Agency Problems, Screening and Increasing Credit Lines. [†]

Yuliy Sannikov[*]

August 14, 2007

ABSTRACT. We propose a model in which an optimal dynamic financing contract for a cash-constrained entrepreneur is a credit line with a growing credit limit. This simple contract, which resembles those used in practice, presents a good benchmark to understand dynamic moral hazard and adverse selection. In our setting the adverse selection problem is that only the agent initially knows the quality of the project. The moral hazard problem is that the agent, who privately observes stochastic cash flows, can manipulate them using hidden savings. It is appealing that the credit line is an incomplete contract: it does not spell out the agent's actions, but gives him full discretion to draw and deposit funds up until the credit limit. The agent has incentives to use discretion in a way that is optimal for the principal.

---

[*] Department of Economics, University of California at Berkeley, and Department of Finance, NYU, sannikov@econ.berkeley.edu.

1. Introduction

This paper derives a credit line as an optimal dynamic financing contract in under moral hazard and adverse selection. Credit lines are important in practice: according to Sufi (2007), 80% of public firms utilize credit lines. Credit lines are also used by consumers, home owners and small businesses. To understand credit lines, one needs a dynamic theory that takes into account not only project fundamentals, such as the level and the variance of cash flows, but also managerial incentives. From the theoretical perspective, this paper develops a new method to study adverse selection in a dynamic contracting environment with moral hazard. Despite the complexity of contract space and the incentive constraints, we make modeling choices that lead to a remarkably simple optimal contract in the form of a credit line. In this contract, the credit limit increases deterministically over time, as the adverse selection problem becomes weaker due to continuing interaction between the lender and the borrower. Our simple solution to a complex dynamic problem presents a clear way to think about many real-world settings.

This paper innovates over the existing literature on dynamic contracting by allowing for adverse selection. Adverse selection is important in practice. An existing firm, which opens a credit line with a bank to have the option to draw funds in the event of losses, may have private information about its future profitability. A home buyer may have private information about her housing preferences and/or future income.[1] Private information is important in the context of micro-lending in developing countries, where credit ratings that normally mitigate adverse selection are not as readily available.[2] Despite the important role of adverse selection, dynamic contracting literature has avoided it due to technical difficulties.

We derive an optimal dynamic contract by solving a maximization problem over the complex space of fully contingent contracts. This approach is more fundamental than simply looking at managerial incentives directly in a specific narrow class of contracts, e.g.

---

[1] See Piskorski and Tchistyi (2006), who apply optimal dynamic contracting methods of DeMarzo and Sannikov (2006) to study mortgages in the setting with only moral hazard.
[2] The 2006 Nobel Prize in Peace, given to Muhammad Yunus and Grameen Bank, underlined the importance of micro-lending for economic and social development.

debt contracts. Optimal contracts have made a significant impact in macroeconomics (for example, see the textbook of Ljungqvist and Sargent (2000)), and they are capable to significantly affect our thinking in finance as well. Unlike the study of incentives created by specific institutional mechanisms, such as specific taxes or financial securities, optimal contracts deliver economic conclusions that are valid across institutional environments. The reason is that the contracting parties are interested in approximating the optimal contract as closely as possible with the legal, organizational and reputational means available to them. Studying contracts from the point of view of optimality instead of looking at specific institutional arrangements is analogous to looking at frictionless efficient markets and abstracting from the institutional details of trading. Clearly, the principle of optimality is important for our understanding of contracts.

Optimal contracts are important not only because they apply to a broad range of environments, but because they can generate important policy implications. In practice legal regulations and organizational norms limit ways in which contracts can be written or make certain contractual provisions difficult to enforce. For example, it may be illegal to deny the homeowner a prepayment option on a mortgage. Also, in the event of bankruptcy, it may be difficult and costly for the lender to repossess the borrower's assets. Optimal contracts can help us evaluate inefficiencies created by such institutional frictions, and assess policy suggestions.

The natural form of our optimal contract, an increasing credit line, is consistent with a number of intuitions we have about dynamic lending under moral hazard and adverse selection. Mirrlees (1974) argues that punishments after poor outcomes are appropriate to deal with moral hazard. In a borrowing context, inefficient liquidation can be a natural form of punishment, as seen, for example, in the two-period model of Bolton and Scharfstein (1990). The seminal paper of Akerlof (1970) has first recognized that adverse selection may make borrowing more difficult, citing the extortionate rates charged by local moneylenders in India. Diamond (1989) finds that adverse selection decreases over time in a model with reputation. In a related problem with bilateral learning, Holmstrom (1999) demonstrates that workers have particularly strong incentives initially to put effort in order to influence market's learning. A similar property is present in our contract as well: while

the credit line is growing, the agent strictly prefers to use cash flows to pay down balance (either now or later) rather than to consume them.

Unlike the papers presented above, which look at static contracts or a series of static contracts, we study a setting in which the lender can commit to a long-term contract. A number of papers have studied dynamic contracts in the context of pure moral hazard. Spear and Srivastava (1987), Abreu, Pearce and Stacchetti (1990), Fudenberg, Holmstrom and Milgrom (1990) and Phelan and Townsend (1991) develop general theories of repeated agency under moral hazard. Albuquerque and Hopenhayn (2004), Clementi and Hopenhayn (2002) and DeMarzo and Fishman (2005) apply these theories to dynamic financing. Sannikov (2004) develops new continuous-time methods to analyze dynamic moral hazard, and DeMarzo and Sannikov (2006), He (2007) and DeMarzo, Fishman, He and Wang (2007) apply these methods in finance. Biais et al (2007) present a connection between discrete and continuous-time methods by a limit argument.

While optimal contracts under dynamic moral hazard are standard, we need to develop a significant methodological innovation to allow for adverse selection. Pure moral hazard settings are tractable due toe the *one-shot deviation principle*, which states that full dynamic incentive compatibility follows whenever the agent has incentives against deviating in just one period. This principle does not work in our setting due to the combination of adverse selection and the agent's ability save secretly. The difficulty is that the agent's incentives and his optimal strategy depend on the type of his project. Moreover, the agent's incentives depend on his past actions and the amount of cash he has available.

Despite these difficulties, we are not only able to solve the problem, but we find a remarkably simple optimal contract. This result, which provides a convenient benchmark to think about credit lines in practice, is surprising. Anyone familiar with dynamic contracts would expect moral hazard and adverse selection to imply such a complex optimal contract that it would be useless for deriving any economic intuitions.

Why is it that that our optimal contract takes such a simple form: a credit line with a credit limit that increases *deterministically* in time, which gives the agent full discretion to draw

and deposit funds in any way he wants? The main reason for simplicity lies behind the great diversity of strategies the agent may follow in our model. Specifically, we allow the agent to manipulate the project's cash flows using his own unobservable savings.[3] While in the real world managers can act in many more ways than any model can describe, this assumption naturally captures the managers' ability to manipulate cash flows by making adjustments in inventories, accelerating payments from customers and borrowing from outside sources. Because the agent can act in many ways to "game" contractual rules, a simple contract becomes optimal. It is a lot harder to give the agent the right incentives with a complex contract, which tempts the agent to find smart deviations to exploit the contractual intricacies. As a result, incentive provision in a complex contract always leads to higher inefficiency than in our simple optimal contract. [4]

Our result is especially attractive because a credit line is not only a simple, but also an *incomplete* contract, which happens to be dynamically optimal. It gives the agent discretion to use available credit but does not spell out the agent's actions. Although the contract is simple, it implements complex dynamic behavior, as the agent has incentives to borrow only to cover losses and pay interest on outstanding balance. This form of a contract creates a connection with contracts used in practice, which cannot physically spell out every contingency, and with the incomplete contracting literature, which studies how simple rules can implement desirable contractual outcomes. For example, Grossman and Hart (1986) argue that the allocation of ownership rights over productive assets between

---

[3] While hidden savings typically make agency problems nearly intractable, on rare occasions they lead to simple and elegant results. Bulow and Rogoff (1989) show that international debt cannot be sustained by reputation alone when countries have access to savings. In a similar spirit, Cole and Kocherlakota (2001) show that in the presence of hidden savings, the best insurance contract against random income shocks relies exclusively on saving and borrowing.

[4] A couple of papers, like Battaglini (2005) and Tchistyi (2005), find tractable models in which the agent's private information is correlated over time. Those settings have a particular form of adverse selection, because when the agent intentionally misreports his current information to the principal, their beliefs about future uncertainty diverge. Closer in spirit to our paper, Tchistyi (2005) finds that an agent must pay higher interest on debt when he is closer to default, a feature known as performance pricing.

two parties helps when optimal production decisions are ex-ante indescribable. Similarly, Hart and Moore (1998) derive an implementation of a borrowing contract that involves renegotiation if the agent fails to pay off debt.

This paper is organized as follows. Section 2 presents the model. Section 3 discusses the benchmark contract, a credit line with a fixed credit limit, which is optimal when moral hazard is the only problem. Section 4 presents an optimal contract with adverse selection. Section 5 describes the new method we develop to justify the optimal contract. Section 6 shows that our results are robust with respect to many important practical concerns, as the agent's ability to borrow from outside investors and the lender's ability to repossess the assets in the event of default. Section 7 concludes.

2. The Model

A risk-neutral agent with limited wealth seeks help from outside investors to run a project. Without outside investment, the agent can derive value R_0 from his financial wealth and the project's assets. The project can be made more profitable if outside investors provide additional capital K and commit funds to cover the project's losses until liquidation. In that case, the project runs continuously until time T, unless it is liquidated prematurely according to the contract between the agent and the investors. Until liquidation the project generates cash flows

$$dX_t = \mu \, dt + \sigma dZ_t,$$

where $\mu > 0$ is the expected cash flow rate and $\{Z_t, t \geq 0\}$ is a standard Brownian motion.[5] At the time of liquidation, the investors obtain a value of L from the project's assets. We normalize the value that the agent gets during liquidation to 0, that is, the agent loses a payoff of R_0 he could obtain before financing.[6] We can interpret R_0 broadly as the amount

---

[5] The model can be easily generalized to let the agent also receive a non-pecuniary benefit from the project, in addition to the cash flows X. For example, see Piskorsky and Tchistyi (2006) who study a related model under pure moral hazard applied to mortgages. In their context, a homeowner receives a non-pecuniary benefit from living in a house.

