Dynamic Contracting with Markovian Cash-Flows

Suvi Vasama

March 3, 2015

Preliminary and incomplete

Abstract

We examine a dynamic principal-agent model in which the firm cash-flow follows a simple Markov process. The cash-flow is unobservable by the principal and must be reported by the agent. The agent has the opportunity to conceal part of the cash-flow from the principal to divert it for his private benefits. If the agent deviates, the players' perceptions about the future cash-flow diverge. Thus, the agent has persistent private information. We prove the one-shot deviation principle in our framework and derive the optimal contract using two state variables: the agent's continuation value and the current cash-flow. We examine how the optimal contract changes if the correlation between the periods changes. We find that when the correlation increases, the incentive problem becomes more severe. In the same time, the probabilistic liquidation threat is a less powerful instrument to provide the agent incentives.

1 Introduction

We examine a dynamic cash-flow diversion model in which the cash-flow follows a Markov process. The agent runs a firm that produces a firm of which cash-flow is unobservable by the principal. The agent reports the cash-flow to the principal and has the possibility to conceal part of it to divert it for his private benefits. Since the cash-flow follows a Markov process, the current cash-flow determines future expectations and is a sufficient statistic for the firm value. If the agent diverts cash-flows, the players' expectations about the future cash-flows diverge. Thus, the agent has persistent private information.

We consider a version of the cash-flow diversion framework in which the agent's private benefit is proportional to the diverted cash-flow. Under this specification, the agent's private benefit from diverting an additional dollar is always independent of the true cash-flow. We show that in our framework, it is necessary and sufficient for incentive compatibility to consider incentive schemes that rule out one-shot deviations by the agent. The one-shot deviation principle simplifies the analysis considerably and guarantees that our model has a similar recursive structure than related models with independently distributed cash-flows.

We assume that the cash-flow follows a simple Markov process with only two possible states of the world. The simplicity of our framework allows us to draw consequences about how the optimal contract changes when the correlation between the periods increases. We find that the incentive problem becomes more severe when the correlation is higher. However, the probabilistic liquidation threat becomes a less valuable instrument.

In the beginning of the relationship, the contract relies on an inefficient termination threat to provide the agent incentives. The payments to the agent are delayed, and the contract accumulates him value for the future. The principal receives the firm cash-flow. With good enough past performance, the agent's continuation value eventually becomes high enough such that the first-best solution can be reached. The firm is never liquidated and the agent is rewarded by delivering him a positive payment whenever he reports a high cash-flow realization.

We characterize the optimal contract using two state variables: the agent's continuation value and the current cash-flow. Spear and Srivastava (1989) showed that the agent's continuation value is a sufficient statistic for the history in the optimal contract. The current cash-flow is a sufficient statistic for the firm value.

1.1 Related Literature

The cash-flow diversion model has turned out to be useful for examining important questions in corporate finance. Our model is a dynamic extension of Bolton and Scharfstein (1990) that was earlier extended by DeMarzo and Fishman (2007) and Biais et al. (2007).[1]

We contribute to the fast growing literature studying cash-flow diversion with serially correlated cash-flows. A similar the model was earlier solved by Tchistyi (2006) who solves a dynamic cash-flow diversion model with serially correlated cash-flows. The main difference is methodological. Tchistyi examines a finite horizon version of the model that he solves using backward induction methods. We study an infinite horizon game, which allows to exploit the stationarity of the problem. Our game has a recursive structure and we can solve the optimal contract by applying dynamic programming methods similar to the models with

[1] DeMarzo and Sannikov (2006) show that the model becomes significantly more tractable in continuous time.

independent cash-flows. Interestingly, the problem becomes much more tractable such that we are able to make significantly more detailed predictions about how the optimal contract changes as we vary the degree of correlation between the periods.

Our paper is related to several recent papers that analyze similar framework in continuous time. Faingold and Vasama (2015) analyze a related model in which the cash-flow follows a standard Brownian motion.[2] While continuous time allows to adopt more states of the world, the simplicity of a two-stage Markov process allows us to derive conclusions about how the optimal contract depends on the degree of correlation between the periods. Besides, discrete time allows to better explore the game-theoretic structure of the model. In particular, we are able to prove the one-shot deviation principle for our game.

2 Setting

We examine a game with two players, a principal and an agent. Both players are risk-neutral and the principal has access to unlimited funds. The agent is protected by limited liability. In our framework, this assumption implies that the contract cannot impose negative payments to the agent.

