Real Options and Dynamic Incentives

Eduardo Faingold Suvi Vasama

March 31, 2014

Abstract

We examine a dynamic principal-agent model in which the output is correlated over time. The optimal contract determines the players' share of the firm cash-flow and a liquidation policy. Incentive compatibility, together with the agent's limited liability, requires that the firm is liquidated following a history of low returns. With correlated outcome, the optimal liquidation decision depends both on the firm profitability and the players' shares of the firm cash-flow. The firm is liquidated more inefficiently if the principal's share is high. The payments to the agent are delayed, and he is rewarded by promising him a higher share of the future returns. Once the agent's share grows high enough, the firm is operated efficiently. In particular, the firm is only liquidated if it is efficient to do so.

1 Introduction

This paper studies dynamic incentives in a real option problem. A principal invests in a firm that produces a stochastic cash-flow. At any point in time, she has the option to terminate the operations and realize a liquidation payoff. As long as the firm is operated, the principal needs to employ an agent to handle the daily operations. The agent has superior knowledge about the firm value and can use the informational advantage to generate a private benefit. To prevent the agent from enjoying private benefits, the principal rewards him by leaving him a share of the firm cash-flow as a compensation, and punishes him by liquidating the firm. The players' shares of the firm cash-flow, as well as its liquidation policy, are determined in a contract on which the players agree in the beginning of the relationship.

Our model describes a natural connection between agency problems and firm profitability. In practice, firm efficiency and leverage are important factors explaining why firms are liquidated under circumstances of financial distress. Such a phenomenon can be interpreted as a natural consequence of an asymmetric information problem between an entrepreneur (agent) and a financier (principal). Because of the agency problem, the financier has to punish the entrepreneur for low outcome realizations, and leave him part of a high surplus as a reward.

In more leveraged firms, the financier has higher claims on the firm cash flow whereas the entrepreneur's share of the output is lower. To provide the agent incentives, the optimal contract threatens to liquidate the firm inefficiently. Moreover, if the agent's share is high, liquidation is not as crucial for incentive provision. However, the firm might still be liquidated if its operations become unprofitable. Firms with low profitability are always liquidated for efficiency reasons. The phenomenon is well documented in Zingales (1998) who provides empirical evidence in tracking industry.

Our starting point is a dynamic principal-agent model in which both players are risk-neutral, but the agent is protected by limited liability. Limited liability implies that the contract cannot impose negative payments to the agent, and that losses have to be covered by the principal. The firm produces a stochastic cash-flow that is unobservable by the principal. The cash-flow is reported by the agent who has the opportunity to conceal part of it from her. Such diversion generates the agent a private benefit, but is related with a social cost. To prevent the agent from diverting cash-flows, the principal has to reward him for reporting high returns.

The principal has two instruments to provide the agent incentives: nonnegative payments and liquidation. If the agent reports high cash-flows, the principal leaves him part of the surplus as a compensation for his report. The agent has no incentives to misreport cash-flows if his share of the cash-flow compensates for the private benefit that he enjoys if he diverts.

The liquidation is irreversible. Firm assets are sold in fire sales and the returns are collected by the principal. With correlated cash-flows the firm value depends on the output, and the liquidation decision corresponds to the realization of a real option. If the cash-flow is high, the expected future returns are high and so is the expected value of future operations. When the output decreases, liquidation becomes more attractive. If the output is low, the firm value is low and the liquidation option relatively more valuable.

With correlated cash-flows past returns predict future outcomes. If the agent reports truthfully, he reveals the expected future cash-flow. If he conceals cash-flows, the players' perceptions of the firm value diverge. The information is payoff relevant. If the agent diverts cash-flows, the principal becomes more pessimistic about the future. Therefore, she has an additional incentive to liquidate the firm following a low cash-flow report.

In our framework, the liquidation also works as an incentive mechanism. If the firm is liquidated, the agent loses the possibility to enjoy future payments or private benefits. The liquidation threat makes low reports less attractive for the agent and saves on the incentive cost in the high states. Thus, the principal has two motives for liquidating the firm: incentive provision and efficiency.

The correlation implies that the agent's action carries information about the current economic environment. If the agent deviates, the contract conditions the incentive scheme on the principal's misperception of the current state. Therefore, the agent's current deviation may distort his incentives in following periods.[1] As a consequence, the agent's optimal strategy might not include truthtelling following a deviation. We need to verify that the agent cannot increase his payoff by deviating and gaining a more beneficial environment.

If the agent deviates, the principal does not know his true continuation value. The agent has persistent private information, which provides a challenge from modeling point of view. To keep the analysis tractable, we restrict the attention on environments in which the agent's private benefit is proportional to the diverted cash-flow, and the outcome follows a Markov process. In such a Markov environment, any divergence in expectations is corrected once the agent reports truthfully. We show that the agent's incentive compatibility conditions can be derived by identifying his payoff of a particular kind of deviation. The deviation strategy prescribes that the agent returns to the equilibrium path after his deviation, and reports truthfully from there onwards. Then we verify that the agent has the right incentives to report truthfully even if he deviated in the past. Interestingly, it turns out that since the agent is risk-neutral, and his private benefit proportional to the diverted cash-flow, he cannot benefit from the principal's misperception of the current state.

Models in which the agent has private information about the underlying economic environment are generally known to be challenging. Adopting an approach that is common in the literature,[2] we solve a relaxed problem where we replace the agent's global incentive compatibility conditions by a weaker necessary condition that we obtain by examining a particular kind of deviation. Then we show that, at the optimum of the relaxed problem, the agent's necessary incentive constraint is satisfied with equality. Finally, we verify that the agent's global incentive compatibility condition is satisfied at the optimum. This implies that the solution of the relaxed problem is the optimum of the unrelaxed problem as well.

The optimal contract admits the following dynamics. In the beginning of the

[1]Similar issues have been examined in adverse selection literature. Besanko (1985), and more recently, Battaglini (2005), Garrett and Pavan (2011, 2012), Pavan, Segal and Toikka (2012), and Eso and Szentes (2013) find conditions under which the optimal contract can be found by using a first-order approach. Battaglini and Lamba (2012) provide a set of counterexamples in which the first-order approach fails.

[2]See, for example, DeMarzo and Sannikov (2011), Pavan, Segal and Toikka (2012) and Eso and Szentes (2013).

relationship, the optimal contract relies on an inefficient liquidation threat to provide the agent incentives. The payments to the agent are delayed, and the agent is rewarded by promising him a higher continuation value for the future. With high performance, the agent's continuation value increases such that the first-best solution becomes attainable. The firm is liquidated efficiently, and the agent is rewarded with immediate payments.

We contribute to the fast growing literature on dynamic principal-agent models with risk-neutrality and limited liability. The seminal continuous-time cash-flow diversion model was presented in DeMarzo and Sannikov (2006). They derive the optimal contract with independently and identically distributed cash-flows, and show how to implement it using a credit line. Biais, Mariotti, Plantin and Rochet (2007) prove that the discrete time counterpart of the model converges to its continuous-time limit, and provide an alternative implementation using cash-reserves.[3] In these models, the firm value is independent of the past performance, and the firm is never liquidated in the first-best solution. We extend the framework by allowing the cash-flow to be correlated such that the efficient solution solves a standard real option problem.

Our solution concepts borrow from DeMarzo and Sannikov (2011) who examine a related model in which the players learn about the unknown mean of the cash-flow process over time. Our framework differs from theirs in several important dimensions. First, in our model, the firm value is known at equilibrium, but stochastically changes over time. Thus our model has no learning at equilibrium, but the players are contracting in a changing environment. Besides, in DeMarzo and Sannikov, the players' outside options depend on the firm value. We assume that the liquidation value is independent of the cash-flow.

The optimal contracting problem can be arbitrarily complex, because the optimal decision might depend on the entire history of past returns. To reduce the complexity, we adopt a classical approach that uses the agent's continuation value as a state variable to summarize the dependence of the optimal contract on the entire history. The reduction allows us to characterize the optimal contract in a complex economic environment, and is justified by Spear and Srivastava (1989) who show that the agent's continuation value is a sufficient statistic for the history in the optimal contract. Besides, because of correlation, we need a second state variable, the current cash-flow, to summarize the firm value.

Our model is also related with literature on principal-agent framework with a risk-averse agent. In continuous time, the benchmark model with independent output in was presented by Sannikov (2008) who derives the optimal contract using

[3] DeMarzo and Fishman (2007) derive the optimal contract in a discrete-time version of the model. Biais, Mariotti, Rochet and Villeneuve (2010) consider Poisson distributed shocks and allow for investment.

the agent's continuation value as a state variable. Our work also greatly benefits from Sannikov (2013) who examines a model in which the agent's actions have long-term consequences. When designing the optimal incentive scheme, the principal has to take into account that the agent's action today affects his incentives in the future, as well as the future output directly. In our framework, the agent's action has no effect on the future output, and any divergence in expectation is corrected once the agent reports truthfully.

To model correlation of output, we adopt an approach that defines cash-flow as the level of a diffusion process, rather than as the increment. In the principal-agent framework, the approach was first considered by Williams (2011) who examines a model in which the cash-flow is persistent, but the agent is risk-averse. Strulovici (2011) extends the model to allow for the players to renegotiate the contract. The models do not consider liquidation. When the agent is risk-neutral, there is no motive for consumption smoothing. The assumption simplifies the problem considerably. The simplicity of the setting allows us to explore the connection between the agency problem and the real option problem in more detail.