[6] Our model can be easily extended to allow the agent to retain a fraction of the assets during liquidation. This issue is of practical importance, because lenders are not always able to repossess all of the borrower's assets in the event of default. See Section 6.

by which the principal can harm the agent by repossessing the assets that secure the loan, compromising the agent's credit rating, involving the agent in time-consuming and costly legal procedures or by other means.[7] Assume that liquidation before time T is inefficient, that is $rL < \mu$, where r is the market interest rate.

There are two informational problems, moral hazard and adverse selection. The moral hazard problem is that neither the cash flows nor the agent's savings are observable to the investors. The agent can secretly divert cash flows to consume or save, and he can fabricate cash flows that did not actually realize from savings.

The source of adverse selection is that the agent, who owns the project before he contracts with the investors, has private information about their productivity. This makes investors concerned about funding a bad unprofitable project. We assume that bad projects generate cash flows and/or private benefits to the agent at an expected rate of $\mu' < \mu$ (for example, according to $dX_t = \mu' dt + \sigma' dZ_t$). The owner of a bad project may imitate an agent with a profitable project by contributing assets R_0 at time 0, in order to obtain a capital investment of K. We assume that the agent with a bad project may have access to unreported sources of wealth, which allow him to fabricate any cash flow path to "buy" more time until liquidation.[8]

We consider market conditions, under which the contract between the principal and the agent must "screen out" bad projects (rather than pool both types of projects together, or bribe agents with bad projects not to finance them). For example, bad projects must be screened out if investors act competitively (see Rothchild and Stiglitz (1976)), or if the proportion of bad projects in the economy is sufficiently large.

The principal can fully commit to any history-dependent contract. By the revelation principle we can focus on truth-telling contracts. Such a contract specifies a termination

---

[7] In the context of micro-lending in developing countries, loans are often made to a group of people who know each other well. A default by one person makes the remaining members of the group responsible for that person's obligation. In this context, the level of R_0 for each borrower is the value of social ties with the remaining members of the group, which become harmed in the event of default.
[8] While both types of agent may have hidden sources of wealth at time 0, only the bad type may benefit from misreporting his wealth. The agent with a good project has incentives to report all wealth to the principal at time 0. By doing so, he creates a possibility for a better contract, with a lower risk of inefficient liquidation.

time $\tau(\hat{X}) \in [0, T]$ and cumulative transfers from the agent to the principal

$\{D_t(\hat{X}), t \in [0, \tau(\hat{X})]\}$ as functions of the agent's reports $\{\hat{X}_s, s \geq 0\}$. Formally, $D(\hat{X})$ is

random process and $\tau(\hat{X})$ is a stopping time progressively measurable with respect to the

agent's reports. We assume that the agent's discount rate is the same as the interest rate on

savings, and so without loss of generality, the agent postpones consumption until

liquidation. Then, formally, the agent's savings evolve according to

$$dS_t = rS_t dt + dX_t - dD_t.$$

The agent's savings must stay nonnegative.

Denote by $W_0 > R_0$ the value, determined by the relative bargaining powers of the agent and

the investor, which an agent with a good project must get from his contract. The problem

of an investor is to find a contract (D, τ) that maximizes profit

$$E\left[\int_0^\tau e^{-rt} dD_t + e^{-r\tau} L \right]$$

subject to three sets of constraints: (1) if the agent with a good project reports truthfully, he

maintains nonnegative savings until time τ and gets value

$$E\left[e^{-r\tau} S_\tau \right] = W_0,$$

(2) if the agent with a good project follows any other reporting strategy $\hat{X}$ that keeps his

savings nonnegative until time τ, he gets value

$$E\left[e^{-r\tau} S_\tau \right] \leq W_0,$$

and (3) if the agent with a bad project follows any reporting strategy $\hat{X}$, he gets value

$$E\left[\int_0^\tau e^{-rt} (\mu' dt - d\hat{X}_t - dD_t) \right] \leq R_0.$$

We say that a contract is *truth-telling* if it satisfies constraint (2) and *screening* if it satisfies constraint (3). Note that the contract that we seek, intended to fund agents with a profitable project and limited initial wealth R_0, must screen out all bad agent regardless of the financial flexibility they have to temporarily fabricate cash flows.

To solve the principal's problem, it is enough to consider only contracts where the agent chooses to keep zero savings, as shown in the following lemma (proved in the Appendix).[9]

Lemma 1. *Given any truth-telling contract, there is another truth-telling contract with the same value to the principal and the agent with a good project, in which the agent does not save until liquidation, when he may receive a payment from the principal. If the former contract screens out bad project, the latter contract can be constructed to do so as well.*

Using the terminology of Hart and Moore (1994), we focus on contracts with the *fastest* repayment path by making the agent to transfer all cash flows to the principal until liquidation. Thus, we consider only contracts with $D_t = \hat{X}_t$ for $t < \tau$.[10] There are many reasons outside our model why such contracts are most attractive. For example, when managers pay out extra cash to investors, it prevents them from empire building and investing in pet projects.

3. The optimal contract under pure moral hazard.

Before addressing the problem of adverse selection and moral hazard, we discuss a benchmark optimal contract under pure moral hazard in this section. This contract can be found by standard methods, which were developed in Green (1987), Spear and Srivastava (1987), Abreu, Pearce and Stacchetti (1990) and, in continuous time, Sannikov (2004). The key insight of these methods is to summarize the agent's incentives using his continuation value, the agent's future expected payoff when he follows the recommended (truth-telling) strategy

---

[9] A simplification of this type is standard in problems with pure moral hazard and hidden savings. For example, see Werning (2002), Kocherlakota (2004) and Williams (2005).

[10] As a mathematical curiosity, the optimal credit line, being an incomplete contract, implements any optimal repayment path.

$$W_t = E_t \left[e^{-r(\tau - t)} S_\tau \mid \hat{X}_s = X_s \right].$$

The optimal contract is found using dynamic programming with W_t as a state variable. For problems of this sort, it is typical for the agent's incentive constraint to be just binding in the optimal contract, since the provision of incentives is costly.

When there is adverse selection, the contracting problem becomes significantly more complicated. Standard methods no longer apply, because one has to take into account the incentives of the agent with a bad project. Therefore, to allow for adverse selection we develop significant innovations over the standard methods. Section 5 discusses in depth our new methods.

For our model under pure moral hazard, DeMarzo and Sannikov (2006) derive the optimal contract using the standard methods. We do not provide a full derivation, but present the contract and explain the key reasons that make this contract incentive-compatible and optimal.

Consider a contract in which the agent gets a credit line with limit μ/r, interest rate r and a starting balance of $M_0 = \mu/r - W_0$. Then the balance evolves according to

$$dM_t = rM_t dt - d\hat{X}_t,$$

where $\hat{X} = D$ are transfers to the principal that are chosen by the agent himself. The balance M_t can be positive or negative, but if it reaches the credit limit before time T, the project is immediately liquidated without a payment to the agent. If the project survives until time T, the agent can collect $W_T = \mu/r - M_T$ from the principal.[11] Before termination the agent has full discretion to transfer funds to and from the credit line.

---
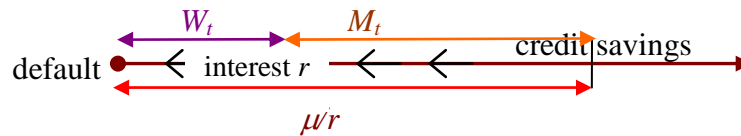
[11] This payment can be interpreted as the value of the collateral, which comes from funds R_0 that the agent contributes and an initial draw on the credit line.

Figure 1: Credit line.

We make two claims about this contract:

1. The agent is indifferent between all strategies. In particular, it is optimal for the agent to pay all cash flows to the principal until termination.

2. This contract is optimal for the principal if the agent follows the latter strategy.

The see that the first claim is true, note the agent can always "cash out" by drawing the remaining credit and triggering termination. Moreover, the agent is always indifferent between cashing out and continuing the project, independently of his current balance or actions, because

- the interest rate on the credit line is the same as on the agent's savings account and

- the interest on the maximal amount of credit equals the expected rate at which cash flows arrive

In particular, the agent maximizes utility by setting $\hat{X} = X$ until liquidation. Note that the remaining credit $W_t = \mu/r - M_t$ is the agent's future expected payoff at time t.

To justify the optimality of this contract, by Lemma 1 we only need to consider alternative contracts in which the agent is required to transfer all reported cash flows to the principal until liquidation, when the principal may make a transfer back to the agent.[12] The key property behind the optimality of the credit line is that the agent has *just enough* incentives to tell the truth. Other truth-telling contracts amplify the agent's incentives, by offering him more than a dollar of future payoff for each dollar reported to the principal. Such amplified incentives precipitate default in the event of losses and delay default in the event of gains relative to the credit line contract. Overall, this causes inefficiency because the likelihood of default is more sensitive to cash flows in the event of losses than in the event of gains.

---

[12] Lemma 1 in the Appendix shows that for any truth-telling contract, there is a payoff-equivalent contract, in which the agent transfers all cash flows to the principal until termination.