The firm needs an initial investment of $I \geq 0$ to be started. The agent has initial wealth $A \geq 0$ that he invests in the firm. If $I > A$ the principal needs to cover the part $I - A$ of the investment cost.

The firm produces a stochastic cash-flow that is unobservable by the principal. In each period, the cash-flow can take two possible values, high or low. We denote the high cash-flow by $x > 0$. The low cash-flow is normalized to 0. We assume that the initial cash-flow is known by the players.

The cash-flow follows a two-stage Markov process. Given the current cash-flow, the next period outcome is distributed as follows

$$\begin{aligned} \Pr(x_t = x | x_{t-1} = x) &= p, & \Pr(x_t = x | x_{t-1} = 0) &= 1 - p, \\ \Pr(x_t = 0 | x_{t-1} = x) &= 1 - p, & \Pr(x_t = 0 | x_{t-1} = 0) &= p. \end{aligned}$$

We assume that the cash-flow is positively correlated, that is, $p \geq 1/2$. Notice that continuation is always efficient since

$$(1 - p)x \geq 0. \tag{1}$$

The cash-flow is unobservable by the principal and is reported by the agent. Let $\hat{x}_t$ denote the cash-flow that the agent delivers to the principal in period t.

[2]DeMarzo and Sannikov (2014) solve a related model in which the players learn about the unknown mean of the cash-flow over time. Williams (2011) and Strulovici (2011) analyze related models in which the agent is risk averse.

The agent has the opportunity to conceal a positive cash-flow from the principal and divert it for his private benefits. If the agent conceals the high cash-flow, he receives a private benefit λx. We assume that $\lambda \in (0,1)$ such that cash-flow diversion is inefficient. Thus, diverting a dollar is related with social cost of $1-\lambda$.

We assume that the agent cannot save privately. Under this assumption, limited liability implies that the agent cannot overreport cash-flows. Thus, $\hat{x}_t \le x_t$ for all t.

As is standard in the literature, the principal has two instruments to provide the agent incentives: nonnegative payments and liquidation. Liquidation is irreversible, and generates the players a payoff that we normalize to 0. At any period t before liquidation, the optimal contract determines a wage to the agent w_t, and a liquidation probability $1-\alpha_t \in (0,1)$. We assume that the players can commit to the contract for its lifetime.

TIMELINE OF EVENTS

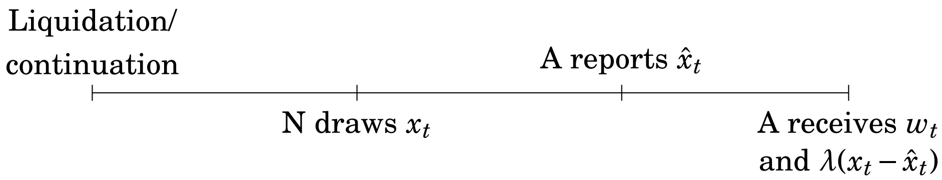

The players are equally patient. We let $\delta \in (0,1)$ denote the common discount factor. With full commitment, we restrict the attention on contracts in which the agent reports the cash-flows truthfully. The restriction is justified by the revelation principle and is without loss of generality.

Public history in period t[3]

$$h^t = \{w_0, \ldots, w_t; \alpha_0, \ldots, \alpha_t; \hat{x}_0, \ldots, \hat{x}_t\}.$$

The agent's private history in period t includes both the public history and actual cash-flow realizations. Formally,

$$\hat{h}^t = \{h^t; x_0, \ldots, x_t\}.$$

A contract determines a payment schedule and a liquidation probability contingent on the public history. We assume that the players can fully commit to the contract.

[3] Since liquidation is irreversible, we define $x_t = 0$ and $\hat{x}_t = 0$ for all periods following the liquidation.

3 One-Shot Deviation Principle

Notice that if the cash-flow is low, the agent has no option but to report truthfully. Therefore, it is without loss of generality to set the payment to 0 following a low cash-flow report. Providing a positive payment would decrease the principal's profit and strengthen the agent's incentive constraint following a high cash-flow report. Let $w(h^t)$ denote the payment to the agent following a high cash-flow report.