2 The Model

We examine a game with two players: a principal and an agent. The principal has access to unlimited funds, but the agent is protected by limited liability. In our framework this assumption implies that the agent cannot make negative payments. Both players discount the future by a common rate $r > 0$.

At time t, the firm produces a cash-flow x_t. The cash-flow is stochastic and evolves according to a Brownian motion. Under this assumption, the cash-flow at period t is

$$x_t = x_0 + \int_0^t dx_s, \tag{1}$$

where x_0 is the initial cash-flow.

The true cash-flow is unobservable by the principal and is reported by the agent. The agent has the opportunity to conceal part of the cash-flow from the principal. Formally, we let $\hat{x}_t$ denote the cash-flow that the agent delivers to the principal. At time t, the agent's report evolves according to

$$d\hat{x}_t = dx_t - l_t dt. \tag{2}$$

From the principal's point of view, the process $\hat{x}_t$ is a Brownian motion. The agent observes the true Brownian motion x_t, and chooses the drift l_t, potentially reducing the principal's payoff. The initial cash-flow x_0 is common knowledge.

Our formulation only allows the agent to choose reporting strategies that are absolutely continuous with respect to the true probability measure. That is, we do not allow the agent's reported process to exhibit jumps. Since we only consider full commitment contracts under which the agent reports truthfully at equilibrium, the restriction is without loss of generality.

If the agent conceals part of the cash-flow, he earns a private benefit that is proportional to the diverted amount. Formally,

$$\lambda(x_t - \hat{x}_t) = \lambda \int_0^t l_s ds.$$

Under this specification, the agent enjoys a private benefit of λ for each dollar that he diverts. We assume that $\lambda \in (0, 1)$ such that cash-flow diversion is inefficient. Hence there is a social loss of $1 - \lambda$ for each diverted dollar. The agent cannot save privately, but has to cover the cost of his report from the current outcome. Thus $\hat{x}_t \leq x_t$ by limited liability.

We solve for the optimal contract that provides the agent the right incentives to report cash-flows truthfully. That is, we restrict the attention on contracts under which the agent optimally chooses $x_t = \hat{x}_t$ for all t. The restriction is without loss of generality, and is justified by the revelation principle.

In the beginning of the relationship, the players agree on a contract. The contract specifies payments $\{w_t \geq 0 : t \geq 0\}$ from the principal to the agent, and a stopping time τ, at which the firm is liquidated. Both the payment process and the stopping time are adapted to the filtration generated by the past reported cash-flows, $\mathcal{F}_t = \sigma\{\hat{x}_s : 0 \leq s \leq t\}$. We assume that both players can fully commit to the contract for its lifetime.

The liquidation is irreversible, and generates a value $L \geq 0$ for the principal.[4] Given that the agent reports truthfully, the principal's total expected profit is

$$v_0 = E\left[\int_0^\tau re^{-rt}(x_t - w_t)dt + e^{-rt}L\right]. \tag{3}$$

Let U_0 denote the agent's promised utility from the contract at time 0. The promise-keeping constraint at time 0 guarantees that the agent receives his expected payoff

$$U_0 = E\left[\int_0^\tau re^{-rt}w_t dt\right] \tag{4}$$

under the contract if he chooses the truthful reporting strategy $\hat{x}_t = x_t$ for all t. The agent's incentive compatibility constraint guarantees that the agent receives at least the same utility

$$U_0 \geq \hat{E}\left[\int_0^\tau re^{-rt}(w_t + \lambda(x_t - \hat{x}_t))dt\right] \tag{5}$$

[4] To simplify the expressions, we normalize $L = r\tilde{L}$.

for reporting truthfully from t onwards than for any arbitrary reporting strategy $\{\hat{x}_s \leq x_s : s \geq t\}$. To simplify expressions, we normalize the agent's liquidation value to 0.

The optimal contract determines an intertemporal payment rule and a liquidation policy that satisfy the agent's promise-keeping condition, the incentive compatibility constraints, and the limited liability constraints, and maximize the principal's expected payoff. Formally, it chooses a nonnegative payment process w and a stopping time τ to maximize (3) subject to the constraints (4) and (5), and the nonnegativity constraints on the payment process.

3 Incentive Compatibility

As is standard in the principal-agent framework, we first solve for conditions for the agent to report the cash-flow truthfully. In this section, we derive a necessary condition for the agent to report truthfully. In the next section we show that it is optimal for the principal to let the necessary condition bind at the optimum. Finally, we verify in Section 5.1 below that the agent's global incentive compatibility condition is satisfied at the optimum.

To derive conditions for the agent to report truthfully, we need to specify how his continuation value U_t depends on the reported cash-flow. Following the standards in the literature,[5] we derive a representation for U_t as a stochastic process, and determine how the process evolves in response to the reported Brownian motion. Then we derive necessary conditions for truthful reporting on the equilibrium path.

At $t \leq \tau$, the agent's promised value from the contract is

$$U_t = E_t \left[\int_t^{\tau} r e^{-r(s-t)} w_s ds \right]. \tag{6}$$

If the agent reports truthfully from t on, such that the true output x_t, and the reported cash-flow $\hat{x}_t$ agree from t onwards, (6) describes the agent's continuation value. Whenever $\hat{x}_t < x_t$, the principal's assessment of the agent's continuation value differs from his actual value.

Moreover, we need to rule out contracts that delay payments to the agent forever. In particular, we allow for the principal only to choose contracts such that $U_t < \infty$ for all t.

The following lemma describes how the agent's promised value U_t evolves in response to his reports

[5]See, for example, Sannikov (2008).

Lemma 1. *Fix a contract $\{w, \tau\}$ with $U_t < \infty$ for all t. The process U_t is the agent's promised value from the contract if and only if the following conditions are satisfied.* (*i*) *At $t \leq \tau$, U_t admits the representation*

$$dU_t = r(U_t - w_t)dt + \beta_t d\hat{x}_t. \tag{7}$$

β is a progressively measurable process in L^2 that describes the sensitivity of the agent's promised value to his reports.[6] (*ii*) *U_t satisfies the transversality condition $\lim_{s \to \infty} E_t[1_{s \leq \tau} e^{-rs} U_{t+s}] = 0$ almost everywhere.*

Proof. See Appendix. □

Next, we determine necessary conditions for the agent to report cash-flows truthfully. The challenge here is that if the agent deviated in the past, the process (6) is not the agent's continuation value. To deal with the challenge, we first determine necessary conditions for incentive compatibility by examining the agent's gain of a particular kind of deviation. To be more concrete, we consider a deviation strategy that prescribes that the agent returns to the equilibrium path after his deviation, and reports truthfully from there onwards. Such a deviation strategy describes the counterpart of a one-shot deviation in our continuous-time, Markov environment. Of course, we need to verify that the conditions we derived are sufficient for incentive compatibility.

First, notice that, because of limited liability, U_t cannot become negative. Therefore, if the agent's continuation value decreases to 0, the only way to provide him incentives to report truthfully is to liquidate the firm. Besides, since the agent is risk-neutral, and cannot gain from a deviation after τ, the principal has no incentive to provide him income after τ. Therefore, it is without loss of generality to concentrate on contracts that determine liquidation as the first time that the agent's continuation value hits 0. Formally, let

$$\tau = \inf\{t : U_t = 0\}. \tag{8}$$

Next, we derive necessary conditions for the agent to report truthfully. The challenge here is that we need to compare the flow value of the agent's private benefit, $\lambda(x_t - \hat{x}_t)dt$, with the instantaneous loss of his continuation value, $\beta_t l_t dt$. To overcome the challenge, we consider a particular deviation strategy. The deviation strategy prescribes that the agent departs from the equilibrium path at time t on a time interval of length Δ, and returns to the equilibrium path on a time interval of length Δ^2. By letting $\Delta \to 0$, we can identify the agent's instantaneous gain from diverting cash-flows.

To be more concrete, consider the following strategy chosen by the agent. At period t, the agent

[6] A process β is in L^2 if $E\left[\int_0^t 1_{s \leq \tau} \beta_s^2 ds\right] < \infty$.

- Departs from the equilibrium path by choosing the drift l on the time interval $[t, t+\Delta]$.
- Returns to the equilibrium path by choosing the drift $-l/\Delta$ on the time interval $(t+\Delta, t+\Delta+\Delta^2]$.
- Reports truthfully from $t+\Delta+\Delta^2$ onwards.

Then at time $t+\Delta+\Delta^2$, $\hat{x}_{t+\Delta+\Delta^2} = x_{t+\Delta+\Delta^2}$. If l is not too large, such that τ is not reached, we can identify the agent's instantaneous gain from diverting cash-flows.

Of course, the agent is only willing to report truthfully if his continuation value increases at least as much as his private benefit would have increased if he had concealed the cash-flow. At t, the agent's instantaneous gain from diverting a dollar is λ. Therefore, it is optimal for the agent to report truthfully if his continuation value changes by $\beta_t \geq \lambda$ for each dollar reported.

