time 0, where λ/r is the slope of $R(\delta)$. Since

$$\bar{b}(\delta)-L(\delta)-R(\delta)$$

increases in Δ, it follows that

$$E\left[e^{-r\tau_1}(\bar{b}(\delta^1_{\tau_1})-L(\delta^1_{\tau_1})-R(\delta^1_{\tau_1}))\,|\,\delta^1_0,W^1_0\right]\geq E\left[e^{-r\tau_2}(\bar{b}(\delta^2_{\tau_2})-L(\delta^2_{\tau_2})-R(\delta^2_{\tau_2}))\,|\,\delta^2_0,W^2_0\right].$$

As a result,

$$\begin{aligned}
&b(R(\delta^1_0)+\varepsilon,\delta^1_0)-b(R(\delta^1_0),\delta^1_0)=\\
&-E\left[e^{-r\tau_1}(\bar{b}(\delta^1_{\tau_1})-L(\delta^1_{\tau_1})-R(\delta^1_{\tau_1}))\,|\,\delta^1_0,W^1_0\right]+\bar{b}(\delta^1_0)-L(\delta^1_0)-R(\delta^1_0)\leq\\
&-E\left[e^{-r\tau_2}(\bar{b}(\delta^2_{\tau_2})-L(\delta^2_{\tau_2})-R(\delta^2_{\tau_2}))\,|\,\delta^2_0,W^2_0\right]+\bar{b}(\delta^2_0)-L(\delta^2_0)-R(\delta^2_0)=\\
&b(R(\delta^2_0)+\varepsilon,\delta^2_0)-b(R(\delta^2_0),\delta^2_0),
\end{aligned}$$

where we used (25) to derive the first and the last inequality. QED

We can use Lemmas 4 and 5 to reach conclusions about how $b_W(W,\delta)$ changes as W increases or as δ and W increase in the direction $dW/d\delta=\lambda(1+\tfrac{\nu}{r})/\nu$.

Lemma 6. *$b_W(W,\delta)$ weakly decreases in W.*

Proof. Let us show that for any $\delta_0 \geq \underline{\delta}$, for any two values $W_0^1 < W_0^2$, $b_W(W_t^1, \delta_0) \geq b_W(W_t^2, \delta_0)$.

Consider the processes (W_t^i,δ_t) (i = 1,2) that follow (24) starting from values (W_0^1,δ_0) and (W_0^2,δ_0) for $\delta_0^1 < \delta_0^2$. Then for any path of Z, we have $W_t^2 - W_t^1 \geq 0$ until time τ_1 when W_t^1 reaches the level of $R(\delta_t)$. The time when W_t^2 reaches the level of $R(\delta_t)$ is $\tau_2 \geq \tau_1$. Since $W^2_{\tau_1}\geq W^1_{\tau_1}=R(\delta_{\tau_1})$, it follows that $\delta_{\tau_2}\leq\delta_{\tau_1}$ and $W^2_{\tau_2}\leq W^1_{\tau_1}$. Using Lemmas 4 and 5,

$$b_W(W^1_0,\delta_0)=E\left[b_W(R(\delta_{\tau_1}),\delta_{\tau_1})\right]\geq E\left[b_W(R(\delta_{\tau_2}),\delta_{\tau_2})\right]=b_W(W^2_0,\delta_0).$$

QED

Lemma 7. *$b_W(W,\delta)$ weakly decreases in the direction, in which W and δ increase according to* $dW/d\delta=\lambda(1+\tfrac{\nu}{r})/\nu$.

Proof. Consider starting values (W_0^i, δ_0^i) that satisfy

$$\delta_0^2 - \delta_0^1 = \Delta > 0 \text{ and } W_0^2 - W_0^1 = \Delta\lambda(1+\tfrac{\nu}{r})/\nu.$$

Starting from those values, let the processes (W_t^i, δ_t^i) (i = 1,2) follow (*). Then for any path of Z, at all times $\delta_t^2 - \delta_t^1 = \Delta$ and $W_t^2 - W_t^1 \geq \Delta\lambda(1+\tfrac{\nu}{r})/\nu$ (with equality after time 0 only if $W_t^2 = W^1(\delta_t^2)$ and $W_t^1 = W^1(\delta_t^1)$). Therefore, the time τ_1 when W_t^1 reaches the level of $R(\delta_t^1)$ occurs at least as soon as the time τ_2 when W_t^2 reaches the level of $R(\delta_t^2)$. Also, since $W_{\tau_1}^2 \geq W_{\tau_1}^1 + \Delta\lambda(1+\tfrac{\nu}{r})/\nu > R(\delta_{\tau_1}) + \Delta\lambda/r$, it follows that $\delta_{\tau_2} \leq \delta_{\tau_1}$ and $W_{\tau_2}^2 \leq W_{\tau_1}^1$. Using Lemmas 4 and 5,

$$b_W(W_0^1, \delta_0^1) = E\left[b_W(R(\delta_{\tau_1}^1), \delta_{\tau_1}^1)\right] \geq E\left[b_W(R(\delta_{\tau_2}^2), \delta_{\tau_2}^2)\right] = b_W(W_0^2, \delta_0^2).$$

QED

8. References.

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

Albuquerque, R. and Hopenhayn, H. A. (2004) "Optimal Lending Contracts and Firm Dynamics," *Review of Economic Studies*, Vol. 72(2), No. 247: 285-315.