The formal argument behind optimality relies on dynamic programming. It turns out that the principal's value function

$$b(W,T-t) = E_t\left[\int_t^\tau e^{-r(s-t)}\mu ds + e^{-r(\tau-t)}L - e^{-r(\tau-t)}W_\tau \mid W_t = W\right],$$

his future expected profit when the remaining lifetime of the project is $T - t$, is concave in the agent's the agent's "continuation value" $W_t = \mu/r - M_t$ (see Lemma 2 in the Appendix). This property formalizes the fact that default is more sensitive to cash flows in the event of losses than in the event of gains. From the law of motion of M_t, the agent's continuation value W_t in the credit line contract follows

$$dW_t = rW_t dt + (d\hat{X}_t - \mu dt),$$

so the agent's marginal payoff from each dollar of reported cash flows is exactly a dollar. There are other contracts that give the agent incentives to tell the truth, in which the agent's continuation value evolves according to

$$dW_t = rW_t dt + \beta_t(d\hat{X}_t - \mu dt),$$

where $\beta_t \geq 1$ is the marginal value the agent gets from reporting cash. The condition $\beta_t \geq 1$ is necessary for incentive compatibility, but not sufficient.[13] However, since the function $b(W,T - t)$ is concave in W, it follows that a contract with $\beta_t > 1$ is suboptimal, since extra variance in W would hurt the principal. Figure 2 shows a typical form of $b(\cdot,T)$.

---

[13] It would be sufficient for incentive compatibility if the agent did not have hidden savings.
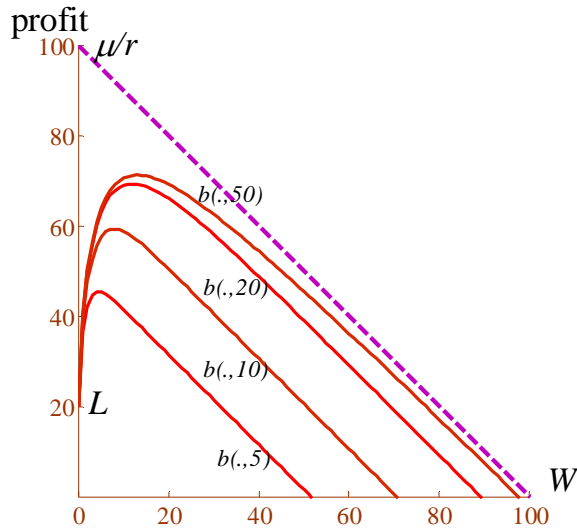
Figure 2: The principal's profit as a function of the agent's value/credit line balance for several values of project lifetime T ($\mu = 10$, $r = 0.1$).[14]

To get a better sense of the significance of the preceding arguments, let us discuss the agent's incentives and the principal's profit in several alternative contracts:

1. A credit line with a fixed credit limit greater than μ/r and interest rate r

2. A credit line with a credit limit μ/r and interest $\gamma > r$ on positive balances

3. A credit line with a fixed credit limit smaller than μ/r and interest rate r

Contracts 1 and 2 do not provide the agent with proper incentives. When the credit limit is greater than μ/r, the agent prefers to draw the credit line and default immediately, since the expected cash flows fall behind the interest on the full amount of credit. The same is true when the credit limit is μ/r but the interest rate is greater than r. In contrast, in contract 3 the agent gets amplified incentives to pay down credit line balance, since drawing credit line to default gives him less value than avoiding default by paying interest. Therefore, contract 3 is truth-telling, but suboptimal.

---

[14] The computation of this figure relies on the following explicit formula: the probability that the project survives until time t is $\Phi(2re^{rt}W_0/(e^{2rt}-1)) - \Phi(-2re^{rt}W_0/(e^{2rt}-1))$, where Φ is the CDF of standard normal distribution.

3.1.   Connection with empirical evidence.

In this subsection we relate our model to empirics. With the optimal contracting approach, one can be more sure of empirical predictions than if one simply compares contracts created by specific securities. Indeed, if one finds that contract A is better than contract B, the empirical predictions about securities used in these two contracts cannot be taken seriously if both of these contracts are far from the optimal one. One needs to know what the optimal contract is, and which of its features matter the most for overall profitability. The features of optimal contracts that have stronger effect on profitability are more likely to be found empirically.

For example, Sufi (2007) finds that firms with larger levels of cash flows open longer lines of credit with banks. However, the positive relationship between cash flow levels and access to credit exists only for firms with a high probability of financial distress, according to Altman's *z-score* (1968). At the same time, our model predicts a credit limit of μ/r, which increases with the level of cash flows for all firms. How can we reconcile this finding with empirical evidence? The reason is that the choice of the credit limit matters the more for the value of the firm when the firm is closer to financial distress.

Consider a firm manager who is considering opening a line of credit to have the option to draw against future potential losses. Figure 3 shows the relationship between the firm's level of cash balances and the payoff that the agent gets if he opens an optimal line of credit with competitive investors. Observe that the agent's value is concave in the level of cash, so that the agent loses value by waiting. Thus, a firm can benefit from opening a line of credit even when it is not currently cash-constraint, and may never need to borrow cash. However, the benefit from opening a line of credit diminishes with the levels of cash. It follows that firms that are far from financial distress do not care as much about having a credit line of optimal length.
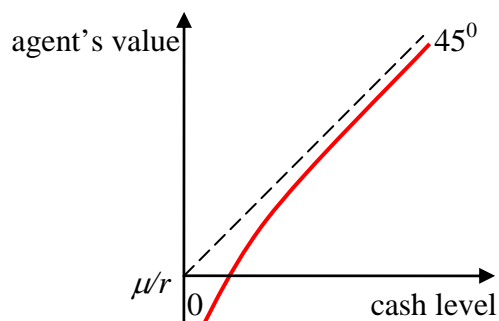
Figure 3: Manager's payoff as a function of the firm's cash level.

Note also that the slope of the agent's value function becomes steeper as he runs out of cash. Thus, the agent's incentives are stronger than those created by the optimal contract. Our discussion of the optimal contract implies that amplified incentives create inefficiency.

The next section, where we add adverse selection to the model, is not **at all** a simple extension of the pure moral hazard setting from the point of view of theoretical analysis. In that setting, the agent's incentives in the optimal contract are significantly different than under pure moral hazard (at least initially while the credit limit is growing). Under adverse selection the marginal value that the agent gets for each dollar of realized cash flows is more than a dollar. How much exactly – that is determined by the fairly complex analysis that justifies the optimal contract. We presented the optimal pure moral hazard contract first mainly because it gives the *eventual* form of the contract under adverse selection when the credit limit is extended fully. Also, with the knowledge of the standard moral hazard model, we can make comparisons and explain the technical innovations needed to analyze adverse selection a lot more clearly in Section 5.

4. The Optimal Contract under Adverse Selection and Moral Hazard

In this section we conjecture the optimal contract under adverse selection and moral hazard. We make brave guesses to derive a remarkably simple contract. We verify that this contract creates the right incentives: for the good agent to tell the truth and for the bad agent not to get financing. Section 5 verifies that this contract is optimal.

The simplicity of the contract may mislead the reader into thinking that the result is trivial (or, worse, an unsurprising extension of DeMarzo and Sannikov (2006)). Thinking this way, one forgets that we consider the full space of complex history-dependent contracts, and that the incentive constraints under adverse selection are nonstandard and significantly more complicated than under pure moral hazard. In fact, after conjecturing the simple optimal contract, we would be crazy to hope that it is optimal given what we know about dynamic contracts. Existing literature finds complex optimal contracts even under pure dynamic moral hazard.[15] If our simple contract does turn out to be optimal, there must be something special about the interaction of moral hazard and adverse selection in our model.

Why, despite the odds, does our simple contract turn out optimal? The fundamental reason has to do with the great flexibility the agent has to manipulate cash flows with hidden savings when the cash flows follow fuzzy Brownian paths. It is because the agent can deviate in so many ways that the optimal contract is simple. In a complicated path-dependent contract, where the agent may try to "game rules" using a smart strategy, it is significantly harder to provide the agent with incentives. Thus, the simple contract creates the agent's incentives at the lowest cost.

This reason for the optimality of an increasing contract in our setting is not at all the same as the reason why a simple credit line is optimal in the pure moral hazard setting of DeMarzo and Sannikov (2006). In their paper, a credit line is optimal even if the agent cannot save. The agent is risk-neutral, and he gets a benefit of exactly $\lambda \in (0, 1]$ for each unit of diverted cash flows. This simple structure makes the credit line optimal in DeMarzo and Sannikov (2006). The agent's flexibility to manipulate cash flows with hidden savings is not at all important for the optimality of a simple credit line in DeMarzo and Sannikov (2006).

We feel intuitively that one of the reasons why simple contracts are used in practice is that the agent has a huge set of strategies he can follow to affect the output. Complicated

---

[15] A couple of exceptions are Holmstrom and Milgrom (1987), who choose the features of the model very carefully to derive their linearity result, and DeMarzo and Sannikov (2006), who assume risk neutrality and the specific form of moral hazard.

contractual clauses would present the agent with opportunities to game rules. This intuition is hard to capture theoretically, because it is difficult to build into a model a sufficiently rich set of strategies for the agent. Our model presents this intuition in a nice way.

4.1. Conjecturing the optimal contract.

Does the optimal contract under pure moral hazard, a credit line with a fixed credit limit μ/r, succeed at screening out bad projects? No. Since $W_0 > R_0$, an agent with a bad project would happily accept financing with such a credit line, and draw the entire balance to default immediately. To screen out bad projects, a contract needs to limit the amount of available funds so that an agent with a bad project finds such a contract unacceptable, even if he has infinite wealth. Specifically, there should be no history of cash flow reports $\{\hat{X}_s, 0 \leq s \leq t\}$, which allows the project to survive until time t, and has value less than the expected value of bad project's cash flows until time t minus R_0. In other words, the project has to be liquidated not later than the first time when

$$\underbrace{\int_0^t e^{-rs} \mu' ds}_{\text{expected value of cash flows}} = \underbrace{R_0 + \int_0^t e^{-rs} d\hat{X}_s}_{\text{funds transferred to investor}} \ .$$

The following lemma, proved in the Appendix, summarizes this finding.

Lemma 3. *Let*

$$R_t = e^{rt}\left(R_0 + \int_0^t e^{-rs} d\hat{X}_s - \int_0^t e^{-rs} \mu' ds \right).$$

A contract fails to screen out bad projects if for some history of reports, for which $R_t < 0$, the project still survives, or the agent gets a payment greater than R_t during termination. A contract succeeds to screen out bad projects if liquidation happens before R_t becomes negative and the agent receives a payment of at most R_t during liquidation.