The results are summarized in the following proposition

Proposition 1. *A necessary condition for truthtelling to be incentive compatible is that $\beta_t \geq \lambda$ for all $t \leq \tau$.*

Proof. See Appendix. □

In the next sections, we derive the optimal incentive compatible contract from the principal's problem. In our framework, choosing the optimal incentive compatible contract amounts to choosing the optimal sensitivity of the agent's continuation value to the reported output, $\{\beta_t : 0 \leq t \leq \tau\}$, and the optimal intertemporal allocation of the payments to the agent $\{w_t : 0 \leq t \leq \tau\}$.[7] Thereby the sensitivity process has to satisfy the incentive compatibility condition $\beta_t \geq \lambda$ for all $t \leq \tau$, and the payment process has to satisfy the agent's limited liability constraint $w_t \geq 0$ for all $t \leq \tau$.

4 First-Best Solution

In the optimal contract, the principal has two ways of rewarding the agent with income: she can either deliver him immediate payments, or increase his continuation value for the future. By the definition of the liquidation policy (8), a higher continuation value U_t implies a later stopping time τ. In our framework, the firm value is persistent, and this choice determines the efficiency of the firm operations. A later stopping time translates to a liquidation of a less profitable firm.

[7]Note that by the definition of τ in (8), together with the representation in Lemma 1, the two choices imply a liquidation policy.

It turns out that once the contract has accumulated a sufficiently high value to the agent, the firm can be operated efficiently. In that case, the liquidation policy (8) reaches the first-best solution. As we will see in the next section, the optimal contract delays payments to the agent until his continuation value is high enough. It remains to determine the first-best optimal liquidation policy and the agent's continuation value that is needed to reach the first-best solution.

At time t, the expected joint profit is the sum of the principal's and the agent's value

$$E\left[\int_t^\tau re^{-r(s-t)}(x_s - w_s)ds + e^{-r(\tau-t)}L\right] + U_t$$
$$= E\left[\int_t^\tau re^{-r(s-t)}x_s ds + e^{-r(\tau-t)}L\right]. \quad (9)$$

The first-best value is the solution of a standard real option problem.[8] The objective is to choose an optimal stopping time τ to maximize (9).

The optimal liquidation policy is Markov in the cash-flow x_t. Since the cash-flow follows a Markov process, the current output reflects future expectations, and is a sufficient statistic for the firm value. The higher the cash-flow is, the more valuable the firm. If the cash-flow falls too low, $x_t < x_L$, the liquidation value becomes relatively more valuable, and the firm is liquidated. The firm is operated so long as the cash-flow stays above the threshold, $x_t \geq x_L$, and it is liquidated as soon as the cash-flow reaches x_L.

From the real option literature we know that the firm is liquidated later than is myopically optimal. The result reflects the option value of liquidation. If the firm is operated further, a positive shock may occur that increases the firm value. If a negative shock occurs, the firm is liquidated. Liquidation, in turn, is irreversible. Indeed, it turns out that the firm is liquidated if the cash-flow reaches

$$x_L = L - (2r)^{-1/2}.$$

The first-best solution is only feasible if the agent's continuation value is high enough. If the agent's value is high enough, the limited liability constraint does not bind. The principal is able to punish the agent for reporting lower outcomes by lowering his continuation value until the first-best optimal stopping time is reached.

When implementing the first-best solution, it is without loss of generality to reward the agent with immediate income. Since both players are risk neutral and discount the future by the common rate r, the intertemporal allocation of wealth is irrelevant for efficiency. The result holds as long as the agent's limited liability constraint is not violated.

[8] See, for example, Dixit and Pindyck (1994).

We determine the agent's smallest possible continuation value $U^{FB}(x_t)$ for which the first-best solution is attained. From the analysis of Section 3 we know that the agent's continuation value has to decrease by at least λ for each dollar that he reports. Next, to reach the first-best solution, the agent's continuation value U_t has to reach 0 at the same time that x_t reaches the efficient liquidation threshold x_L. It turns out that if the agent's continuation value is at least

$$U^{FB}(x_t) = \lambda(x_t - x_L),$$

the optimal contract avoids inefficient liquidation and attains the first-best solution.

The intuition behind the result is straightforward. Notice that besides the limited liability condition, $\hat{x}_t \leq x_t$, we do not impose any restriction on the agent's strategy l_t. In particular, the agent always has the opportunity to react instantaneously, and return to the equilibrium path. Therefore, the agent's deviation does not impose him any *additional* risk of termination at some level $x_\tau > \hat{x}_\tau = x_L$.

Next, the agent's liquidation payoff is 0, whereas continuation has nonnegative value. Therefore, the agent never diverts cash-flows such that the firm is liquidated at a level $x_\tau > \hat{x}_\tau = x_L$. As a consequence, the principal only has to compensate the agent for reporting a cash-flow of $x_t \geq x_L$. Furthermore, we know from Proposition 1 that the agent has to earn at least λ for each dollar he reports in the continuation region.

The optimal contract can be summarized in the following proposition

Proposition 2. *Under the efficient solution, the firm is operated as long as* $x_t > x_L$*, and is liquidated as soon as* x_t *reaches* x_L*, where*

$$x_L = L - (2r)^{-1/2}. \tag{10}$$

The principal's payoff is $v^{FB}(x) = s(x) - U^{FB}(x)$*, where*

$$s(x) = x - x_L e^{-\sqrt{2r}(x-x_L)} + L e^{-\sqrt{2r}(x-x_L)} \tag{11}$$

$$U^{FB}(x) = \lambda(x - x_L) \tag{12}$$

if $x \geq x_L$*, and* $s(x) = L$ *and* $U(x) = 0$ *otherwise.*

Proof. See Appendix. □

5 Optimal Contract

In this section, we derive the optimal contract heuristically. We show that the principal's optimality requires that whenever the agent's report does not trigger liquidation, the incentive constraints bind. That is, the principal leaves the agent the smallest possible compensation that is compatible with truthful reporting.

Moreover, if the agent's continuation value U_t reaches 0, the limited liability constraint becomes binding. As discussed in Section 3, the only way to provide the agent incentives is to liquidate the firm. Notice that the incentive constraints *do not bind* at the cash-flow levels that are not reached before U_t reaches 0.[9]

The optimal contract relies on an inefficient termination threat to provide the agent incentives in the beginning of the relationship. The payments to the agent are delayed and he is rewarded by promising him a higher continuation value. With good enough past performance, the contract has accumulated a high enough value to the agent. Then the first-best solution is implemented.

Using the results obtained in Section 3, we can rewrite the optimal contracting problem (3)-(5) as

$$v(U,x) = \max_{w,\beta} E\left[\int_0^\tau re^{-rt}(x_t - w_t)dt + e^{-r\tau}L\right],$$

subject to the law of motion of the state variables x_t and U_t in (1) and (7). Thereby the sensitivity β of the agent's continuation value to the Brownian motion has to satisfy the incentive compatibility constraints, $\beta_t \geq \lambda$ for all $t \leq \tau$, and the payment process w the nonnegativity constraints, $w_t \geq 0$ for all $t \leq \tau$.

The principal's optimal choice of $\{w, \beta\}$ at any point can be derived from her Hamilton-Jacobi-Bellman equation

$$rv(U,x) = \max_{w\geq 0,\beta\geq\lambda} \Big\{ r(x-w) + r(U-w)v_U(U,x) + \frac{\beta^2}{2}v_{UU}(U,x) + \beta v_{Ux}(U,x) + \frac{1}{2}v_{xx}(U,x) \Big\}. \quad (13)$$

The boundary condition is $v(0,x) = L$.[10]

First, we discuss the principal's optimal choice of the agent's income w. The principal can either reward the agent with immediate payments, or increase his continuation value for the future. The marginal cost of providing immediate income to the risk-neutral agent is -1, whereas the cost of increasing the agent's continuation value can be higher or lower. The opportunity to provide immediate income ensures that at the optimum, the marginal cost of providing incentives can never exceed the marginal cost of providing immediate income.

[9] In general, we cannot assume that the incentive constraints bind, and indeed, they do not bind everywhere. See Rochet and Chone (1998) for a related discussion in a multidimensional screening setting.

[10] Notice that the boundary condition is only determined in terms of the state variable U. The boundary value of the state variable x is free and must be determined as part of the solution.

Indeed, by optimizing (13) with respect to w, we find that if

$$v_U(U, x) \geq -1, \tag{14}$$

it is optimal to set $w = 0$. In the beginning of the relationship, the payments to the agent are delayed, and the principal rewards the agent by promising him a higher continuation value for the future. The optimal contract provides the agent incentives by threatening to liquidate the firm inefficiently following a period of low cash-flow realizations.

The marginal cost of delaying payments depends on the agent's continuation value relative to the cash-flow level. If the cash-flow is high, but the continuation value is low, the risk of inefficient termination is high. Delaying payments increases the agent's continuation value, which protects the firm against liquidation. If $U_t < U^{FB}(x_t)$, the principal can gain by increasing the agent's continuation value. The payments to the agent are delayed until his continuation value reaches $U^{FB}(x_t)$, and the first-best solution is attainable. Thereafter $U_t = U^{FB}(x_t)$, the firm is liquidated at the efficient level, and the agent is rewarded with immediate income.

