# Myopic Agency and Bonus Bank Contracts

John Y. Zhu[1]

September 23, 2013

**Abstract**

I consider a dynamic principal-agent setting in which the agent repeatedly chooses between hidden "long-term" and "short-term" actions. Relative to the long-term action, the short-term action boosts output today but hurts output tomorrow. The optimal contract inducing long-term actions is explicitly characterized. The key component of the contract is a cliff-like arrangement that can be implemented using bonus banks. Bonus banks feature prominently in the compensation contracts recently introduced at UBS, Morgan Stanley, and other major financial institutions specifically to combat short-termism.

JEL Codes: C73, D86, J33, J41, M52

[1]zhuyiran@wharton.upenn.edu, The Wharton School, University of Pennsylvania. I thank Yuliy Sannikov for suggesting a number of improvements in a private correspondence about the paper. I thank Daniel Neuhann for excellent research assistance, and Richard Kihlstrom for providing valuable guidance throughout the development of the paper. I am grateful to Robert Anderson, Sugato Battacharya, Alex Edmans, Mike Fishman, Itay Goldstein, Ben Hermalin, Steven Matthews, Jianjun Miao, Andy Postlewaite, Andy Skrzypacz, Rob Stambaugh and Bilge Yilmaz for comments and advice. Lastly I would like to thank the seminar participants at the University of Pennsylvania, 2012 NBER/NSF/CEME Conference on Mathematical Economics and General Equilibrium Theory, Northwestern Finance, 2013 NSF/CEME Decentralization Conference, and 2013 China International Conference in Finance for helpful feedback.

# 1 Introduction

How can a firm's owners prevent a manager from taking hidden actions that look good today but hurt long-run profitability? The large literature on moral hazard has surprisingly little to say about this problem. I call this agency problem faced by the owners "myopic" agency. In this paper, I investigate myopic agency in a dynamic principal-agent setting. At each date, the agent takes a hidden action that has persistent effects on firm performance. There are two actions: long-term and short-term, with the agent suffering an effort cost if he chooses the long-term action. The long-term action maintains a certain benchmark level of expected output. The short-term action causes current expected output to rise above the benchmark and future expected output to drop below. The drop is assumed to be sufficiently large relative to the rise so that the principal prefers the long-term action. I then explicitly characterize and study the optimal incentive contract that always induces the long-term action from the agent.

The optimal contract's key feature is a cliff-like arrangement that can be implemented using a bonus bank system: Each day, the agent receives a bonus for producing high output. The bonus pays out a portion today with the rest deferred and paid out over future dates with interest. As the agent continues to produce high output, he receives more and more bonuses and consequently his total bonus pay each date increases: Not only does he receive a portion of today's bonus, but also portions of the bonuses he received the previous dates. However, if at any date the agent fails to produce high output, then not only does he not receive a bonus today, but also all of the remaining unpaid portions of the previous bonuses get wiped out.

Why does the bonus bank arrangement consistently induce the long-term action? When the agent's history of sustained high output production is short - either because he just started the job or because he recently produced low output - most of the large rewards are backloaded: The agent has not yet received a lot of bonuses that can contribute to his pay each date. In order to accrue a healthy set of bonuses and attain those large backloaded rewards, the agent needs to build up a streak of high outputs. This requires taking the long-term action since the short-term action only helps produce high output now, not high output over and over again. Thus, at least initially, the agent is motivated to take the long-term action. As the agent continues to produce high output and receive new bonuses, eventually, the rate at which old bonuses vest and new bonus payments come in balance out. This signals that the previously backloaded large rewards are now paid upfront. At this point, interestingly, the agent is *still* motivated to try and maintain his high output streak and therefore take the long-term action. Why? Because now if the agent takes the short-term action his hard fought high output streak will likely be broken, in which case all of his not-yet-fully-vested bonuses get wiped out, and he will have to start all over again.

The bonus bank concept was pioneered by the management consulting firm Stern Stewart & Co and first adopted by Coca-Cola in 1988 and Briggs and Stratton in 1989 (Boeri, Lucifora, and Murphy, 2013). According to Stern Stewart & Co, the system is designed to help stretch managerial horizons from the short-term to the long-term. In the aftermath of the financial crisis, a number of large banks have transformed their bonus pay systems into versions of the bonus bank arrangement. Also, in 2012 the FDIC approved a proposed bonus-bank-type implementation of Section 956 of Dodd-Frank regulating incentive pay. In each of these cases, the intention behind introducing the bonus bank system was to combat managerial myopia.