To find conditions for the agent to be willing to report the cash-flows truthfully, we need to determine how his continuation value has to change if he reports cash-flows truthfully. In period t, the agent's continuation value is

$$u(h^t) = E\left[\sum_t^{\infty} \alpha(h^t)w(h^t)\middle| h^{t-1}\right]. \tag{2}$$

If the agent diverts a high cash-flow realization, he earns a private benefit of λx. For the contract to be incentive compatible, the agent's continuation value has to increase by the same λx if he reports the high cash-flow truthfully. The increase in the continuation value is independent of the state of the world. As a consequence, the agent's incentives to report truthfully are the same on and off the equilibrium path. This implies that the agent has the right incentives to report the cash-flow truthfully, regardless of if he diverted in the past. Thus, the one-shot deviation principle holds in our framework.

The following proposition proofs the one-shot deviation principle in our game

Proposition 1. *Consider a contract with $u_t < \infty$. Suppose that $x_t = x$. A necessary and sufficient condition for incentive compatibility is that*

$$w(h^t) + \delta u(h^{t+1}) \geq \lambda x + \delta u(\bar{h}^{t+1}) \tag{3}$$

for all t and all truthful histories h^t and all alternative histories $\bar{h}^t$ such that $\hat{x}_t < x_t$ only in period t.

Proof. See Appendix. □

4 Dynamic Programming Formulation

We characterize the optimal contract using two state variables: the agent's continuation value u_t and the cash-flow x_t. The intuition behind the recursive structure is similar to the uncorrelated case. The agent's continuation value is sufficient statistic for the history of reports in the optimal contract. Moreover, since the cash-flow follows a Markov process, the current cash-flow is a sufficient statistic for the firm value.

We decompose the problem in two parts, depending on if the cash-flow is high or low. Notice, in particular, that there are two possible pairs denoting the current state of the world: (u^h, x) and $(u^l, 0)$, where u^h and u^l denote the agent's continuation value starting from the high and the low state of the world. We let w^h and w^l denote the payments to the agent following a high cash-flow report when starting from the high and the low state.

The agent's state-contingent continuation values have to satisfy the following functional equations

$$u^h = \alpha^h \left[pw^h + \delta(pu^{hh} + (1-p)u^{hl}) \right]. \tag{4}$$

and

$$u^l = \alpha^l \left[(1-p)w^l + \delta((1-p)u^{lh} + pu^{ll}) \right]. \tag{5}$$

The promise keeping constraints (4) and (5) are book-keeping constraints that guarantee that the agent receives his expected value at the optimal contract when he reports truthfully.

Proposition 1 implies that it is sufficient to restrict the attention on one-shot deviations by the agent at the optimal contract. Starting from the high state of the world, the incentive compatibility condition can be written as

$$w^h + \delta u^{hh} \geq \lambda x + \delta \hat{u}^{hl}, \tag{6}$$

where $\hat{u}^{hl}$ is the agent's continuation value if he reports a low cash-flow when the true cash-flow was high.

Notice that the agent's continuation value following a deviation differs from the agent's continuation value following a truthful low cash-flow report in terms of probability distribution. Here the liquidation probabilities, the payments and the agent's future continuation values are determined by the principal's problem following the corresponding truthful history.

Similarly, the incentive constraints starting from a low state can be written as

$$w^l + \delta u^{lh} \geq \lambda x + \delta \hat{u}^{ll}, \tag{7}$$

where $\hat{u}^{ll}$ denotes the agent's continuation value following a deviation.

Starting from each state of the world, the principal chooses a liquidation policy, a nonnegative payment and the agent's continuation values for the next period. Apart from the incentive and the promise-keeping constraints, the choice variables have to satisfy the following feasibility and limited liability constraints

$$(\alpha^h, \alpha^l, w^h, w^l, u^{hh}, u^{hl}, u^{lh}, u^{ll}) \in [0,1]^2 \times \mathbb{R}^6 \tag{8}$$

Starting from the high state of the world, the principal's value function satisfies her Bellman equation

$$v(u^h, x) = \max\left\{\alpha^h \left[p(x - w^h) \ \ + \delta(pv(u^{hh}, x) + (1-p)v(u^{hl}, 0))\right]\right\}, \tag{9}$$

and starting from the low state of the world

$$v(u^l, 0) = \max\left\{\alpha^l \left[(1-p)(x - w^l) \ \ + \delta((1-p)v(u^{lh}, x) + pv(u^{ll}, 0))\right]\right\} \tag{10}$$

It is useful to rewrite the problem in terms of maximizing the firm value subject to the feasibility constraints. Define

$$s(u^h, x) \equiv v(u^h, x) + u^h,$$

and

$$s(u^l, 0) \equiv v(u^l, 0) + u^l.$$

Then we can rewrite the problem as one of maximizing the firm value that is more suitable for calculations