Next, we discuss the principal's optimal choice of the sensitivity process β. To show that the incentive constraints bind in the continuation region, we adopt the following approach, developed in DeMarzo and Sannikov (2011). We first conjecture that, at the optimum, $\beta = \lambda$. Then we compare the principal's profit from this contract with her profit from any other incentive compatible contract with the sensitivity $\beta \geq \lambda$. We show that the contract with $\beta = \lambda$ attains the highest feasible profit for the principal. Therefore, the principal optimally lets the agent's incentive compatibility constraints bind for any cash-flow report that does not trigger liquidation.

For $\beta = \lambda$, the principal's Hamilton-Jacobi-Bellman equation (13) can be written as

$$\begin{aligned} rv(U, x) = \max_{w \geq 0} \Big\{ & r(x - w) + r(U - w) v_U(U, x) \\ & + \frac{\lambda^2}{2} v_{UU}(U, x) + \lambda v_{Ux}(U, x) + \frac{1}{2} v_{xx}(U, x) \Big\}. \end{aligned} \tag{15}$$

By comparing (13) and (15), we find that $\beta = \lambda$ is optimal for the principal only if

$$\frac{1}{2}(\beta - \lambda)^2 v_{UU}(U, x) + (\beta - \lambda)(\lambda v_{UU}(U, x) + v_{Ux}(U, x)) \leq 0. \tag{16}$$

First, as we can see from (7), β determines the volatility of the agent's continuation value U_t. If U_t varies more, the risk that it reaches 0 increases. As in the framework with uncorrelated cash-flows, excess volatility of the agent's continuation value increases the risk of unnecessarily early liquidation. Increasing the probability of inefficient liquidation is costly for the principal. Therefore,

$$(\beta - \lambda)^2 v_{UU}(U, x) \leq 0, \tag{17}$$

which is maximized at $\beta = \lambda$. The principal optimally imposes the firm with the minimal inefficient termination risk that is necessary to sustain incentives.

Next, the shocks are persistent in our framework. Increasing β today increases the agent's share of the total cash-flow in the future. Conversely, the principal's share of the cash-flow decreases. In the future, the agent has to receive a compensation of λ for each dollar that he reports. Formally,

$$(\beta - \lambda)(\lambda v_{UU}(U, x) + v_{Ux}(U, x)) \leq 0. \tag{18}$$

At the optimum, the principal only increases the agent's share of the cash-flow so much that the incentive constraint just binds. That is, (18) is maximized at $\beta = \lambda$. Finally, (16) together with (17) and (18) imply that the principal optimally sets the sensitivity at its lowest admissible level $\beta = \lambda$.

The intuition behind the results is straightforward. In the beginning of the relationship, the principal can extract rents from the agent by threatening to liquidate the firm inefficiently if he reports low cash-flow returns. To minimize the risk of inefficient termination, the principal sets the variance of the agent's continuation value at the lowest admissible level. The result is consistent with the results obtained earlier in the framework with independently distributed cash-flows.

Next, the principal can profit by delaying payments to the agent. Instead of rewarding the agent with immediate payments, she can promise him a higher share of the returns in the future. As is standard in the principal-agent framework, the firm can be operated more efficiently if the agent owns a higher share of it. Then his income varies more with the fluctuations of the output, and he cares more about efficiency. This relaxes the incentive constraint. Increasing the agent share is profitable since the firm value increases towards efficiency. The principal is able to extract all the surplus from the increase from the agent.

The results are summarized in the following proposition

Proposition 3. *Starting from the initial cash-flow x_0, and the agent's initial value $U_0 \in [0, U^{FB}(x_0)]$, the optimal contract attains the profit $v(U_0, x_0)$ for the principal. The players' values evolve stochastically in response to the fluctuations of the output, and admit the following dynamics:*

1. *When the agent's continuation value U_t is on the interval $(0, U^{FB}(x))$, it evolves according to*

$$dU_t = rU_t dt + \lambda dx_t. \tag{19}$$

The payments to the agent are delayed, and his flow payment is set to $w_t = 0$. *The principal's expected profit at any point is* $v(U, x)$, *which is the unique solution of the following partial differential equation*

$$rv(U,x) = x + rUv_U(U,x) + \frac{\lambda^2}{2}v_{UU}(U,x) + \lambda v_{Ux}(U,x) + \frac{1}{2}v_{xx}(U,x) \quad (20)$$

with the boundary conditions $v_U(U^{FB}(x), x) = -1$, $v(U^{FB}(x), x) = v^{FB}(x)$, *and* $v(0, x) = L$.

2. *When* U_t *reaches* $U^{FB}(x_t)$, *the first-best solution is implemented. The agent receives a flow payment* $w_t = \lambda(x_t - x_L)$, *where* $x_L = L - (2r)^{-1/2}$. *At any point, the principal's expected payoff is* $v(x) = x - U^{FB}(x) - x_L e^{-\sqrt{2r}(x - x_L)} + Le^{-\sqrt{2r}(x-x_L)}$, *where* $U^{FB}(x) = \lambda(x - x_L)$.

3. *When* U_t *reaches* 0, *the contract is terminated. The players receive their liquidation payoffs* $v(0, x) = L$, *and* $U = 0$.

5.1 Full Incentive Compatibility

In this section, we verify heuristically that the agent's global incentive compatibility constraint is satisfied at the optimal contract. The formal proof follows by showing that the if $\beta_t = \lambda$ for all $t \leq \tau$, the agent's global incentive compatibility is satisfied at the optimum, and is delegated to the Appendix. The result guarantees that the agent has the right incentives to report truthfully starting from an arbitrary history, possibly off the equilibrium path.

Before τ is reached, the agent's incentives to report truthfully are independent of his past actions. Since the agent is risk-neutral and the private benefit is proportional to the cash-flow, the agent's incentives to divert an additional dollar are independent of the cash-flow level. Therefore, the conditions that guarantee that the agent reports truthfully at equilibrium, guarantee that the agent cannot gain from additional deviations off the equilibrium path.

Starting from a history that is on the equilibrium path, the agent's continuation value increases by $\beta_t = \lambda$ for each additional dollar that he reports. The increase compensates him for his private benefit of λ that he would have received if he would have concealed the cash-flow. Similarly, we can consider the agent's incentive to divert starting from a history that is off the equilibrium path. Again, reporting an additional dollar increases the agent's continuation value by $\beta_t = \lambda$, which exactly offsets the value of concealing the cash-flow. The agent's gain of concealing an additional dollar, λ, is the same both on and off the equilibrium path.

Next, notice it is never optimal for the agent to choose a strategy such that the reported cash-flow $\hat{x}_t$ and the actual output x_t do not agree at $t = \tau$. The result follows from the fact that the agent's liquidation value is 0 whereas the contract has nonnegative value to the agent.

The result is summarized in the following proposition

Proposition 4. *The optimal contract with $\beta_t = \lambda$ for all $t \leq \tau$ satisfies the agent's global incentive compatibility constraint.*

Proof. See Appendix. □

5.2 Initialization of the Contract

In this section, we discuss how the contract is initialized. That is, we discuss the principal's optimal choice of the agent's initial value U_0 in (4), given that the initial cash-flow is x_0. We call the time $t = 0$ the contracting stage. We assume that the principal makes a take-it-or-leave-it offer to the agent.

The principal's optimal choice of U_0 is a standard constrained optimization problem. At $t = 0$, the principal solves

$$\max_{U_0 \geq 0} v(U_0, x_0) \tag{21}$$

subject to the players' participation constraints

$$U_0 \geq \underline{U} \tag{22}$$

and

$$v(U_0, x_0) \geq \underline{v}, \tag{23}$$

where $\underline{v}$ and $\underline{U}$ denote the principal's and the agent's reservation utilities at the contracting stage. We can examine different divisions of the players' bargaining power, or allow for investment cost, by varying the players' reservation utilities.

To focus attention on the interesting cases, we assume that $x_0 > x_L$, and $\underline{v} + \underline{U} < s(x_0)$, where $s(x_0)$ is the first-best optimal value of the firm at time 0, as defined in Proposition 2. Under these conditions, contracting is always efficient at time 0.[11]

If $\underline{U} \geq U^{FB}(x_0)$, the first-best solution can be attained at time 0. The agent receives an upfront payment of size $\underline{U} - U^{FB}(x_0)$ as a compensation for participating in the contract. If $\underline{U} < U^{FB}(x_0)$, the firm is eventually liquidated inefficiently in the beginning of the relationship.

Next, suppose that the optimal solution is such that the constraint (22) does not bind. Let U_0^* denote the solution of (21). Then the maximal pledgeable

[11] Of course, contracting may become inefficient after time 0 if x_t reaches x_L.

income is $v(U_0^*, x_0)$. The contract is initialized only if the principal's participation constraint (23) is satisfied at the optimum. That is, only if

$$v(U_0^*, x_0) \geq \underline{v}. \tag{24}$$

If the condition (24) fails, the contract is not initialized even if it would be efficient to do so. The moral hazard problem is so severe that it prevents contracting altogether.