We conjecture that the optimal contract under both informational problems arises if we *superimpose* the adverse selection restriction of Lemma 3 onto the optimal pure moral hazard contract. That is, we give the agent a credit line with a starting balance M_0 that depends on the agent's value W_0, and liquidate when R_t reaches 0 or when M_t reaches μ/r, whichever happens sooner. It is not hard to verify that this simple and intuitive contract creates proper incentives to both types of agents (as we do in Lemma 4 at the end of this section). However, it is truly unexpected and remarkable that this contract is *optimal*. Due to adverse selection, it is completely unclear how to check the optimality of this contract with standard methods. We develop a new approach to verify that this contract is optimal in the next section. Fundamentally, the key reason for the optimality of this contract lies behind the great flexibility the agent has to manipulate cash flows.

Let us show that R_t reaches 0 when M_t reaches a critical level $\bar{M}_t$, so that the contract we obtain is a credit line with credit limit $\min(\mu/r, \bar{M}_t)$. Since the expression for R_t includes the value of cash flows that the agent reports, the credit line balance can be expressed in terms of R_t

$$ dM_t = rM_t dt - d\hat{X}_t \Rightarrow M_t = e^{rt}\left(M_0 - \int_0^t e^{-rs} d\hat{X}_s \right) = e^{rt}\left(M_0 + R_0 - \int_0^t e^{-rs}\mu' ds \right) - R_t. $$

Therefore, R_t reaches 0 when the balance M_t reaches

$$ \bar{M}_t = e^{rt}\left(M_0 + R_0 - \int_0^t e^{-rs}\mu' ds \right). $$

When $M_0 + R_0 > \mu'/r$, this credit limit increases in time and reaches μ/r in a finite moment of time T^*. Figure 5 shows a typical path that the credit limit takes.
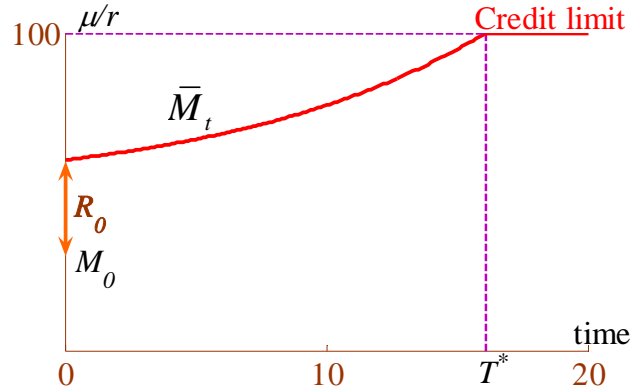
Figure 5. The evolution of the credit line when $M_0 = 30$, $R_0 = 30$, $\mu = 10$, $\mu' = 5$ and $r = .1$.

One prominent property of this contract is that the credit limit evolves deterministically in time, independently of the agent's actions. The reason for this simple property is that before time T^*, the contract makes funds available to the agent up until the limit where the agent with a bad project is just indifferent between taking financing or not. The point of indifference depends only on the lifetime of the bad project and outstanding debt, and not the movement of funds between the credit line and savings, which have the same interest rate. Therefore, the credit limit before time T^* does not depend on the agent's actions. The agent with a bad project is indifferent between not funding the project and funding it to follow any course of actions that leads to default not later than time T^*.

We are ready to formulate our main result, which is justified in the next section.

Main Theorem. *The optimal under adverse selection and moral hazard takes the form of a credit line with interest rate r and credit limit $\bar{M}_t$ that grows to μ/r until time T^*, and stays at value μ/r thereafter. If the balance M_t reaches the credit limit before time T, liquidation occurs immediately and the agent gets no payment from the principal. If the project survives until time T, the agent gets a payment of $\mu/r - M_T$.*

We will refer to this contract as the *increasing credit line* contract for values (W_0, R_0). Unlike in a credit line with a fixed limit μ/r, the precise relationship between the starting balance M_0 and the value of the agent with a good project W_0 is somewhat complicated. Since an initial restriction on funds available to the agent hurts his value, $W_0 < \mu/r - M_0$. We

denote by $U(R_0,T^*)$ the function that gives the agent's value W_0 as a function of the amount of funds R_0 initially available to the agent and the amount of time T^* it takes the credit line to increase fully. Function U plays an important role in our justification of the optimal contract: it tells us the continuation value $W_t = U(R_t,T^*-t)$ of the agent with a good project at any time $t \leq T^*$ and his marginal value of cash flows $U_1(R_t,T^*-t) > 1$. That is, the agent gets *more* than a dollar of payoff for each dollar of cash flows he receives before time T^*, because a reported dollar not only reduces balance, but also allows the credit line to grow longer. The reason for this property is adverse selection: extra funds help the agent signal good project quality by paying interest on the outstanding balance. Figure 6 shows a typical form of function U, which is increasing in T^* and concave in R_0. The range of W_0 for which the principal's problem has a solution is obtained by varying T^* from 0 to T.
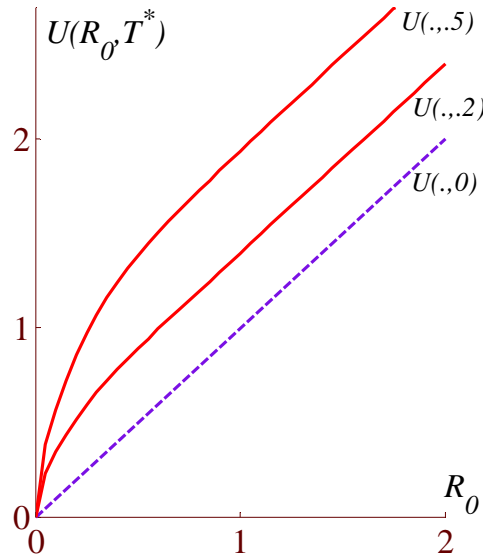
Figure 6: The agent's value as a function of initially available funds R_0, for $T^*=0$, 2 and 5
$(\mu = 10,\ \mu' = 8,\ r = 0.1)$.[16]

In the remainder of this section, we only demonstrate that agents with both types of projects have proper incentives in the increasing credit line contract. Regarding bad projects, Lemma 3 already showed that the increasing credit line contract successfully screens them out. Lemma 4, which is proved in the appendix, shows that this contract also gives the agent with a good project incentives to tell the truth.

---

[16] This figure is computed using Monte Carlo simulations.

Lemma 4. *In the increasing credit line contract for values (W_0, R_0) the agent with a good project gets value W_0 from any strategy, under which he never reaches default before time T^* with positive savings. If the agent sometimes ends up in default with positive savings before time T^*, he gets value less than W_0 in expectation.*

In particular, Lemma 4 implies that if the agent makes truthful reports, he maximizes his expected payoff under the increasing credit line contract and obtains a value of W_0. The proof of Lemma 4 relies on three observations. The first two imply that any strategy in which the agent exhausts his savings before time T^* achieves the same payoff for the agent.

1. Conditional upon avoiding default until time T^*, the payoff that the agent gets for his money is independent of his strategy, as if the credit limit was μ/r throughout.

2. Before time T^* default is determined only by the total value of reported cash flows and not the timing of reports.

Therefore, the time of default before time T^* (as a function of true cash flows) and the agent's payoff if the project survives until time T^* are the same for any strategy that exhausts the agent's savings to report cash flows in the event of a default before time T^*.

Lastly, it is strictly optimal for the agent to use savings fully in order to avoid default before time T^*:

3. When the balance on the credit line approaches $\bar{M}_t$ and $t < T^*$, the agent strictly prefers to avoid default (e.g. by paying interest) whenever he has funds. The reason is that expected cash flows exceed interest payments, and in addition the agent gets an extension to the credit line.[17]

5. Justification.

This section proves that the increasing credit line contract described in Section 4 is optimal. As we discussed in Section 3, standard optimal contracting methods do not apply

---

[17] Here the logic is the same as in the previous section, in which we argued that the agent gets amplified incentives in a credit line with a limit shorter than μ/r.

to our setting because of adverse selection. Therefore, we develop a new method to verify the optimality of our contract. Our verification procedure takes several steps.

1. For an *arbitrary* contract, **define** key dynamically evolving variables.

2. **Formulate** conditions that these variables must satisfy in any contract that motivates the agent with a good project to report cash flows truthfully and screens out bad projects. These are *necessary* incentive constraints. They hold for our conjectured optimal contract, which provides complete incentives for the agent to tell the truth, and screens out bad projects.

3. **Verify** that among the contracts that satisfy these necessary incentive constraints, our conjectured optimal contract attains the maximal profit for the principal.

Clearly, these steps do not provide a recipe, but just a set of guidelines to follow for problems of this kind. To implement these steps, one has to be creative and to face substantial uncertainty about whether initial guesses ultimately lead to a successful argument. First, one has to guess correctly which variables can be used to formulate the incentive constraints that generate inefficient liquidation in the optimal contract. Second, to formulate the necessary incentive constraints one has to find a *dynamic* strategy that gives the agent a higher payoff than telling the truth each time the constraints are violated. Third, to verify the optimality of the conjectured contract, one has to characterize the principal's profit under this contract, and show that it cannot be improved upon while maintaining the necessary incentive constraints. This step involves a verification of inequalities on the second-order derivatives of the principal's profit function.

These steps are significantly more challenging to implement in comparison with the standard methods that have been developed for pure moral hazard. However, one should not be discouraged from tackling related important problems by the level of difficulty of our analysis. The greatest challenge is to develop a new method and carry out this type of analysis the first time through. For the next problem, one still faces a lot of uncertainty about making the right analytical choices, but at least one can borrow a lot from the general line of reasoning here. Eventually, after enough problems are solved, there will be

a standard method to understand dynamic contracting problems where other informational asymmetries interfere with the agent's incentives.