Lemma 1. *The functions $s(u^h, x)$ and $s(u^l, 0)$ solve the following program*

$$s(u^h, x) = \max\{\alpha^h[px + \delta[ps(u^{hh}, x) + (1-p)s(u^{hl}, 0)]]\}, \tag{11}$$

$$s(u^l, 0) = \max\{\alpha^l[(1-p)x + \delta[(1-p)s(u^{lh}, x) + ps(u^{ll}, 0)]]\}, \tag{12}$$

subject to the feasibility and limited liability constraints

$$(\alpha^h, \alpha^l, u^{hh}, u^{hl}, u^{lh}, u^{ll}) \in [0,1]^2 \times \mathbb{R}^4, \tag{13}$$

the promise-keeping constraints that guarantee that the agent receives his promised value from the contract if he reports the cash-flows truthfully

$$u^h \geq \alpha^h[pu^{hh} + (1-p)u^{hl}], \tag{14}$$

$$u^l \geq \alpha^l[(1-p)u^{lh} + pu^{ll}], \tag{15}$$

and the incentive constraints that guarantee that the agent receives at least the same value for truthful reporting than for diverting the high cash-flow realization

$$u^h \geq \alpha^h[p\lambda x + \delta(u^{hl} + p(\hat{u}^{hl} - u^{hl}))], \tag{16}$$

$$u^l \geq \alpha^l[(1-p)\lambda x + \delta(u^{ll} + p(\hat{u}^{ll} - u^{ll}))]. \tag{17}$$

Proof. (11) follows by adding up (4) and (9).

Rewrite (6) to obtain

$$w^h \geq \lambda x - \delta[u^{hh} - \hat{u}^{hl}].$$

Then we can rewrite (4) as

$$\begin{aligned} u^h &\geq \alpha[w^l + p\lambda x - \delta[p(u^{hh} - \hat{u}^{hl})] + \delta[p(u^{hh} - u^{hl})] + \delta u^{hl}] \\ &\geq \alpha^h[p\lambda x + u^{hl} + p(\hat{u}^{hl} - u^{hl})]. \end{aligned}$$

(16) follows by reorganizing. (17) follows reasoning along the same lines.

(14) and (15) follow by using the limited liability constraints $w^h \geq 0$ and $w^l \geq 0$ to reduce (4) and (5). □

The first two terms in the incentive constraints (16) and (17) are standard, and would describe how much the agent's continuation value had to decrease following a high report for the contract to be incentive compatible if the agent had no persistent private information.[4] The last term is new, and describes the agent's additional private information about the distribution of the future cash-flows. If the agent deviates, the players' perceptions about both the future cash-flow and the agent's continuation value diverge. For the agent to be willing to report truthfully, the principal has to provide a compensation for the additional information rent.

The one-shot deviation principle, that we proved in Proposition 1, implies that the agent cannot profit from the additional private information. Moreover, notice that since $\hat{u}^{hl} \geq u^{hl}$, the right hand sides of (16) and (17) increase as p increases. Whenever the cash-flow is high, the principal has to compensate the agent for both revealing the higher cash-flow today and the higher continuation value in the future. The additional information rent increases as p increases. The observation suggests that the incentive problem is more severe when the correlation is higher.

Next, it follows from (16) and (17) that whenever $u^h < p\lambda x$ or $u^l < (1-p)\lambda x$, liquidation must occur with positive probability. The result is in line with the earlier results obtained with independently distributed cash-flows.[5] However, notice that if p is higher, the agent's continuation value decreases more if the low cash-flow is realized. As a consequence, the liquidation probability has to increase more following a low cash-flow realization. Indeed, (17) suggests that the probabilistic liquidation threat becomes a less powerful tool when the p is increases.

[4]Notice that in the independent case, $\hat{u}^{kl} = u^{kl}$, $k \in \{h, l\}$ and the last term of (16) and (17) vanishes, cf. Biais et al. (2007).

[5]See, for example, Clementi and Hopenhayn (2006), Biais et al. (2007) or DeMarzo and Fishman (2007).

5 First-Best Solution

In this section, we discuss how to implement the first-best solution, and solve for the players' value functions under the first-best regime. Notice that inefficient liquidation is the only source of inefficiency in the optimal contract. Because of the agent's limited liability, the contract cannot impose him negative payments. To save on incentive cost to the agent, the optimal contract uses an inefficient liquidation threat to provide the agent incentives.