6 Conclusions

This paper examines a real option problem in which the investor needs to employ an agent to handle the daily operations of the firm she is investing in. To ensure that the agent runs the firm efficiently, the investor has to reward him for delivering high returns and punish him if the output falls. To provide the agent incentives, the principal has the opportunity to deliver him nonnegative payments, or to liquidate the firm and receive a liquidation value. In the beginning of the relationship, the players write a contract that determines the optimal incentive scheme. We assume that the players can fully commit to the contract.

We find that the optimal contract is of the following form. In the beginning of the relationship, the payments to the agent are delayed, and the contract relies on liquidation threat to sustain incentive compatibility. The firm is eventually liquidated inefficiently to save on the incentive costs for the principal. Payments to the agent are delayed until his continuation value becomes sufficiently high. Then the first-best solution is implemented.

Our results deliver interesting insights about the connection between moral hazard and efficiency of the firm liquidation policy. In particular, we find that the firm is liquidated more efficiently if the agent's share of the cash-flow is high. If the agent's share is high, he receives a higher compensation for his report, which increases his incentives for truthful reporting. Conversely, if the agent's share is low, the liquidation threat has to be severe. Otherwise the agent will engage in inefficient activities to gain private benefits.

7 Appendix

Proof of Lemma 1. First, we prove that U_t admits the representation (7). Consider the process

$$A_t = \int_0^{t\wedge\tau} re^{-rs}w_s ds + e^{-rt\wedge\tau}U_{t\wedge\tau}. \tag{25}$$

From the principal's point of view, the process A_t is a martingale. By the Martingale Representation Theorem, there exists a progressively measurable process β in L^2 such that for $t \leq \tau$,

$$dA_t = e^{-rt}\beta_t d\hat{x}_t. \tag{26}$$

Applying Itô's lemma on (25), we can write

$$dA_t = re^{-rt}w_t dt - re^{-rt}U_t + e^{-rt}dU_t = e^{-rt}\beta_t d\hat{x}_t,$$

where the last equality follows from (26). By reorganizing, and multiplying by e^{rt}, we obtain (7).

To prove that the transversality condition holds, notice that since $U_\tau = 0$,[12]

$$\int_0^{t\wedge\tau} re^{-rs}w_s ds \to \int_0^{\tau} re^{-rs}w_s ds$$

as $t \to \infty$. Then, since $0 \leq U_t < \infty$, $0 \leq \int_0^\tau re^{-rs}w_s ds < \infty$. Hence by the Dominated Convergence Theorem

$$E[1_{s\leq\tau}e^{-rt}U_t] = E\left[\int_0^{\tau} re^{-rs}w_s\right] - E\left[\int_0^{t\wedge\tau} w_s ds\right] \to 0$$

as $t \to \infty$. The argument that $\lim_{s\to\infty} E_t[1_{s\leq\tau}e^{-rs}U_{t+s}] \to 0$ follows similarly.

Finally, suppose that U_t satisfies the conditions (i) and (ii). Then from (26), we can see that the process A_t is a martingale. Therefore,

$$U_0 = A_0 = E[A_t] = E\left[\int_0^{t\wedge\tau} re^{-rs}w_s ds\right] + e^{-rt\wedge\tau}U_{t\wedge\tau} \underset{t\to\infty}{\to} E\left[\int_0^{\tau} re^{-rs}w_s\right].$$

since $U_\tau = 0$, and the transversality condition holds. The argument for $t > 0$ follows similarly. □

Proof of Proposition 1. Recall from the Section 2 that the contract $\{w, \tau\}$ is incentive compatible if and only if for all t and for all strategies such that $\hat{x}_t = x_t$ after t,

$$E\left[\int_0^{\tau} re^{-rs}w_s ds\right] \geq \hat{E}\left[\int_0^{\tau} re^{-rs}(w_s + \lambda(x_s - \hat{x}_s))ds\right]. \tag{27}$$

[12] The results can easily be extended to the case in which $U_\tau > 0$. See, for example, DeMarzo and Sannikov (2011).

The operator $\hat{E}$ denotes the agent's expectation under the strategy $\hat{x}$.

To derive necessary conditions for incentive compatibility, we consider a particular kind of deviation. Under the deviations strategy, the agent departs from the equilibrium strategy at time 0 and returns to truthful reporting at $\Delta + \Delta^2$.[13] Notice that at $\Delta + \Delta^2$, the process (7) again describes the agent's continuation value. Therefore, we can write (27) as

$$U_0 \geq \hat{E}\left[\int_0^{\Delta+\Delta^2} re^{-rs}(w_s + \lambda(x_s - \hat{x}_s))ds + e^{-r(\Delta+\Delta^2)}U_{\Delta+\Delta^2}\right]. \tag{28}$$

Next, notice that

$$\begin{aligned} d(e^{-rt}U_t) &= -re^{-rt}U_t dt + e^{-rt}dU_t \\ &= -re^{-rt}U_t dt + re^{-rt}(U_t - w_t)dt + e^{-rt}\beta_t d\hat{x}_t, \end{aligned} \tag{29}$$

where the first equality follows by Itô's Lemma, and the second by (7). Integrating (29) from 0 to t, we obtain

$$e^{-rt}U_t = U_0 - \int_0^t re^{-rs}w_s ds + \int_0^t e^{-rs}\beta_s d\hat{x}_s. \tag{30}$$

Using (30) with $t = \Delta + \Delta^2$, we can rewrite (28) as

$$\begin{aligned} U_0 \geq\; & U_0 + \hat{E}\left[\int_0^{\Delta+\Delta^2} re^{-rs}(w_s + \lambda(x_s - \hat{x}_s))ds\right] \\ & - \hat{E}\Bigg[\int_0^{\Delta+\Delta^2} re^{-rs}w_s ds - \int_0^{\Delta+\Delta^2} e^{-rs}\beta_s \underbrace{d\hat{x}_s}_{=dx_s - l_s ds}\Bigg] \\ =\; & U_0 + \hat{E}\left[\int_0^{\Delta+\Delta^2} re^{-rs}\lambda\int_0^s l_u du - e^{-rs}\beta_s l_s ds\right] \end{aligned} \tag{31}$$

The last equality uses the fact that from the agent's point of view, x_t is a standard Brownian motion, and therefore,

$$\hat{E}\left[\int_0^{\Delta+\Delta^2} e^{-rs}\beta_s dx_s\right] = 0.$$

Next, we consider the following deviation strategy. At the time interval $s \in [0, \Delta]$, the agent departs from the equilibrium path and diverts cash-flows at rate

[13] A similar argument shows the necessary conditions for incentive compatibility at $t > 0$.

l. On the time interval $s \in [\Delta, \Delta + \Delta^2]$, he returns to the equilibrium path by choosing a drift $-l/\Delta$. Notice that

$$x_{\Delta+\Delta^2} - \hat{x}_{\Delta+\Delta^2} = \int_0^{\Delta} l ds + \int_{\Delta}^{\Delta+\Delta^2} \frac{l}{\Delta} ds = 0$$

such that the agent has returned to the equilibrium path at time $s = \Delta + \Delta^2$. The principal chooses a sensitivity of β.