5.1.   The state variables.

For an arbitrary contract, define the following two state variables:[18]

- the *limit on funds*

$$R_t = e^{rt}\left(R_0 + \int_0^t e^{-rs} d\hat{X}_s - \int_0^t e^{-rs} \mu' ds \right)$$

that adverse selection imposes on the contract (see Lemma 3)

- the *continuation payoff* of the agent with a good project when he tells the truth,

$$W_t = E_t[e^{-r(\tau-t)} W_\tau],$$

where W_τ is the payment that the agent receives during liquidation. Recall that by Lemma 1, we restrict attention to contracts, in which the agent's savings are zero when he reports the cash flows truthfully.

In any contract, variables W_t and R_t evolve dynamically with the agent's reports $\hat{X}$. We need a representation of these variables in terms of $\hat{X}$ in order to formulate the necessary incentive constraints in the next section. Differentiating the definition of R_t, we find that

$$dR_t = rR_t dt + (d\hat{X}_t - \mu' dt).$$

The following lemma gives a representation for W_t:

Lemma 5. *Given any contract, there exists a stochastic process $\{\beta_t, t \le \tau\}$ in L^2 such that[19]*

---

[18] Following Fernandez and Phelan (2000), these two variables do not present a full recursive structure of our problem, as they do not summarize the agent's incentives to save. We focus on W and R only, as these variables allow us to formulate the necessary incentive constraints that we use to justify the optimality of the conjectured contract.

[19] L^* is the space of processes β for which $E[\int_0^t \beta_s^2 ds] < \infty$ for all $t < \infty$.

$$dW_t = rW_t dt + \beta_t (d\hat{X}_t - \mu dt).$$

Proof. Note that W_t is the agent's continuation value when he tells the truth, i.e. $\hat{X}_t = X_t$. Then $e^{-rt}W_t$ is a martingale. By the Martingale Representation Theorem, there exists a process $\{\beta_t, t \le \tau\}$ in L^* such that

$$d(e^{-rt}W_t) = e^{-rt}\beta_t dZ_t = e^{-rt}\beta_t(dX_t - \mu dt).$$

This expression is equivalent to the desired representation above. QED

As W_t and R_t evolve dynamically, the difference $W_t - R_t$ reflects the degree of adverse selection that remains to be dealt with. While initially $W_0 > R_0$, W_t and R_t must become equal at time τ or sooner. Indeed, W_t and R_t follow continuous paths, and the agent cannot get a payment of $W_\tau > R_\tau$ during liquidation by Lemma 3. When W_t becomes equal to R_t for the first time, the adverse selection problem goes away, since the continuation value of the good agent no longer exceeds the upper bound on the value of the bad agent. At that time, it is optimal for the "continuation contract" to be a pure moral hazard contract with value $W_t=R_t$. Such a continuation contract would give the bad agent value of at most R_t, and would also minimize incentives of the good type to build up savings to manipulate cash flows after time t. (The good agent gets exactly a dollar of value for each dollar of saved cash flows after time t). From now on, we can assume that the contract becomes a pure moral hazard contract from the time when W_t and R_t become equal, which always happens before liquidation.

5.2.   A necessary incentive-compatibility condition.

In this subsection we formulate a necessary condition for truth-telling on the marginal value β_t that the agent gets from each realized dollar of cash flows. Lemma 5 defines the process β_t for an arbitrary contract, which depends on the entire past history of cash flows.

In the increasing credit line contract, β_t is completely determined by W_t and R_t, which are in turn determined by the current credit limit and balance. Ito's lemma implies that the value of β_t in the increasing credit line contract as a function of W_t and R_t is given by

$$\sigma(W, R) = U_1(R, s).$$

At time T^* when the full credit limit is reached, $W_{T^*} = R_{T^*}$ and $\beta_{T^*} = \sigma(W_{T^*}, R_{T^*}) = 1,$ as in the optimal contract under pure moral hazard. Before time T^*, $W_t > R_t$ and $\beta_t = \sigma(W_t, R_t) > 1$. Indeed, the agent gets more than a dollar of value for each dollar of cash flows before time T^*: reported cash flows not only reduce the balance, but also allow the credit line to grow longer.

In an arbitrary truth-telling screening contract under adverse selection, how is β_t related to W_t and R_t? If the agent does not report a dollar of cash flows when it arrives, he can save it and use it to fabricate cash flows when the time is convenient. Hidden savings give the agent great flexibility to manipulate reports in many ways. The increasing credit line contract limits the agent's opportunities to gain through cash flow manipulations, because *the agent's payoff is the same as long as he never reaches default with positive savings before time T^**. This important point worth repeating: an increasing credit line counts balances in such a straightforward way that the agent gains the same value from an entire range of reporting strategies. This simplicity minimizes the agent's benefit from diverting a dollar. In an alternative more complicated contract, if the agent diverted a dollar, he would seek out a way to gain the most from it by a sophisticated reporting strategy. As a result, it turns out that the agent's marginal value from realized cash flows, as a function of W_t and R_t, is the lowest in the increasing credit line contract of any other contract.

This intuition motivates our necessary condition on β_t for an agent to tell the truth in an arbitrary screening contract.

Proposition 1. *In any truth-telling screening contract $\beta_t \geq \sigma(W_t, R_t)$.*

The formal proof is in the Appendix. Unlike in standard dynamic contracting problems, where one derives the incentive constraints by looking at one-shot deviations, to justify the

necessary condition $\beta_t \geq \sigma(W_t, R_t)$ one has to look at fully dynamic strategies of cash flow manipulations. For an arbitrary screening contract that violates the necessary incentive-compatibility constraint, one has to find a dynamic strategy that gives the agent a payoff greater than W_t. Thus, we develop fundamentally new methods to prove Proposition 1.

Although we need to deal with complicated strategies to prove Proposition 1, our basic idea is simple. To find a profitable deviation, the agent has to divert cash when it is beneficial to do so (i.e. when β_t is low) and to fabricate cash flows otherwise (when β_t is high). Integrating the benefits that the agent gets over time, and taking the expectation, we show that a profitable deviation from truth-telling exists whenever $\beta_t < \sigma(W_t, R_t)$ on a set of positive measure.

5.3. Optimality.

In this subsection to verify that the increasing credit line contract is optimal. Denote by $f(W, R, s)$ the profit from the increasing credit line contract for values (W, R) when the lifetime of the project is s. We show any contract that satisfies the necessary incentive-compatibility condition $\beta_t \geq \sigma(W_t, R_t)$ achieves profit of at most $f(W_0, R_0, T)$. Technically, our verification argument uses properties of the second-order partial derivatives of f, which we derive analytically in Lemma 6 in the Appendix from the definition of f. Proposition 2 carries out the verification argument using those properties.

To see the logic behind our verification argument, let us draw analogy with the standard case of pure moral hazard. The necessary incentive-compatibility condition $\beta_t \geq \sigma(W_t, R_t)$, which holds with equality in the increasing credit line contract, is analogous to the condition $\beta_t \geq 1$ that holds with equality in the optimal contract under pure moral hazard. Under pure moral hazard, we argued that it is inefficient to amplify the agent's incentives by setting $\beta_t > 1$, because the principal's profit $b(W,s)$ is concave in the agent's continuation value W. In other words, giving the agent access to *more funds* in the event when he reports good cash flows, and *taking away funds* in the event that he reports losses causes inefficiency overall. The formal verification argument relies on martingale techniques using the law of motion of W_t and Ito's lemma.

The case of adverse selection is based on similar logic. We must show that $f(W,R,s)$, the profit from the increasing credit line contract for values (W,R) when the lifetime of the project is s, is the maximal profit that the principal can achieve. That is, we need to show that a contract with $\beta_t > \sigma(W_t, R_t)$ achieves lower profit than the increasing credit line contract. The basic structure of our verification argument is the same as for the case of pure moral hazard. However, because variable R_t evolves dynamically together with W_t, this conclusion follows formally from more complicated properties of the principal's value function than in the case of pure moral hazard, i.e. from the inequalities

$$f_{11}(W,R,s) \le 0 \quad \text{and} \quad \sigma(W,R)f_{11}(W,R,s) + f_{12}(W,R,s) \le 0,$$

as can be seen from the proof of Proposition 2. Lemma 6 in the Appendix introduces a number of technical innovations to prove these properties.

Intuitively, why is a contract with $\beta_t > \sigma(W_t,R_t)$ is suboptimal? Such a contract would take the agent further from default in the already good event of high cash flows at the expense of precipitating default in the already bad event of low cash flows. Overall, this causes inefficiency.

Proposition 2. *Consider an alternative truth-telling screening contract with value W_0 to an agent with a good project and R_0 to an agent with a bad project. The principal's profit from this contract is at most $f(W_0,R_0,T)$. Therefore, the increasing credit line contract is optimal.*

Proof.[20] Let $\hat{\tau}$ be the earliest time when $W_t = R_t$. As we argued earlier, $\hat{\tau} \le \tau$. In an arbitrary truth-telling screening contract, the processes W_t and R_t follow

$$dW_t = rW_t dt + (\sigma(W_t, R_t) + \varsigma_t)dZ_t \quad \text{and} \quad dR_t = (rR_t + \mu - \mu')dt + dZ_t$$

with $\varsigma_t \ge 0$. In an increasing credit line contract we have $\varsigma_t = 0$.

---

[20] The proof is technically quite involved, and may be skipped. The first part of the proof follows a standard dynamic-programming approach to show that the value function $f(W, R, T)$ is unimprovable, while the second half (Lemma 6) demonstrates a key property of f using non-standard techniques. The argument of the proof is very much connected with the verbal discussion at the beginning of the section, hence the choice to leave the proof in text.

Let the process $\{G_t,\ t \le \hat{\tau}\}$ be defined by $G_t = \int_0^t e^{-rs}\mu ds + e^{-rt} f(W_t, R_t, T - t)$.