Notice that by (1), liquidation is never efficient. Thus, and the first-best solution never liquidates the firm such that $\alpha^h = \alpha^l = 1$. Moreover, since both players are risk-neutral, it is (weakly) optimal to provide the agent incentives by rewarding him with am immediate payment following a high cash-flow report.

In the first-best regime, the agent's continuation value only depends on the current cash-flow that is the only payoff relevant state of the world. That is, $u^h = u^{hh} = u^{lh} = \hat{u}^{hl} = \hat{u}^{ll}$ and $u^l = u^{hl} = u^{ll}$. Let $u^h \equiv u^{h,fb}$ and $u^l \equiv u^{l,fb}$ denote the agent's continuation values in both states in the first-best regime. Substituting into (6) and (7), we find that the agent's compensation has to be at least

$$w^h + \delta u^{h,fb} = \lambda x + \delta u^{h,fb},$$
$$w^l + \delta u^{h,fb} = \lambda x + \delta u^{h,fb},$$

from which we can immediately conclude that the agent receives a payment $w \equiv w^h = w^l = \lambda x$.

Substituting for the payments in (4) and (5), the promise keeping constraints simplify to

$$u^{h,fb} = p\lambda x + \delta(pu^{h,fb} + (1-p)u^{l,fb}), \tag{18}$$
$$u^{l,fb} = (1-p)\lambda x + \delta((1-p)u^{h,fb} + pu^{l,fb}). \tag{19}$$

Solving the system of equations, we find that the agent's continuation values are

$$u^{h,fb} = \frac{p(1-\delta)+\delta(1-p)}{(1-\delta)(1-\delta p+\delta(1-p))}\lambda x, \tag{20}$$
$$u^{l,fb} = \frac{1-p}{(1-\delta)(1-\delta p+\delta(1-p))}\lambda x. \tag{21}$$

Next, we can determine the first-best optimal firm value. If the previous period cash-flow was high, the cash-flow is high with probability p. In that case, the next-period value function takes the same value in the next period. With probability $1-p$, the cash-flow is low, and the firm value in the next period is low. Thus the firm value in the high state satisfies

$$s^{fb}(x) = px + \delta(ps^{fb}(x) + (1-p)s^{fb}(0)), \tag{22}$$

Starting from the low state of the world, the firm value satisfies

$$s^{fb}(0) = (1-p)x + \delta(ps^{fb}(0) + (1-p)s^{fb}(x)). \tag{23}$$

Solving for the system of equation (22) and (23), we find the following expressions

$$s^{fb}(x) \equiv \frac{p(1-\delta)+\delta(1-p)}{(1-\delta)(1-\delta p+\delta(1-p))}x, \tag{24}$$

$$s^{fb}(0) \equiv \frac{1-p}{(1-\delta)(1-\delta p+\delta(1-p))}x. \tag{25}$$

In every period, the principal receives the cash-flow net of the payment to the agent. Subtracting (20) and (21) from (24) and (25), we find that, under the first-best regime the principal earns the state-contingent value

$$v^{fb}(x) \equiv \frac{p(1-\delta)+\delta(1-p)}{(1-\delta)(1-\delta p+\delta(1-p))}(1-\lambda)x, \tag{26}$$

starting from the high state, or

$$v^{fb}(0) \equiv \frac{1-p}{(1-\delta)(1-\delta p+\delta(1-p))}(1-\lambda)x. \tag{27}$$

starting from the low state.

6 Optimal Contract

In this section, we derive the optimal contract from the principal's problem. We show that the optimal contract relies on an inefficient liquidation threat to provide the agent incentives in the beginning of the relationship. The payments to the agent are delayed and the optimal contract provides the agent incentives by promising a higher continuation value for the future.

With good enough past performance, the agent's state-contingent continuation value eventually reach the first-best optimal levels (20) and (21). Thereafter, the firm is never liquidated, and it is without loss of generality implement the contract in which the agent is rewarded for high reports by providing an immediate payment.

We start by proving some useful properties of the value function $s(u, x_t)$. First in the range of the agents continuation values that are high enough, such that the first-best solution is attainable, the value function $s(u_t, x_t)$ is linear in u_t.

Lemma 2. *The function $s(u, x_t)$ is continuous, concave and increasing. For $k \in \{l, h\}$, the solution satisfies the following properties.*

(i) *When* $u^k \in (0, u^{k,c})$, *the firm is liquidated with probability* $1 - u/u^{k,c}$.