The agent's deviation strategy can be illustrated in the following phase diagram

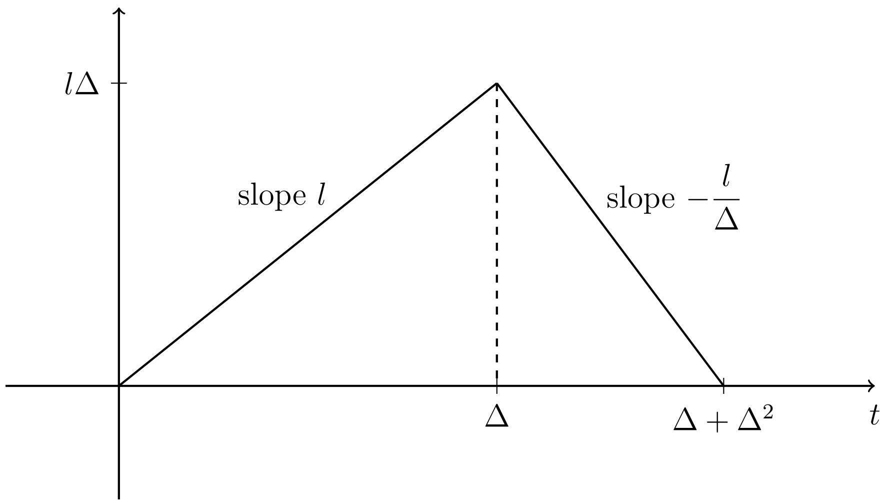

Using the particular deviation strategy, we can rewrite (31) as

$$\begin{aligned}
&\hat{E}\left[\int_0^{\Delta} re^{-rs}\lambda \int_0^{s} l du - e^{-rs}\beta l ds\right] - \hat{E}\left[\int_{\Delta}^{\Delta+\Delta^2} re^{-rs}\lambda \int_{\Delta}^{s} \frac{l}{\Delta} du - e^{-rs}\beta \frac{l}{\Delta} ds\right] \\
=& \int_0^{\Delta} re^{-rs}\lambda s l - e^{-rs}\beta l ds - \int_{\Delta}^{\Delta+\Delta^2} re^{-rs}\frac{s-\Delta}{\Delta} - e^{-rs}\beta \frac{l}{\Delta} ds \\
=& -\lambda l e^{-rs} s\Big|_0^{\Delta} + \lambda l \int_0^{\Delta} e^{-rs} ds + \beta l \left.\frac{e^{-rs}}{r}\right|_0^{\Delta} + \lambda l \left.\frac{e^{-rs} s}{\Delta}\right|_{\Delta}^{\Delta+\Delta^2} - \lambda l \int_{\Delta}^{\Delta+\Delta^2} \frac{e^{-rs}}{\Delta} ds \\
& - \lambda l \, e^{-rs}\Big|_{\Delta}^{\Delta+\Delta^2} - \beta l \left.\frac{e^{-rs}}{r\Delta}\right|_{\Delta}^{\Delta+\Delta^2} \\
=& -\lambda l \Delta e^{-r\Delta} - \lambda l \left.\frac{e^{-rs}}{r}\right|_0^{\Delta} - \beta l \frac{1-e^{-r\Delta}}{r} + \lambda l (1+\Delta) e^{-r(\Delta+\Delta^2)} - \lambda l e^{-r\Delta} \\
& - \lambda l e^{-r(\Delta+\Delta^2)} + \lambda l e^{-r\Delta} + \lambda l \left.\frac{e^{-rs}}{r\Delta}\right|_{\Delta}^{\Delta+\Delta^2} - \beta l \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta} \\
=& -\lambda l \Delta e^{-r\Delta} + \lambda l \frac{1-e^{-r\Delta}}{r} - \beta l \frac{1-e^{-r\Delta}}{r} + \lambda l \Delta e^{-r(\Delta+\Delta^2)}
\end{aligned}$$

$$+ \lambda l \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta} - \beta l \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta}$$
$$= -(\beta - \lambda) l \left(\frac{1 - e^{-r\Delta}}{r} + \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta} \right) + \lambda l (e^{-r(\Delta+\Delta^2)} - e^{-r\Delta})\Delta \leq 0.$$

Dividing both sides by Δ^2 yields

$$-(\beta - \lambda) l \left(\frac{1 - e^{-r\Delta}}{r\Delta^2} + \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta^3} \right) + \lambda l \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{\Delta} \leq 0.$$

Taking the limits as $\Delta \to 0$, we find that the condition becomes

$$-\frac{r}{2}(\beta - \lambda) l \leq 0.$$

The result follows since by L'Hôpital's rule

$$\lim_{\Delta \to 0} \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{\Delta} = \lim_{\Delta \to 0} \frac{-r(1+2\Delta)e^{-r(\Delta+\Delta^2)} + re^{-r\Delta}}{1} = 0$$

such that the last term vanishes. Moreover,

$$\lim_{\Delta \to 0} \frac{1 - e^{-r\Delta}}{r\Delta^2} + \frac{e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta^3} = \lim_{\Delta \to 0} \frac{\Delta(1 - e^{-r\Delta}) + e^{-r(\Delta+\Delta^2)} - e^{-r\Delta}}{r\Delta^3}$$
$$= \lim_{\Delta \to 0} \frac{1 - e^{-r\Delta} + r\Delta e^{-r\Delta} - r(1+2\Delta)e^{-r(\Delta+\Delta^2)} + re^{-r\Delta}}{3r\Delta^2}$$
$$= \lim_{\Delta \to 0} \frac{2re^{-r\Delta} - r^2\Delta e^{-r\Delta} - 2re^{-r(\Delta+\Delta^2)} + r^2(1+2\Delta)^2 e^{-r(\Delta+\Delta^2)} - r^2 e^{-r\Delta}}{6r\Delta}$$
$$= \lim_{\Delta \to 0} \frac{1}{6r}(-3r^2 e^{-r\Delta} + r^3\Delta e^{-r\Delta} + 6r^2(1+2\Delta)e^{-r(\Delta+\Delta^2)}$$
$$- r^3(1+2\Delta)^3 e^{-r(\Delta+\Delta^2)} + r^3 e^{-r\Delta})$$
$$= \frac{1}{6r}(-3r^2 + 6r^2) = \frac{r}{2}.$$

□

Proof of Proposition 2. From the real option literature, we know that the first-best value (11) is the unique solution of the ordinary differential equation

$$rs(x) = rx + \frac{1}{2}s_{xx}(x) \tag{32}$$

with the boundary conditions $s(x_L) = L$ and $s_x(x_L) = 0$.

We argue that given that the agent earns the value $U^{FB}(x)$, the maximal attainable value for the principal is $v^{FB}(x)$. Towards this end, consider the process

$$S_t = \int_0^t e^{-rs} r x_s ds + e^{-rt} s(x_t).$$

We show that S_t is a supermartingale. Using Itô's lemma, we find that

$$e^{rt}dS_t = rx_t dt - rs(x_t)dt + \frac{1}{2}s_{xx}(x_t)dt + s_x(x_t)dx_t.$$

By (32), S_t is a martingale when $x_t \geq x_L$ and a supermartingale if $x_t < x_L$. The principal's profit at time 0 satisfies

$$E\left[\int_0^\tau re^{-rt}(x_t - w_t)dt + e^{-r\tau}L\right] = E\left[\int_0^\tau re^{rt}x_t dt + e^{-r\tau}L\right] - U_0$$
$$\leq E[S_\tau] - U_0 \leq S_0 - U_0 = s(x_0) - U_0$$

with equality only if the principal chooses the optimal stopping time such that τ is reached as x_t reaches x_L.

Next, we derive the lower bound $U^{FB}(x)$ of the agent's continuation value that is required to implement the first-best solution. Since both players are risk-neutral and discount the future at a common rate, it is without loss of generality to concentrate on contracts in which the agent is compensated with immediate income. Furthermore, the agent's continuation value reaches 0 in the same time that x_t reaches x_L.

Since the agent is rewarded with immediate payments, his flow income is $rw_t dt = rU_t dt$. Substituting in (7) with $\beta_t = \lambda$, we find that the agent's continuation value U_t solves

$$dU_t = \lambda dx_t.$$

Solving the stochastic differential equation with the boundary conditions such that U_t hits 0 at the same time that x_t hits x_L, we find that the agent's value at the point $x_t = x$ is given by (12). □

Next, we verify that the contract conjectured in Section 5 is indeed optimal. The proof proceeds as in DeMarzo and Sannikov (2011).

Before proving the dynamics of the optimal payment schedule, we examine some key properties of the principal's equilibrium value function. We then use the properties to prove the results in Proposition 3.

We show that the principal's value function satisfies the following conditions

$$v_{UU}(U, x) \leq 0,$$

and

$$\lambda v_{UU}(U, x) + v_{Ux}(U, x) \leq 0.$$

for all (U, x). Moreover, since $v_{UU}(U_t, x_t) \leq 0$ and $v_U(U_t, x_t) = -1$ if $U_t \geq U^{FB}(x_t)$, $v_U(U_t, x_t) \geq -1$ for all U_t.[14]

[14]The result can be seen graphically: Since $v_U(U_t, x_t)$ is (weakly) decreasing, and we know that it is -1 from $U_t \geq U^{FB}(x_t)$ onwards, we must have $v_U(U_t, x_t) \geq -1$. Otherwise $v_U(U_t, x_t)$ would not be decreasing.

The state variable U_t evolves according to (19) so long as $U_t \leq U^{FB}(x_t)$, and takes the value $U_t = U^{FB}(x_t)$ as soon as it reaches $U^{FB}(x_t)$. The value of the state variable x_t is given in (1).

To understand the idea behind the proof, notice that the principal's marginal value $v_U(U_t, x_t)$ changes as the state variables U_t and x_t evolve in two directions: in the direction of volatilities and in the direction of drifts. Both the agent's continuation value and the cash-flow change in the direction of volatilities, in which $dU/dx = \lambda$. For each unit that the cash-flow increases, the agent's continuation value increases by λ. In particular, both processes are driven by the same Brownian motion.

Moreover, the agent's continuation value increases in the direction of drifts. Intuitively speaking, the optimal contract accumulates value to the agent to protect the firm from an inefficient liquidation. That is, whenever the payments to the agent are delayed, his continuation value grows at the rate $rU_t dt$. Notice that if the agent's continuation value is higher, it grows at a (weakly) higher rate.

We need to show that the principal's value function is (weakly) concave in both of the directions. The different directions can be illustrated in a phase diagram.

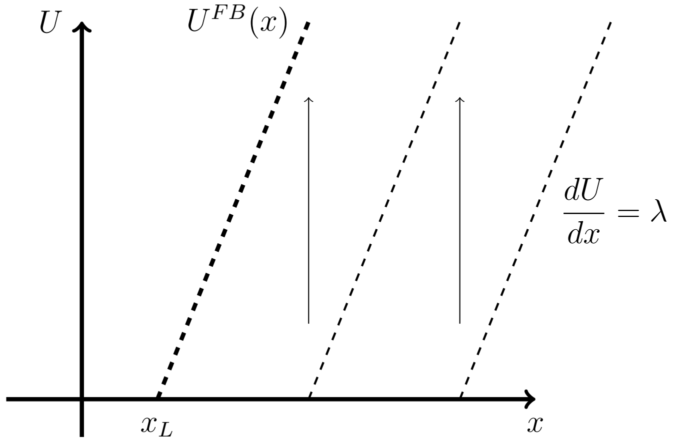

Lemma 2. *Given that the state of the world (U_t, x_t) evolves according to (1) and (19), $v_U(U_t, x_t)$ is a martingale.*

Proof. Differentiating (20) with respect to U, we find that

$$rv_U(U,x) = rv_U(U,x) + rUv_{UU}(U,x) + \frac{\lambda^2}{2}v_{UUU}(U,x) + \lambda v_{UUx}(U,x)$$

or,

$$0 = rUv_{UU}(U,x) + \frac{\lambda^2}{2}v_{UUU}(U,x) + \lambda v_{UUx}(U,x).$$

Applying Itô's lemma with (19) on $v_U(U_t, x_t)$, we can see that the right hand side corresponds to its drift when $U_t \leq U^{FB}(x_t)$.