Allen, F. and R. Michaely, 2003, Payout Policy, in G. Constantinides, M. Harris, and R. Stulz, eds., Handbooks of Economics, North-Holland.

Atkeson, Andrew (1991) "International Lending with Moral Hazard and the Risk of Repudiation," *Econometrica* 59, 1069-1089.

Atkeson, Andrew, and Harold Cole (2005) "A Dynamic Theory of Optimal Capital Structure and Executive Compensation," NBER working paper No. 11083.

Baron, David P. and David Besanko (1984) "Regulation and information in a continuing relationship," *Information Economics and Policy*, 1:267-302.

Benartzi, Shlomo, Roni Michaely, and Richard Thaler, 1997, Do changes in dividends signal the future or the past? *Journal of Finance* 52: 1007-1043

Benartzi, Shlomo, Gustavo Grullon, Roni Michaely, and Richard Thaler, 2005, Dividend changes do not signal changes in future profitability, *Journal of Business* 78: 1659-1682

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

Biais, B., T. Mariotti, J.-C. Rochet and S. Villeneuve (2007) “Environmental Risk Insurance under Dynamic Moral Hazard,” working paper, Université de Toulouse

Brav, A., Graham, J., Harvey, C. and R. Michaely, 2005, Payout Policy in the 21st Century, *Journal of Financial Economics*

Clementi, G. L. and Hopenhayn, H. A. (2006) “A Theory of Financing Contracts and Firm Dynamics,” *Quarterly Journal of Economics*, 121, 229-265.

DeMarzo, Peter, and Michael Fishman (2007) “Optimal long-term financial contracting,” *Review of Financial Studies*, 20:2079-2128.

DeMarzo, Peter, and Michael Fishman, (2007) “Agency and optimal investment dynamics,” *Review of Financial Studies*, 20:151-188.

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.

Dixit, A. K., and R.S. Pindyck (1994), *Investment under Uncertainty*. Princeton University Press.

Fama, Eugene and Harvey Babiak, 1968, Dividend policy: an empirical analysis, *Journal of the American Statistical Association* 63: 1132-1161

Fama, Eugene F. and Kenneth R. French, 2002, “Testing trade-off and pecking order predictions about dividends and debt,” *Rev. Financ. Stud.* 15: 1-33.

Freixas, X., R. Guesnerie and J. Tirole (1985) “Planning under Incomplete Information and the Ratchet Effect,” *Review of Economic Studies*, 52(2), 173-191.

Grullon, Gustavo and Roni Michaely, 2002, Dividends, share repurchases and the substitution hypothesis, *Journal of Finance*, 57 , 1649-84.

Grullon, Gustavo, Roni Michaely and Bhaskaran Swaminathan, 2002, Are dividend changes a sign of firm maturity? *The Journal of Business*, 75, 387-424.

Healy, Paul M. and Krishna G. Palepu, 1988, Earnings information conveyed by dividend initiations and omissions, *Journal of Financial Economics* 21: 149-176

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

Laffont, J-J. and J. Tirole (1988) “The Dynamics of Incentive Contracts,” *Econometrica* 56(5), 1153-1175.

Lintner, John., 1956, Distribution of Incomes of Corporations Among Dividends, Retained Earnings, and Taxes, *American Economic Review*, *46(2)*, 97-113.

Ljungqvist, L. and T. Sargent, 2004, *Recursive Macroeconomic Theory*, MIT Press.

Miller, Merton H. and Franco Modigliani (1961), “Dividend Policy, Growth, and the

Valuation of Shares," *The Journal of Business* 34(4), 411-433.

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

Philippon, T. and Y. Sannikov (2007) "A Dynamic Agency Model of IPOs," working paper, New York University.

Piskorski, T., and A. Tchistyi (2007): "Optimal Mortgage Design," working paper, New York University.

Praveen Kumar, Bong-Soo Lee: Discrete dividend policy with permanent earnings Financial Management, Autumn, 2001

Roberts, Michael R. and Michaely, Roni, "Corporate Dividend Policies: Lessons from Private Firms" (February 22, 2007). Available at SSRN: http://ssrn.com/abstract=927802

Rogerson, W. P. (1985) "Repeated Moral Hazard," *Econometrica* 53, 69-76.

Sannikov, Yuliy (2007a) "A Continuous-Time Version of the Principal-Agent Problem," forthcoming in the *Review of Economic Studies*.

Sannikov, Yuliy (2007b) "Agency Problems, Screening and Increasing Credit Lines," working paper, University of California at Berkeley.

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

Stokey, N. and R. E. Lucas, Jr., 1989, *Recursive Methods in Economic Dynamics*, Harvard University Press.

Williams, Noah (2007) "On Dynamic Principal-Agent Problems in Continuous Time," working paper, Princeton University.