(ii) *When* $u^k \in (0, u^{k,fb})$, $s(u, x_t)$ *is strictly increasing.*

(iii) *When* $u^k \geq u^{k,fb}$, $s(u, x_t) = s^{fb}(x_t)$ *is constant.*

Proof. See Appendix. □

In the beginning of the relationship, the payments to the agent are delayed and the optimal contract relies on an inefficient liquidation threat to provide the agent incentives. Concavity of the value function implies that it is optimal to let the incentive constraints (16) and (17) bind. The optimal contract punishes the agent by decreasing his continuation value at the smallest feasible rate to minimize the risk of inefficient liquidation.

With good enough past performance, the contract has eventually accumulate sufficiently high value to the agent such that the first-best solution can be attained. If the state-contingent continuation values reach (20) and (21), the first-best solution can be reached. The agent is rewarded by delivering him an immediate income of λx following a high cash-flow report.

The optimal contract is summarized in the following proposition

Proposition 2. *The optimal incentive compatible contract is characterized by three regimes. Let* $k \in \{l, h\}$ *and* $x_t \in \{0, x\}$.

(i) *When* $u^k \in (0, u^{k,c})$, $k \in \{l, h\}$, *the firm is liquidated with probability* $1 - u/u^{k,c}$. *The payments to the agent are delayed such that the players receive the payments*

$$w_t = 0 \tag{28}$$

$$x_t - w_t = x_t. \tag{29}$$

(ii) *When* $u^k \in [u^{k,c}, u^{k,fb})$, *the firm is continued with probability* 1. *The payments to the agent are delayed such that the players receive the payments (28) and (29).*

(iii) *When* $u^k \geq u^{k,fb}$, *the optimal contract attains the first-best solution. The players receive the payments*

$$w_t = \lambda x_t \tag{30}$$

$$x_t - w_t = (1 - \lambda) x_t. \tag{31}$$

7 Initialization of the Contract

In this section, we examine conditions under which the firm receives funding. As is standard in the principal-agent framework, the optimal contract depends on the relative bargaining power of the players. We consider the two extreme cases: the one in which the agent has all bargaining power and the one in which the principal has all the bargaining power.

The initial cash-flow is given and known by the players at the contracting stage. It remains to determine the agent's initial value that maximizes the players' profits, given the players' relative bargaining powers and subject to their participation constraints. To the latter constraints, the players' expected profits have to cover the cost of investment.

If the agent has all the bargaining power, the contract maximizes the agent's value given that the principal has to earn a value in the contract that covers her cost of investment $I-A$. Formally,

$$u_0 = \max\{u \in \mathbb{R}_+ | v(u,x_0) \geq I-A\}$$

If the principal has all the bargaining power, the contract maximizes her value given that the agent has to earn at least his share of the cost of investment A. Formally,

$$u_0 \in \arg\max_{u \in \mathbb{R}_+} \{v(u,x_0) | u_0 \geq A\}.$$

Notice that the agent's participation constraint might not bind at the equilibrium. Increasing the agent's value reduces the risk of inefficient liquidation. Decreasing the agent's value allows the principal to extract more rents from the agent. Thus, there is a trade-off between efficiency and rent extraction. Notice, in particular, that since the principal's value is concave in the agent's value and the agent earns a positive value at the optimal contract, the first-best regime cannot be reached at the contracting state.

The pledgeable income is $\max_{u_0 \in \mathbb{R}_+}\{v(u_0,x_0)\}$. If the moral hazard problem is very severe, the optimal contract might not be able to cover the cost of investment for the principal. As in static framework, moral hazard may lead to credit rationing.

8 Conclusion

9 Appendix

Proof of Proposition 1. We first prove the result for a finite horizon game and then show that the game is continuous at infinity. This allows us to extent the result for the infinite horizon game.

First, assume that the game is finite, and let $T < \infty$ denote the last period. We show that (3) is a necessary and sufficient condition for incentive compatibility. Necessity is trivial. To proof sufficiency, we show that, if the one-shot incentive condition holds, longer deviations are not optimal for the agent.

The proof follows by contradiction. Assume that the condition (3) holds, but there is a profitable deviation that induces a public history $\bar{\bar{h}}^t$ such that $\hat{x}_s < x_s$ for more that one period $s \geq t$. Let t' denote the last period in which $\hat{x}_t < x_t$. Since $T < \infty$ such a period t' exists.