Furthermore, if $U_t \geq U^{FB}(x_t)$, $v_U(U_t, x_t) = -1$. By combining the results, we conclude that $v_U(U_t, x_t)$ is a martingale. □

The next lemma shows that the marginal value of continuing is always weakly higher if the cash-flows are terminated at a higher level. The result follows from the fact that the liquidation value is constant while the cash-flow is higher if the current cash-flow is higher. Therefore, the marginal gain from continuation is higher for higher levels of x_t.

Lemma 3. *$v_U(0, x)$ weakly increases in x.*

Proof. Consider two processes $(U_t^i, x_t^i)_{s \geq t}$, starting from the values $U_0^1 = U_0^2 = \varepsilon$, and $x_0^1 = x_0^2 + \delta$, for some $\varepsilon > 0$ and $\delta > 0$.

Since $U_0^1 = U_0^2$, we have by (19) for any path of the Brownian motion that

$$U_t^1 - U_t^2 = \int_0^t r\underbrace{(U_s^1 - U_s^2)}_{=0}ds + \int_0^t \lambda\underbrace{(dx_s - dx_s)}_{=0} = 0$$

and, therefore, $U_t^1 = U_t^2$. Therefore, $\tau(U^1) = \tau(U^2) = \tau$. Moreover, by (1), $x_t^1 - x_t^2 = \delta$. Substituting in the principal's value function, we find that

$$\begin{aligned} v(x_0^1, \varepsilon) - v(x_0^2, \varepsilon) &= E\left[\int_0^\tau re^{-rt}(x_t^1 - x_t^2)dt\right] + E\left[e^{-r\tau}(L - L)\right] - \varepsilon + \varepsilon \\ &= E\left[\int_0^\tau re^{-r\tau}\delta dt\right] \geq 0, \end{aligned}$$

which implies further that

$$\begin{aligned} & v\left(\varepsilon, x_0^1\right) \geq v\left(\varepsilon, x_0^2\right) \\ \iff & v\left(\varepsilon, x_0^1\right) - L \geq v\left(\varepsilon, x_0^2\right) - L \\ \iff & v\left(\varepsilon, x_0^1\right) - v\left(0, x_0^1\right) \geq v\left(\varepsilon, x_0^2\right) - v\left(0, x_0^2\right). \end{aligned}$$

The result follows since δ and ε were chosen arbitrarily. □

The next lemma implies that $v_{UU}(U, x) \leq 0$. That is, the principal's value is concave in U.

Lemma 4. *$v_U(U, x)$ weakly decreases in U.*

Proof. We show that for any x_0, and for any two values $U_0^1 > U_0^2$, $v_U(U_t^1, x_t) \leq v_U(U_t^2, x_t)$.

Consider two processes $(U_s^i, x_s)_{s\geq 0}$, $i = 1, 2$, starting from the values $U_0^1 = U_0^2 + \delta$ and $x_0^i = x_0$. That is, both processes start from the same initial cash-flow, but the first process starts at a higher promised value to the agent. Let $\tau^1 \equiv \tau(U^1)$ and $\tau(U^2) \equiv \tau^2$ denote the first time at which each process reaches 0.

Since $U_0^1 > U_0^2$, it holds by (19) for any path of the Brownian motion

$$U_t^1 - U_t^2 = \int_0^t r\underbrace{(U_s^1 - U_s^2)}_{\geq 0} ds + \int_0^t \lambda \underbrace{(dx_s - dx_s)}_{=0} \geq 0$$

Therefore, for any path of the Brownian motion, the process U^1 never reaches 0 earlier than the process U^2. This implies that $\tau^1 \geq \tau^2$, and, therefore, $x_{\tau^1}^1 \leq x_{\tau^2}^1$.

Furthermore, by (1), $x_{\tau^2}^1 = x_{\tau^2}^2$. Combining the two observations, we find that $x_{\tau^1}^1 \leq x_{\tau^2}^2$.

Using Lemmas 2 and 3, we obtain

$$v_U(U_0^1, x_0) = E\left[v_U(0, x_{\tau^1}^1)\right] \leq E\left[v_U(0, x_{\tau^2}^2)\right] = v_U(U_0^2, x_0). \qquad \square$$

The next lemma proves that $v_U(U, x)$ weakly decreases in the direction of volatilities.

Lemma 5. *$v_U(U, x)$ weakly decreases in the direction in which U and x increase according to $dU/dx = \lambda$.*

Proof. Consider the processes $(U_s^i, x_s^i)_{s\geq 0}$, $i = 1, 2$, that follow (1) and (19) starting from the values that satisfy

$$x_0^1 - x_0^2 = \delta > 0 \text{ and } U_0^1 - U_0^2 = \lambda\delta.$$

Again, let $\tau^1 \equiv \tau(U^1)$ and $\tau^2 \equiv \tau(U^2)$.

Again, we can see from (19) that

$$U_t^1 - U_t^2 = \int_0^t r\underbrace{(U_s^1 - U_s^2)}_{\geq \lambda\delta} ds + \int_0^t \lambda \underbrace{(dx_s - dx_s)}_{=0} \geq 0$$

Hence, for any path of the Brownian motion and $t > 0$,

$$U_t^1 - U_t^2 \geq \lambda\delta.$$

Thus, we find that $\tau^1 > \tau^2$.

Next,

$$\begin{aligned} &U^1_{\tau^2} - \underbrace{U^2_{\tau^2}}_{=0} \geq \lambda\delta \\ \iff &U^1_{\tau^2} - \underbrace{U^1_{\tau^1}}_{=0} \geq \lambda\delta, \end{aligned}$$

or,

$$U^1_{\tau^2} - U^1_{\tau^1} = \int_{\tau^2}^{\tau^1} (rU^1_t + dx_t) \leq -\lambda\delta$$

by rewriting we find that

$$x^1_{\tau^1} - x^1_{\tau^2} = \int_{\tau^2}^{\tau^1} dx_t \leq -\delta - \int_{\tau^2}^{\tau^1} \frac{rU^1_t}{\lambda} dt \leq -\delta$$

from which it follows that $x^1_{\tau^2} - x^1_{\tau^1} \geq \delta$. Moreover, $x^2_{\tau^2} = x^1_{\tau^2} - \delta$ by (1). By combining the results, we can conclude that $x^2_{\tau^2} \geq x^1_{\tau^1}$.

Using Lemmas 2 and 3,

$$v_U(U^1_0, x^1_0) = E\left[v_U(0, x^1_{\tau^1})\right] \leq E\left[v_U(0, x^2_{\tau^2})\right] = v_U(U^2_0, x^2_0). \qquad \square$$

Proof of Proposition 3. Consider the process

$$P_t = e^{-rt} v(U_t, x_t) + \int_0^t re^{-rs}(x_s - w_s) ds. \tag{33}$$

The state variables x_t and U_t evolve according to (1) and (7).

We show that given $\beta_t \geq \lambda$ and $w_t \geq 0$, P_t is a supermartingale. It is a martingale only if $\beta_t = \lambda$, and $w_t = 0$ whenever $v_U(U_t, x_t) > -1$.

Using Itô's lemma on (33), taking the expectations, and multiplying by e^{rt}, we can write

$$\begin{aligned} \frac{e^{rt} E\left[dP_t\right]}{dt} = &\, r(x_t - w_t) - rv(U_t, x_t) + r(U_t - w_t) v_U(U_t, x_t) \\ &+ \frac{\beta_t^2}{2} v_{UU}(U_t, x_t) + \beta_t v_{Ux}(U_t, x_t) + \frac{1}{2} v_{xx}(U_t, x_t). \end{aligned}$$

Adding (20) we find that

$$\begin{aligned} \frac{e^{rt} E\left[dP_t\right]}{dt} = &-r(1 + v_U(U_t, x_t)) w_t \\ &+ \frac{1}{2}\left(\beta_t^2 - \lambda^2\right) v_{UU}(U_t, x_t) + \beta_t \left(\beta_t - \lambda\right) v_{Ux}(U_t, x_t). \end{aligned}$$

or,

$$\frac{e^{rt}E\left[dP_t\right]}{dt} = -r(1+v_U(U_t,x_t))w_t + \frac{1}{2}\left(\beta_t-\lambda\right)^2 v_{UU}(U_t,x_t) + \left(\beta_t-\lambda\right)\left(\lambda v_{UU}(U_t,x_t)+v_{Ux}(U_t,x_t)\right). \quad (34)$$

Since $w_t \geq 0$, $v_U(U,x) \geq -1$ implies that

$$-r(1+v_U(U_t,x_t))w_t \leq 0,$$

with equality only if $w_t = 0$ whenever $v_U(U_t,x_t) > -1$.