In an increasing credit line, G_t is time-t expectation of the principal's profit, and thus a martingale, i.e. the drift of G_t is 0. Then Ito's lemma implies that

$$\mu - rf + rWf_1 + (rR + \mu - \mu')f_2 - f_3 + \tfrac{1}{2}\sigma(W, R)^2 f_{11} + \sigma(W, R)f_{12} + \tfrac{1}{2}f_{22} = 0.$$

In an alternative contract, G_t is a *submartingale* because

$$\mu - rf + rWf_1 + (rR + \mu - \mu')f_2 - f_3 + \tfrac{1}{2}(\sigma(W, R) + \varsigma_t)^2 f_{11} + (\sigma(W, R) + \varsigma_t)f_{12} + \tfrac{1}{2}f_{22} \le 0.$$

(We used the facts that $f_{11}, \sigma(W, R)f_{11} + f_{12} \le 0$, see Lemma 6 in the Appendix).

Therefore, the principal's profit from an alternative contract is less than or equal to

$$E[\int_0^{\hat{\tau}} e^{-rs}\mu ds + e^{-r\hat{\tau}} \underbrace{f(W_{\hat{\tau}}, W_{\hat{\tau}}, T - \hat{\tau})}_{=b(W_{\hat{\tau}}, T - \hat{\tau})}] = E[G_{\hat{\tau}}] \le G_0 = f(W_0, R_0, T).$$

This completes the proof. QED

6. Properties of the Optimal Contract.

In this section we discuss the applicability of our main results to various institutional environments. We argue that our result is holds strong to many adjustments to the model that are of practical concern, as the agent's ability to borrow from outside investors or the investors' ability to repossess all of the firm's assets in the event of liquidation. We also discuss the value of the assumption that an agent with a bad project has full flexibility to imitate any cash flow path from the point of view of robustness.

Competition between lenders is of practical concern. Our model applies both to environments in which investors compete to lend to the agent at time 0, and it is robust to environments in which the agent may secure additional sources of funding from outside investors after the contract with the original investor has been signed. At time 0, competition between investors is affects only the relative bargaining powers between the agent and the investors, which is reflected in the agent's starting value W_0. After time 0,

one may be worried that by borrowing from outside investors, the agent may make the original loan riskier. This should not be a concern for three reasons:

- since the project's assets secure the loan with the original investor, the agent would not be able to borrow again in the model because he has no new collateral

- even if the agent has additional assets, using them as a collateral to borrow from a new lender to avoid or delay liquidation only benefits the original lender[21]

- the contract is robust to the agent's ability to produce additional funds while the project is running, as long as the agent still maximizes his total discounted consumption. For example, see the constraint of Lemma 3 related to the ability of the bad type to produce extra funds.

Second, let us argue that if the agent is able to retain some of the project's assets in the event of liquidation, the optimal contract still takes the form of a credit line with an increasing credit limit. Of course, the credit limit and the starting balance become different, and overall the contract becomes less efficient.

Suppose that of the liquidating value of the project's assets $L = L_1 + L_2$, the investors can only repossess L_1, and L_2 goes to the agent. Then, under pure moral hazard the optimal credit line has length $\mu/r - L_2$. It makes the agent fully indifferent between all strategies, since the interest rate on the full amount of credit equals expected cash flow minus the opportunity cost of postponing L_2 for the agent (see Section 3). Under adverse selection, the project has to be terminated when the balance reaches $\mu/r - L_2$ or

$$R_t = e^{rt}\left(R_0 + \int_0^t e^{-rs}d\hat{X}_s - \int_0^t e^{-rs}\mu'ds\right)$$

reaches L_2, whichever happens sooner. This gives rise to an optimal contract of the same form as in our basic model: a credit line with a deterministically increasing credit limit, albeit with different starting balance and credit limits.

---

[21] Of course, if the agent and investor were aware of additional assets usable to secure the loan at time 0, they would be used in the optimal contract. A syndicate loan could be used if additional lenders are involved.

Lastly, let us discuss the assumption that the agent with a bad project has access to unlimited funds. It is important not to interpret this assumption literally, otherwise one may ask, why investor cannot repossess those additional funds in the event of liquidation? The assumption captures the notion of robustness, making the contract incentive-compatible even in the worst-case scenario. This issue is of particular concern when investors are competitive. If one of a population of investors offers a contract that does not screen a particular bad type of an agent, he is going to pick up the whole population of such bad types.

The assumption that allows the agent with a bad project to have unlimited funds also plays a simplifying role analytically, creating a deterministically increasing credit limit that can be computed easily. Even when one is doubtful of a specific assumption, the result can still be valuable as a benchmark for comparisons with other settings. For example, the result of Modigliani and Miller (1958) that the choice of capital structure has on effect on the firm's value has great value as a benchmark, even though the assumptions behind this result are far from true in practice. For our model, after one knows our optimal contract, one can think a lot better about related settings. For example, for the setting in which the agent with a bad project has limited funds, one may guess that the optimal credit limit would be increasing faster when the agent carries lower balances.

To conclude, this paper considers a simple dynamic principal-agent setting with moral hazard, adverse selection and hidden savings. We derive a remarkably simple optimal contract: a credit line with a credit limit that increases deterministically in time. This form of a contract resembles many arrangements used in practice, and can serve as a natural benchmark to understand them. Also, like actual contracts, a credit line is an *incomplete* contract. It specifies interest and credit limit, and gives the agent full discretion to use available credit. Although the specification of the contract is very simple, it implements a complex outcome, as liquidation and payments to the agent become functions of the entire past history of cash flows.

Appendix.

Proof of Lemma 1. Consider an arbitrary contract, in which the agent is required to make transfers $\{D_t(\hat{X})\}$ such that

$$\int_0^t e^{-rs} dD_s(\hat{X}) \le \int_0^t e^{-rs} d\hat{X}_s$$

as long as the project survives. When the agent has incentives to report truthfully, he consumes

$$S_{\tau(X)} = e^{r\tau(X)} \left(\int_0^{\tau(X)} e^{-rs} dX_s - \int_0^{\tau(X)} e^{-rs} dD_s(X) \right)$$

Consider an alternative contract, which requires the agent to transfer all reported cash flows to the principal until termination, when the principal gives the agent a payment of

$$e^{r\tau(\hat{X})} \left(\int_0^{\tau(\hat{X})} e^{-rs} d\hat{X}_s - \int_0^{\tau(\hat{X})} e^{-rs} dD_s(\hat{X}) \right).$$

Effectively, the principal modifies the contract to do savings for the agent.

This change extracts the same or greater amount of funds from the agent after all histories, which implies that the agent (with a good or a bad project) cannot become better off. Yet, if the agent with a good project tells the truth, he ends up with zero savings before liquidation, but the same amount $S_{\tau(X)}$ to consume at the liquidation time as before. Therefore, the modified contract gives the agent the same payoff as the old contract and incentives to report truthfully. Moreover, because the liquidation occurs at the same time $\tau(X)$ under both contracts, the principal's profit is the same. QED

Lemma 2. The function $b(W,T\text{-}t)$ is concave in W.

Proof. First, note that $b(W,T\text{-}t)$ is weakly decreasing in t since

$$b(W, T-t) + W = E_t \left[\int_t^\tau e^{-r(s-t)} \mu ds + e^{-r(\tau-t)} L \mid W_t = W \right]$$

and the project has less time to run in expectation for larger t. Moreover, since $b(0, T\text{-}t) = L$, it follows that $b_1(0, T\text{-}t)$ is weakly decreasing in t. Note that since

$\int_0^t e^{-rs} \mu ds + e^{-rt} b(W_t, T-t)$ is a martingale and $dW_t = rW_t dt + dZ_t$, Ito's lemma implies that

$b(W, T\text{-}t)$ satisfies the partial differential equation

$$rb(W, s) = \mu + rWb_1(W, s) - b_2(W, s) + \tfrac{1}{2} b_{11}(W, s).$$

Differentiating with respect to W, we obtain

$$0 = rWb_{11}(W, s) - b_{12}(W, s) + \tfrac{1}{2} b_{111}(W, s),$$

so $b_1(W_t, T-t)$ is a martingale and $b_1(W, T-t) = E_t[b_1(W_\tau, T-\tau) \mid W_t = W]$. Since $b_1(0, T\text{-}t)$ is weakly decreasing in t and $b_1(W, 0) = -1$ for all W, it follows that $b_1(W, T\text{-}\tau)$ is weakly decreasing in the initial condition $W_t = W$ for all paths of X. Therefore, $b_1(W, T\text{-}t)$ is decreasing in W and $b(W, T\text{-}t)$ is concave in W. QED

Proof of Lemma 3. Suppose that for some history of reports, the project survives until time t when $R_t(\hat{X}) < 0$. Then, if the agent with a bad project produces this history of reports, he collects an expected value of $\int_0^t e^{-rs} \mu' ds$ from the project's cash flows while transferring a value of $R_0 + \int_0^t e^{-rs} d\hat{X}_s$ to the principal. As a result, the agent's payoff is $-e^{-rt} R_t(\hat{X}) > 0$ and the contract fails to screen out bad projects. Similarly, if he gets a payment of $W_\tau > R_\tau$ after a history of reports $\hat{X}$, then agent with a bad project would get an expected payoff of

$e^{-r\tau} W_\tau - R_0 - \int_0^\tau e^{-rs} d\hat{X}_s > 0$ by fabricating that history.