Now consider the agent's incentives to deviate at period t'. Since (3) implies that

$$w(\bar{\bar{h}}^{t'}) + \lambda x + u(\bar{\bar{h}}^{t'}) \leq w(\bar{\bar{h}}^{t'-1}, x_t) + u(\bar{\bar{h}}^{t'-1}, x_t),$$

diverting cannot be optimal in period t', a contradiction. Now if $t' = t+1$, we are done. Otherwise, we construct a strategy that includes diverting for more than one period, but agrees with truthful reporting in period t'. We can repeat the argument to show that diverting is suboptimal in the last period and so on. The argument proves that the agent cannot profit from deviations of finite length if the one-shot incentive condition (3) is satisfied.

Next, we extent the argument to allow for a infinite time horizon. That is, we need to show that the agent cannot gain from some infinite sequence of deviations.

First, notice that, since $\delta < 1$ and $u(h^t) < \infty$, the transversality condition

$$\lim_{t \to \infty} E[\delta^t u(h^t)] = 0$$

holds. Therefore,

$$\sup_{h, \bar{\bar{h}} \text{ s. t. } h^t = \bar{\bar{h}}^t} E|\delta^t (u(h^t) - u(\bar{\bar{h}}^t))| = \sup_{h, \bar{\bar{h}} \text{ s. t. } h^t = \bar{\bar{h}}^t} \delta^t E|(u(h^t) - u(\bar{\bar{h}}^t))| \to 0 \tag{32}$$

as $t \to \infty$. The result follows by dominated convergence theorem since $0 \leq u(h^t) < \infty$ for all t.

Suppose again that the one-shot deviation condition (3) holds, but the contract is not incentive compatible. Now, there is a period t and a history $\bar{\bar{h}}$ such that $\hat{x}_t$ and x_t do not agree for more infinitely many periods, and that yields a high payoff than truthful reporting. Let $\hat{x}$ denote the agent's full strategy, that includes the deviation, and $\varepsilon > 0$ the improvement when compared to the truthtelling strategy.

By (32) there is an alternative strategy $\hat{x}'$ such that $\hat{x}$ and $\hat{x}'$ that agree until period t', includes truthtelling from period t' onwards and improves the agent's payoff by $\varepsilon/2$. However, this is a contradiction to the result that the agent cannot benefit from deviations of finite length. $\square$

Proof of Lemma 2. (iii) follows directly from (24) and (25). It remains to prove (i) and (ii).

The proof uses an argument similar to Clementi and Hopenhayn (2006) and Biais et al. (2007). Let T denote the Bellman operator associated to (11) and (12). We prove the result starting from the state of the world (u^h, x). The goal is to prove that the fixed point

$$s(u^h, x) = Ts(u^h, x)$$

admits the desired properties.

Consider a bounded, continuous function $s \in B(\mathbb{R}_+)$. Notice that the function obtains its maximum at $u^{h,fb}$ where it takes the value (24). Thus, Ts is bounded from above. Moreover, $Ts \geq 0$ such that it is bounded from below.

Since $u^{hh} \in [0, u^{h,fb}]$ and $u^{hl} \in [0, u^{fb,l}]$, Berge's maximum theorem applies and s is continuous.

Next, we prove that s is nondecreasing. Let $\bar{u}^h \geq u \geq 0$. Let $(\alpha^h, u^{hh}, u^{hl})$ be the optimal choice in the program that implies $Ts(u^h, x)$. Since the constraint set (14)-(17) is linear, $(\alpha^h, u^{hh}, u^{hl})$ is feasible in the program starting from the state $(\bar{u}^h, x_t)$. Therefore, $Ts(\bar{u}, x) \geq Ts(u, x)$ such that Ts is nondecreasing.

We prove that s is concave. Decompose (11) - (15) into two subproblems: with and without liquidation.

Let T^c denote the Bellman operator for the problem upon continuation.

$$T^c s(u^h, x) = \max\left\{px + \delta[ps(u^{hh}, x) + (1-p)s(u^{hl}, 0)]\right\} \tag{33}$$

subject to the constraints

$$(u^{hh}, u^{hl}) \in \mathbb{R}^2, \tag{34}$$

$$u^h \geq p\lambda x + \delta \hat{u}^{hl}, \tag{35}$$

and

$$u^h \geq pu^{hh} + (1-p)u^{hl}, \tag{36}$$

Assume that $s(u^h, x)$ is a concave function. We prove that $T^c s(u^h, x)$ is concave. Suppose that $\bar{u}^h \geq u \geq p\lambda x$. Let $(\bar{u}^{hh}, \bar{u}^{hl})$ and (u^{hh}, u^{hl}) denote the optimal choices under the both programs. Furthermore, let $u^h_\theta = \theta\bar{u}^h + (1-\theta)u^h$, where $\theta \in [0,1]$. Let $(u^{hh}_\theta, u^{hl}_\theta)$ denote the optimal choice in the program starting from the state (u^h_θ, x).