Moreover, since

$$v_{UU}(U_t,x_t) \leq 0$$

by Lemma 4, the second term of (34) is nonpositive, and it is 0 only if $\beta_t = \lambda$. Finally, since $\beta_t \geq \lambda$, and

$$\lambda v_{UU}(U_t,x_t) + v_{Ux}(U_t,x_t) \leq 0$$

by Lemma 5, the last term of (34) is always nonpositive, and it is 0 only if $\beta_t = \lambda$. Therefore, P_t is a supermartingale for an arbitrary incentive compatible contract, and a martingale if the optimal contract is chosen.

Finally, notice that for $U_t \geq U^{FB}(x_t)$, $v_U(U_t,x_t) = -1$, $v_{UU}(U_t,x_t) = v_{Ux}(U_t,x_t) = 0$. Therefore, (20) can be rewritten as

$$r\underbrace{(v(U_t,x_t)+U_t)}_{=s(x_t)} = x + \frac{1}{2}\underbrace{v_{xx}(U_t,x_t)}_{=s_{xx}(x_t)}.$$

From the analysis of Section 4 we know that the principal optimally lets U_t hit 0 at the same time that x_t hits x_L. That is, whenever U_t reaches $U^{FB}(x_t)$, it is optimal to implement the first-best solution as described in Proposition 2.

The next step is to evaluate the principal's profit for an arbitrary incentive compatible contract. That is,

$$E\left[\int_0^\tau re^{-rs}(x_s-w_s)ds + e^{-r\tau}L\right] = E\left[P_\tau\right] \leq P_0 = v(U_0,x_0),$$

with equality if and only if the optimal contract is chosen. This proves that the principal earns the highest feasible value if she chooses the contract as described in Proposition 3. □

Finally, we verify that restricting the attention to payment schedules of the form $dw_t = w_t dt$ is without loss of generality.

Lemma 6. *The optimal income process w is absolutely continuous with respect to t.*

Proof. We show that the principal chooses the income process $\{w_t : t \geq 0\}$ such that it does not exhibit jumps. We show that adding a jump in the income process weakly decreases the principal's value.

Suppose that at t, the principal can profit by adding a jump of size $dw_t = \Delta$ in the agent's income process. Note that then $dU_t = -dw_t = -\Delta$. Intuitively, the principal increases the agent's income by a lump-sum today, which decreases his continuation value by the same lump-sum for tomorrow.

To be more precise, let w_t jump from w_t to $w_t + \Delta$. Then U_t jumps from $U_t + \Delta$ to U_t. Again, we can evaluate the effect of the principal's value by looking at the process P_t in (33). P_t jumps by

$$\begin{aligned} e^{rt} dP_t &= \underbrace{v(U_t, x_t)}_{\text{Principal's value after jump}} - \underbrace{(v(U_t + \Delta) + \Delta)}_{\text{Principal's value before the jump}} \\ &= -\Delta v_U(U_t, x_t) - \Delta \\ &= -\Delta(v_U(U_t, x_t) + 1) \\ &\leq 0. \end{aligned}$$

The second equality follows from the fact that jumps have bounded variation, and the last inequality follows from the fact that $v_U(U_t, x_t) \leq -1$. Note that for $U_t \geq U_t^{FB}$, $v_U(U_t, x_t) = -1$ such that it is only weakly suboptimal to let the agent's income process exhibit jumps if the first-best solution is reached. This reflects the fact that the implementation of the first-best solution is not unique. □

Proof of Proposition 4. Consider an arbitrary, full strategy l chosen by the agent. The agent's expected value from the strategy is

$$U_0(l) = \hat{E}\left[\int_0^\tau re^{-rt}\left(w_t + \lambda \int_0^t l_s ds\right) dt\right]. \quad (35)$$

We compare the agent's expected value from the strategy l with his expected value for the truthful reporting strategy U_0. Notice that $U_\tau = U_\tau(l)$. Reasoning along the same lines than in the proof of Proposition 1, we find that

$$U_0(l) = U_0 - \hat{E}\left[\int_0^\tau e^{-rt}\left(\beta_t l_t - \lambda r \int_0^t l_s ds\right) dt\right].$$

By comparing the agent's expected value under the different strategies, we find that truthtelling is optimal only if

$$U_0(l) - U_0 = -\hat{E}\left[\int_0^\tau e^{-rt}\left(\beta_t l_t - \lambda r \int_0^t l_s ds\right) dt\right]$$

$$
\begin{aligned}
&= -\hat{E}\left[\int_0^\tau \underbrace{\left(e^{-rt} - e^{-r\tau} + e^{-r\tau}\right)}_{=\int_t^\tau re^{-rs}ds + e^{-r\tau}} \beta_t l_t - re^{-rt}\lambda \int_0^t l_s ds dt\right] \\
&= -r\hat{E}\left[\int_0^\tau \int_t^\tau e^{-rs} ds \beta_t l_t - e^{-rt}\lambda \int_0^t l_s ds dt\right] \\
&\quad - \hat{E}\left[\int_0^\tau e^{-r\tau}\beta_t l_t dt\right] \\
&= -r\hat{E}\left[\int_0^\tau e^{-rt}\int_0^t \beta_s l_s ds - e^{-rt}\lambda \int_0^t l_s ds dt\right] \\
&\quad - \hat{E}\left[\int_0^\tau e^{-r\tau}\beta_t l_t dt\right] \\
&= -r\hat{E}\left[\int_0^\tau e^{-rt}\int_0^t (\beta_s - \lambda) l_s ds\right] - \hat{E}\left[e^{-r\tau}\int_0^\tau \beta_t l_t dt\right]. \qquad (36)
\end{aligned}
$$

The forth line uses Fubini's Theorem to exchange the order of integration.

Using the result that at equilibrium, $\beta_t = \lambda$ for all $t \leq \tau$, we can write (36) as

$$
\begin{aligned}
U_0(l) - U_0 &= \underbrace{-r\hat{E}\left[\int_0^\tau e^{-rt}\int_0^t (\lambda - \lambda) l_s ds\right]}_{=0} - \hat{E}\left[e^{-r\tau}\int_0^\tau \lambda l_t dt\right] \\
&= -\lambda\hat{E}\left[e^{-r\tau} \underbrace{\int_0^\tau l_t dt}_{=x_\tau - \hat{x}_\tau \geq 0 \text{ by limited liability}}\right] \leq 0 \text{ a.s.} \qquad (37)
\end{aligned}
$$

The last term of (37) arises because of persistence. It describes the agent's expected loss from choosing his strategy such that the true cash-flow x_t and the reported value $\hat{x}_t$ do not agree as τ is reached. □

References

[1] Battaglini, M. (2005): Long-term Contracting with Markovian Consumers, American Economic Review, 95, 637–658

[2] Battaglini, M. and Lamba, R. (2012): Optimal Dynamic Contracting, mimeo, Princeton University

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

[4] Biais, B., Mariotti, T., Rochet, J-C. and Villeneuve, S (2010): Large Risks, Limited Liability, and Dynamic Moral Hazard, Econometrica, Vol. 78, No. 1, 73–118

[5] DeMarzo, P.M. and Fishman, M.J. (2007): Optimal Long-Term Financial Contracting, Review of Financial Studies, 20, 2079–2128

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

[7] DeMarzo, P.M., and Sannikov, Y. (2011): Learning, Termination, and Payout Policy in Dynamic Incentive Contracts, mimeo

[8] Dixit, A.K. and Pindyck, R.S. (1994): Investment under Uncertainty, Princeton University Press

[9] Eso, P. and Szentes, B. (2013): Dynamic Contracting with Adverse Selection: An Irrelevance Result.

[10] Garrett, D. and Pavan, A. (2011): Dynamic Managerial Compensation: On the Optimality of Seniority-based Schemes, working paper, Northwestern University

[11] Garrett, D. and Pavan, A. (2012): Managerial Turnover in a Changing World, Forthcoming in Journal of Political Economy

[12] Pavan, A., Segal, I. and Toikka, J. (2012): Dynamic Mechanism Design, mimeo, MIT, Northwestern University, and Stanford University

[13] Rochet, J.-C. and Choné, P. (1998): Ironing, Sweeping, and Multidimensional Screening. Econometrica, 66, No. 4, 783–826

[14] Sannikov, Y. (2008): A Continuous-Time Version of the Principal-Agent Problem. Review of Economic Studies, Vol. 75, No. 3, 957–984

[15] Sannikov, Y. (2013): Moral Hazard and Long-Run Incentives.

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

[17] Strulovici, B.H. (2011): Renegotiation-Proof Contracts with Moral Hazard and Persistent Private Information, mimeo, Northwestern University

[18] Zingales, L. (1998): Survival of the Fittest or the Fattest? Exit and Financing in the Trucking Industry, Journal of Finance, Vol. 53, No. 3, 905 – 938

[19] Williams, N. (2011): Persistent Private Information, Econometrica, 79(4), 1233–1274