Conversely, if liquidation always happens at or before the time when R_t reaches 0, and if the agent receives payment at most R_t at that time, then the value he gets from transfers is at most

$$e^{-rt}R_t - R_0 - \int_0^t e^{-rs}d\hat{X}_s.$$

Letting τ be the time when the project is liquidated for a given strategy of an agent with a bad project, his total expected value is at most

$$E[e^{-rt}R_t - R_0 - \int_0^t e^{-rs}d\hat{X}_s + \int_0^\tau e^{-rt}dX_t] = 0.$$

Therefore, the contract succeeds at screening out bad projects. QED

Proof of Lemma 4. Consider a strategy such that whenever liquidation happens before time T^*, the agent has zero savings. Then the value of reported cash flows must equal the value of true cash flows at the liquidation time $\tau(\hat{X})$ when $\tau(\hat{X}) < T^*$, i.e.[22]

$$\int_0^{\tau(\hat{X})} e^{-rs}d\hat{X}_s = \int_0^{\tau(\hat{X})} e^{-rs}dX_s \Rightarrow R_{\tau(\hat{X})}(\hat{X}) = R_{\tau(\hat{X})}(X) = 0.$$

Thus, $\tau(\hat{X}) = \tau(X)$ when $\tau(\hat{X}) < T^*$. Then the agent's payoff from strategy $\hat{X}$ is

$$E\left[e^{-rT^*}1_{\tau(X) \geq T^*}\left(\overbrace{\underbrace{\frac{\mu}{r} - e^{rT^*}\left(M_0 - \int_0^{T^*} e^{-rs}d\hat{X}_s \right)}_{\text{remaining credit } M_{T^*}(\hat{X})} + \underbrace{\int_0^{T^*} e^{r(s-T^*)}(dX_s - d\hat{X}_s)}_{\text{savings } S_{T^*}(\hat{X})}}^{W_{T^*}(X)} \right) \right] = W_0.$$

Now, consider an alternative strategy, under which there is a history of reports $\{\hat{X}_s, s \in [0,t]\}$ that results in termination at time $t < T$ with positive savings S_t. Let us show that by depositing savings immediately and reporting cash flows truthfully thereafter, the agent gets value greater than S_t. If the agent does so, then $M_s - \bar{M}_s$ evolves as

$$d(\bar{M}_s - M_s) = r(\bar{M}_s - M_s)ds + (dX_s - \mu'ds),$$

---

[22] Recall the assumption that the agent does not consume until liquidation, which we maintain throughout.

and so

$$\bar{M}_{\hat{\tau}} - M_{\hat{\tau}} = e^{r(\hat{\tau}-t)}S_t + \int_t^{\hat{\tau}} e^{r(s-t)}(dX_s - \mu' ds) \Rightarrow E[e^{-r(\hat{\tau}-t)}(\bar{M}_{\hat{\tau}} - M_{\hat{\tau}})] > S_t$$

for any stopping time $\hat{\tau}$ after t, e.g. the earlier of the liquidation time or time T^*. Therefore, if with positive probability the agent ends up in default with positive savings before time T^*, he is following a suboptimal strategy. QED

Proof of Proposition 1 (continued). Let us show that an alternative screening contract with $\beta_t < \sigma(W_t, R_t)$ on a set of positive measure does not provide the agent with incentives to tell the truth. Denote by

$$V(W,R,S) = U(R+S,t^*), \text{ where } t^* \text{ is given by } W = U(R,t^*),$$

the agent's value in an increasing credit line contract when $W_t=W$, $R_t=R$ and the agent has savings $S_t=S$. Consider an arbitrary fabrication strategy $\{\gamma_t\}$, defined by

$$dS_t = (rS_t - \gamma_t)dt \quad \text{and} \quad d\hat{X}_t = \gamma_t dt + dX_t.$$

By Lemma 4, in an increasing credit line contract $e^{-rt}V(W_t,R_t,S_t)$ is a martingale for any strategy γ_t, as long as the agent exhausts savings before liquidation. Using Ito's lemma and setting the drift of $e^{-rt}V(W_t,R_t,S_t)$ to 0, we get

$$(rW_t + \sigma(W_t,R_t)\gamma_t)V_1 + (rR_t + (\mu - \mu') + \gamma_t)V_2 + (rS_t - \gamma_t)V_3 - rV +$$
$$\tfrac{1}{2}(\sigma(W_t,R_t)^2 V_{11} + 2\sigma(W_t,R_t)V_{12} + V_{22}) = 0 \Rightarrow \sigma(W_t,R_t)V_1 + V_2 - V_3 = 0.$$

Consider an alternative contract, where $\beta_t < \sigma(W_t, R_t)$ on a set of positive measure. If the agent spends his savings at rate γ_t, then the drift of $e^{-rt}V(W_t,R_t,S_t)$ is

$$e^{-rt}\left(\gamma_t(\beta_t V_1 + V_2 - V_3) + rW_t V_1 + (rR_t + (\mu - \mu'))V_2 + rS_t V_3 - rV + \tfrac{1}{2}(\beta_t^2 V_{11} + 2\beta_t V_{12} + V_{22})\right)$$

Let us show that the agent can always make $e^{-rt}V(W_t,R_t,S_t)$ a submartingale by an appropriate choice of γ. Note that $V_1(W,R,S) = U_2(R+S,t^*)/U_2(R,t^*)\big|_{W=U(R+S,t^*)} > 0.$

When $\beta_t < \sigma(W_t, R_t)$ (and so $\beta_t V_1 + V_2 - V_3 < 0$) the agent can make the drift of

$e^{-rt}V(W_t, R_t, S_t)$ positive by making γ_t sufficiently negative. Similarly, unless $S_t = 0$, when

$\beta_t > \sigma(W_t, R_t)$ (and so $\beta_t V_1 + V_2 - V_3 < 0$) the agent can make the drift of $e^{-rt}V(W_t, R_t, S_t)$

positive by making γ_t sufficiently large and positive. Finally, (1) if $\beta_t = \sigma(W_t, R_t)$ then

$e^{-rt}V(W_t, R_t, S_t)$ is driftless for all γ_t, and (2) when $S_t = 0$ and $\beta_t > \sigma(W_t, R_t)$ then

$e^{-rt}V(W_t, R_t, 0) = e^{-rt}W_t$ is driftless when the agent sets $\gamma_t = 0$.

If the agent sets γ_t according to the guidelines above until time $\hat{\tau}$ when W_t becomes equal

to R_t for the first time (this must happen before liquidation, as we argued at the beginning

of Section 5), then he would earn a value of at least

$$E[e^{-r\hat{\tau}}(W_{\hat{\tau}} + S_{\hat{\tau}})] = E[e^{-r\hat{\tau}}V(W_{\hat{\tau}}, \underbrace{R_{\hat{\tau}}}_{W_{\hat{\tau}}}, S_{\hat{\tau}})] > V(W_0, R_0, \underbrace{S_0}_{0}) = W_0,$$

where $W_{\hat{\tau}} + S_{\hat{\tau}} = V(W_{\hat{\tau}}, W_{\hat{\tau}}, S_{\hat{\tau}})$ is the agent's value in a fixed credit line when the remaining

credit is $W_{\hat{\tau}}$ and the agent's savings are $S_{\hat{\tau}}$. We conclude that $\beta_t \geq \sigma(W_t, R_t)$ is a necessary

condition for any screening contract to be truth-telling. QED

Lemma 6. For all $W \in [R, U(R,s)]$, $\sigma(W,R)f_{11}(W,R,s) + f_{12}(W,R,s) \leq 0$ and $f_{11}(W,R,s) \leq 0$.

Proof. Letting t^* be the time such that $W = U(R, t^*)$, we have

$$\underbrace{\sigma(W,R)}_{U_1(R,t^*)} f_{11}(W,R,s) + f_{12}(W,R,s) = \frac{df_1(U(R,t^*), R, s)}{dR}.$$

We need to show that $f_1(U(R, t^*), R, s)$ is decreasing in R. Note that the principal's profit

in an increasing credit line contract for values $(U(R, t^*), R)$ is

$$f(U(R, t^*), R, s) = E\left[L + \int_0^{t^*} e^{-ru} 1_{\tau \geq u}(\mu - rL)du + e^{-rt^*} 1_{\tau \geq t^*}(b(R_{t^*}, s - t^*) - L) \mid R_0 = R \right]$$

and $U(R, t^*) = E\left[e^{-rt^*} 1_{\tau \geq t^*} R_{t^*} \right]$.

We have $f_1(U(R,t^*),R,s) = \dfrac{\partial f(U(R,t^*),R,s)/\partial t^*}{\partial U(R,t^*)/\partial t^*}$, where

$$\frac{\partial f(U(R,t^*),R,s)}{\partial t^*} = E[e^{-rt^*}1_{\tau \geq t^*}(\mu - rL - r(b(R_{t^*},s-t^*)-L) + (rR_{t^*} + \mu - \mu')b_1(R_{t^*},s-t^*)$$

$$-b_2(R_{t^*},s-t^*) + \tfrac{1}{2}b_{11}(R_{t^*},s-t^*) \mid R_0 = R] = E\left[e^{-rt^*}1_{\tau \geq t^*}(\mu - \mu')b_1(R_{t^*},s-t^*) \mid R_0 = R\right],$$

from $dR_t = (rR_t + \mu - \mu')\,dt + dZ_t$, Ito's lemma and equation $\mu - rb + rWb_1 - b_2 + \tfrac{1}{2}b_{11} = 0$. We

also have $\partial U(R,t^*)/\partial t^* = E\left[e^{-rt^*}1_{\tau \geq t^*}(\mu - \mu') \mid R_0 = R\right]$. It follows that

$$f_1(U(R,t^*),R,s) = E\left[b_1(R_{t^*},s-t^*) \mid R_0 = R, \tau \geq t^*\right].$$

Since b_1 is decreasing in the first argument by Lemma 2 and, as shown in Lemma 7, the

conditional distributions of R_{t^*} given $R_0 = R, \tau \geq t^*$ are ordered by first order stochastic

dominance in R, it follows that $f_1(U(R,T^*),R,T)$ is decreasing in R, as required.