It follows from the linearity of the constraint set that $\theta(\bar{u}^{hh},\bar{u}^{hl})+(1-\theta)(u^{hh},u^{hl})$ is a feasible choice in the program starting from the state (u^h_θ,x). Since $s(u,x)$ is concave, we find that

$$\begin{aligned} T^c s(u^h_\theta,x) &= \max\left\{px+\delta[ps(u^{hh}_\theta,x)+(1-p)s(u^{hl}_\theta,0)]\right\} \\ &\geq \max\left\{px+\delta[p[\theta s(\bar{u}^{hh},x)+(1-\theta)s(u^{hh},x)]+(1-p)[\theta s(\bar{u}^{hl},0)+(1-\theta)s(u^{hl},0)]]\right\} \\ &= \theta T^c s(\bar{u}^h,x)+(1-\theta)T^c s(u^h,x) \end{aligned}$$

i.e., $T^c s(u^h,x)$ is concave.

Next, taking into account liquidation, we can write the program as

$$Ts(u^h,x)=\max\{\alpha T^c s(u^{h,c},x)\} \tag{37}$$

subject to

$$(\alpha^h,u^{h,c})\in[0,1]\times\left[p\lambda x,\frac{p(1-\delta)+\delta(1-p)}{(1-\delta)(1-\delta p+\delta(1-p))}\lambda x\right] \tag{38}$$

and

$$u^h=\alpha^h u^{h,c}. \tag{39}$$

Solving for α^h from (39) and substituting, we can rewrite (37) and (38) as

$$Ts(u^h,x)=\max\left\{\frac{u^h}{u^{h,c}}T^c s(u^{h,c},x)\right\} \tag{40}$$

for all $u^h\geq 0$ subject to

$$u^{h,c}\geq\max\{p\lambda x,u^h\}. \tag{41}$$

Next, notice that since $T^c s(u^h,x)$, is continuous, $T^c s(u^h,x)/u^{h,c}$ is a continuous function. Therefore, it reaches its maximum on the interval $[p\lambda x,u^{h,fb}]$. Since $T^c s(u^h,x)$ is concave, the set of maximizers is an interval $[\underline{u}^{h,c},\overline{u}^{h,c}]$, possibly a point. It follows that

$$Ts(u^h,x)=\begin{cases}\dfrac{u^h}{u^{h,c}}T^c s(u^h,x) & \text{if} \quad u^h\leq\overline{u}^{h,c} \\ T^c s(u^h,x) & \text{if} \quad u^h>\overline{u}^{h,c}\end{cases}$$

which implies that $Ts(u^h,x)$ is concave.

It remains to proof that $s(u^h,x)$ us strictly increasing on $(u^{h,c},u^{h,fb})$. The proof follows by contradiction. Suppose that there are two values u^h and $\bar{u}^h$ such that $u^h<\bar{u}^h<u^{h,fb}$, but $s(u^h,x)=s(\bar{u}^h,x)$. However, by concavity of s, we must have that $s(u^h,x)=s(\bar{u}^h,x)=s(u^{h,fb},x)$, a contradiction. □

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] Bolton, P. and Scharfstein, D. (1990): A Theory of Predation Based on Agency Problems in Financial Contracting, American Economic Review, 80 (1), 93–106

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

[9] Faingold, E., and Vasama, S. (2015): Real Options and Dynamic Incentives

[10] Fernandes, A. and Phelan, C. (2000): A Recursive Formulation for Repeated Agency with History Dependence, Journal of Economic Theory, Volume 91, Issue 2, April 2000, 223–247

[11] Green, E. (1987): Contractual Arrangements for Intertemporal Trade, in Minnesota Studies in Macroeconomics, edited by Edward Prescott and Neil Wallace. Minneapolis: University of Minnesota Press, 1987.

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

[13] Garrett, D. and Pavan, A. (2012): Managerial Turnover in a Changing World, Journal of Political Economy, 120 (5), October 2012, 879-925.

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

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

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

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

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

[19] Tchistyi, A. (2006): Security design with correlated hidden cash flows: The optimality of performance pricing

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