Finally, $f_{11}(W,R,s) = \dfrac{\partial f_1(U(R,t^*),R,s)/\partial t^*}{\partial U(R,t^*)/\partial t^*}$, and since $\partial U(R,t^*)/\partial t^* > 0$, we need to

show that $\partial f_1(U(R,t^*),R,s)/\partial t^* < 0$. We have

$$\frac{\partial f_1(U(R,t^*),R,s)}{\partial t^*} = \frac{E\left[b_1(R_{t^*},s-t^*)1_{\tau \geq t^*} \mid R_0 = R\right]}{\Pr(\tau \geq t^*)} =$$

$$E[\underbrace{((rR+(\mu-\mu'))b_{11}(R_{t^*},s-t^*) - b_{12}(R_{t^*},s-t^*) + \tfrac{1}{2}b_{111}(R_{t^*},s-t^*))}_{(\mu-\mu')b_{11}(R_{t^*},s-t^*)<0}1_{\tau \geq t^*} \mid R_0 = R]/\Pr(\tau \geq t^*)$$

$$+\frac{b_1(0,s-t^*)}{\Pr(\tau \geq t^*)}\frac{d\Pr(\tau \geq t^*)}{dt^*} - \frac{E\left[b_1(R_{t^*},s-t^*)\mid R_0 = R\right]}{\Pr(\tau \geq t^*)^2}\frac{d\Pr(\tau \geq t^*)}{dt^*} < 0$$

since

$$b_1(0,s-t^*) > E\left[b_1(R_{t^*},s-t^*)\mid R_0 = R, \tau \geq t^*\right] = \frac{E\left[b_1(R_{t^*},s-t^*)\mid R_0 = R\right]}{\Pr(\tau \geq t^*)}$$

by Lemma 2 and $d\Pr(\tau \geq t^*)/dt^* < 0$. QED

Lemma 7. If $0 < R' < R$, the conditional distribution of R_{T^*} given $R_0 = R, \tau \geq T^*$ first

order stochastically dominates the conditional distribution of R_{T^*} given $R_0 = R', \tau \geq T^*$.

Proof. Let $\bar{R}$ be an arbitrary positive number. Define a process $\tilde{R}$ by

$$d\tilde{R}_t = (r\tilde{R}_t + \text{sign}(\tilde{R}_t)(\mu - \mu'))dt + dZ_t.$$

Then $P_t = E_t[1_{\tilde{R}_{T^*}>0} - 1_{\tilde{R}_{T^*}<0}]$ and $Q_t = E_t[1_{\tilde{R}_{T^*}>\bar{R}} - 1_{\tilde{R}_{T^*}<-\bar{R}}]$ are martingales. Define functions P and Q by

$$P(\tilde{R}_t, T^* - t) = E_t[P_{T^*} \mid \tilde{R}_t] \text{ and } Q(\tilde{R}_t, T^* - t) = E_t[Q_{T^*} \mid \tilde{R}_t].$$

Then

$$P(R, T^*) = E[P_{\min(\tau, T^*)}] = \Pr[\tau \geq T^* \mid R_0 = R] \text{ and}$$

$$Q(R, T^*) = E[Q_{\min(\tau, T^*)}] = \Pr[R_{T^*} \geq \bar{R}, \tau \geq T^* \mid R_0 = R],$$

since P_τ and Q_τ are 0 in the event that $\tau < T^*$. We need to show that

$$h(R, T^*) = Q(R, T^*) / P(R, T^*) = \Pr[R_{T^*} \geq \bar{R} \mid \tau \geq T^*, R_0 = R]$$

is monotonically increasing in R for $R{\geq}0$. Since $P(\tilde{R}_t, T^* - t)$ and $Q(\tilde{R}_t, T^* - t)$ are martingales, P and Q satisfy the same partial differential equations

$$(rR + \text{sign}(R)(\mu - \mu'))P_1 - P_2 + \frac{1}{2}P_{11} = 0 \text{ and } (rR + \text{sign}(R)(\mu - \mu'))Q_1 - Q_2 + \frac{1}{2}Q_{11} = 0$$

but with different boundary conditions at $t = T^*$: $P(R, 0) = 1_{R>0} - 1_{R<0}$ and $Q(R, 0) = 1_{R>\bar{R}} - 1_{R<-\bar{R}}$. Plugging $Q(R, s) = h(R, s)P(R, s)$ into the equation for Q, we get

$$(rR + \text{sign}(R)(\mu - \mu'))(h_1 P + hP_1) - h_2 P - hP_2 + \frac{1}{2}(h_{11}P + 2h_1 P_1 + hP_{11}) = 0 \Rightarrow$$

$$h\underbrace{((rR + \text{sign}(R)(\mu - \mu'))P_1 - P_2 + \frac{1}{2}P_{11})}_{0} + \left(rR + \text{sign}(R)(\mu - \mu') + \frac{P_1}{P}\right)h_1 - h_2 + \frac{1}{2}h_{11} = 0.$$

Also, h satisfies the boundary conditions $h(R, 0) = 1_{R \notin (-\bar{R}, \bar{R})}$. It follows that $h(\hat{R}_t, T^* - t)$ is a martingale when $\hat{R}$ follows

$$d\hat{R}_t = \left(r\hat{R}_t + \operatorname{sign}(\hat{R}_t)(\mu - \mu') + P_1(\hat{R}_t, T^* - t) / P(\hat{R}_t, T^* - t) \right) dt + dZ_t.$$

Now, let $\hat{R}$ and $\hat{R}'$ be two such processes that start from a initial conditions $\hat{R}_0 = R$ and $\hat{R}_0' = R'$. Then $\hat{R}_t \geq \hat{R}_t'$ pathwise for all $t \in [0, T^*]$. Let $\hat{\tau}$ the earliest time when $\hat{R}_t' = -\hat{R}_t$ for the first time. Then

$$h(R, T^*) - h(R', T^*) = E[1_{\hat{\tau} < T^*} \underbrace{(h(\hat{R}_{\hat{\tau}}, T^* - \hat{\tau}) - h(\hat{R}_{\hat{\tau}}', T^* - \hat{\tau}))}_{\text{=0, by symmetry}} + 1_{\hat{\tau} \geq T^*} \underbrace{(h(\hat{R}_{T^*}, 0) - h(\hat{R}_{T^*}', 0))}_{>0, \text{ since } \hat{R}_{T^*} > |\hat{R}_{T^*}'|}] > 0.$$

QED

References

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

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

Allen, F. (1985) "Repeated Principal-Agent Relationships With Lending and Borrowing," *Economic Letters,* 17, 27-31

Akerlof, G. (1970) "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, Vol. 84, No. 3. pp. 488-500

Atkeson, A. (1991) "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica*, Vol. 59, No. 4, pp. 1069-1089.

Altman, E. (1968) "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy," *Journal of Finance*, 23, 589-609.

Battaglini, M. (2005) "Long-Term Contracting with Markovian Consumers," *American Economic Review*, Vol. 95, n.3, June 2005, pp. 637-658

Biais, B., T. Mariotti, G. Plantin and J.C. Rochet (2004) "Optimal Design and Dynamic Pricing of Securities," working paper.

Bolton, Patrick and David Scharfstein (1990) "A Theory of Predation Based on Agency Problems in Financial Contracting," *American Economic Review* 80, 94-106.

Bulow, J. and K. Rogoff (1989) "Sovereign Debt: Is to Forgive to Forget?" *American Economic Review*, 79.

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

Cole, H. L. and N. R. Kocherlakota (2001) "Efficient Allocations with Hidden Income and Hidden Storage" *The Review of Economic Studies*, Vol. 68, 3, pp. 523-542

DeMarzo, P. and Fishman, M. (2003) "Optimal Long-Term Financial Contracting with Privately Observed Cash Flows," working paper, Stanford University.

DeMarzo, P. and Fishman, M. (2003b) "Agency and Optimal Investment Dynamics," working paper, Stanford University.

DeMarzo, P. and Y. Sannikov (2005) "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," forthcoming in *Journal of Finance*

Diamond, Douglas W. (1989) "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, Vol. 97, No. 4, pp. 828-862.

Fernandes, A. and C. Phelan (2000) "A Recursive Formulation for Repeated Agency with History Dependence" *Journal of Economic Theory* 91, pp. 223-247

Fudenberg, D., B. Holmstrom and P. Milgrom (1990) "Short-Term Contracts and Long-Term Agency Relationships." *Journal of Economic Theory* Vol. 51(1), 1-31.

Green, E. (1987) "Lending and the Smoothing of Uninsurable Income," in E. Prescott and N. Wallace, Eds., *Contractual Arrangements for Intertemporal Trade*, Minneapolis: University of Minnesota Press.

Grossman, S. and O. Hart (1986) "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *The Journal of Political Economy* 94, 691-719.

Hart, O. and Moore, J. (1994) "A Theory of Debt Based on the Inalienability of Human Capital" *Quarterly Journal of Economics*, 109, 841-79.

Hart, O., and J. Moore (1998) "Default and Renegotiation: a Dynamic Model of Debt," *Quarterly Journal of Economics*, 113, 1, 1-41.

Holmstrom, B. (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies* Vol. 66, No. 1, Special Issue: Contracts, pp. 169-182.

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

Karatzas, I. and Shreve, S. (1991): *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York.

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

Jensen, M. C. (1986) "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers," *American Economic Review* 76, 323-39

Ljungqvist, L. and Sargent, T. J. (2000): *Recursive Macroeconomic Theory*, MIT Press.

Mirrlees, J.A. (1974) "Notes on welfare economics, information and uncertainty." In M. Balch, D. McFadden and S. Wu (eds.), *Essays in Equilibrium Behavior Under Uncertainty,* Amsterdam: North-Holland.

Modigliani, F. and M. Miller (1958) "The Cost of Capital, Corporation Finance and the Theory of Investment," *American Economic Review*.

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

Rothschild, M. and J. Stiglitz (1976) "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics,* 90, 629-650

Sannikov, Y. (2003) "A Continuous-Time Version of the Principal-Agent Problem," working paper, Stanford University.

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

Sufi, A. (2007), "Bank Lines of Credit in Corporate Finance: An Empirical Analysis," *Review of Financial Studies*,

Tchistyi, A. (2005), "Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing," Ph.D. dissertation, Stanford University.

Werning, I. (2001) "Repeated Moral-Hazard with Unmonitored Wealth: A Recursive First-Order Approach," working paper, MIT

Williams, N. (2004) "On Dynamic Principal-Agent Problems in Continuous Time," working paper, Princeton University.