In what settings will a principal face a repeated myopic agency problem like in the model? A natural setting is a firm's R&D department. Channeling resources to R&D can help increase the long-term profitability of the firm provided management exerts effort to determine the right projects to support. The short-term action of not fostering R&D will both save the management effort costs and help boost profits and dividends today, but may cause the firm to become obsolete in the future. More generally, settings where the manager must make investment decisions are vulnerable to myopic agency if effective investment requires that the manager exerts hidden effort to decide how best to invest the funds.

Financial markets are also rife with myopic agency problems. Subprime lending is a perfectly legitimate long-term action, but doing it prudently requires effort to carefully vet the borrowers and assess the complex associated risks. Without the proper incentives, an agent may engage in indiscriminate lending under terms overly favorable to the borrower. While deviating to this short-term action is an easy way to inflate business today, the long-term effects can be disastrous. Hedge funds are also susceptible to myopic agency. Here, the desired long-term action involves the manager's exerting effort to turn his innate skill into generating alpha. The undesirable short-term action can be employing strategies that are essentially equivalent to writing a bunch of puts, which inflates net asset value today but exposes the fund to significant future tail risk.

A common property in these examples is that the moral hazard problem is played out over time. It cannot be properly modeled when actions in date $t$ affect the outcome only in date $t$. This action persistence makes the design of optimal contracts tricky.

To get a feel for the potential pitfalls of contracting under repeated myopic agency, consider the problem of trying to induce a manager to take the long-term action today. The usual view of moral hazard, familiar in an insurance setting, tells us to reward good outcomes and punish bad outcomes. But rewarding high output today will only encourage the manager to take the short-term action. A more sensible strategy is to wait until tomorrow (when the long-term effects of today's action have been realized) and reward the agent only if high output is produced then. While this strategy works in a one-shot model of myopic agency, in a dynamic setting, this arrangement only serves to pass today's agency problem onto tomorrow. Facing such a contract with delayed rewards and punishments, a sophisticated manager will simply behave today and wait until tomorrow to take the short-term action.

When actions are non-persistent, rewarding high output, as a general rule, helps alleviate the agency problem. However, in a myopic agency setting, rewarding high output can exacerbate the agency problem just as much as alleviate it. Divorced from the past and the future, a stochastic output today communicates very little information about the agent's decisions. As a result, the principal must pay careful attention to the pattern of production across time. The structure of the optimal dynamic contract reflects this requirement.

The first salient feature of the optimal contract is how it decides when the agent is doing a good job. A typical optimal dynamic contract will formulate an endogenous measure(s) of good performance based on the history of outcomes. It will then use this measure to determine the spot contract that the agent faces today. In the myopic agency model, high output is a priori an ambiguous signal of performance. High output today is a positive signal only when tomorrow's output is high. Otherwise, it really looks like the agent took the short-term action today. Similarly, tomorrow's high output is a positive signal only if the day after tomorrow's output is high. Thus, the surest indicator that the agent is behaving is if there is an *unbroken string of high outputs.* This means the optimal contract should care about sustainability. Indeed, the endogenous good performance measure of the optimal contract tracks the number of consecutive high output dates leading up to today. The optimal contract shows that there is a myopic agency-driven contribution to the economic value of consistent performance.

On the flip side, because the optimal contract values sustaining high output, low output after a string of high outputs can have a seemingly disproportionate effect on the agent's compensation. This creates a cliff-like arrangement where, the longer the high output streak, the higher the agent ascends on the contractual cliff but also the farther he falls when the streak is broken. The cliff-arrangement is isomorphic to the bonus bank arrangement described above and is a natural way of dealing with the double-edged sword of rewarding high output in a myopic agency setting.

## 1.1 A Comparative Discussion

The optimal dynamic contract under myopic agency differs in significant ways from many of those that arise under traditional, non-persistent moral hazard. Unlike the consistency-based measure of good performance used by the myopic agency optimal contract, traditional dynamic contracts typically have a measure that resembles counting total high output dates net low output dates. The aggregation is insensitive to the specific timing of performance. This type of approximately order-independent aggregation of performance over time has been the dominant measure of performance quality in the dynamic contracting literature since the subject's inception. Early examples include the cumulated performance $S$ of Radner (1981, 1985) and the linear aggregator of Holmstrom and Milgrom (1987).

More generally, myopic agency forces us to change how we think about what makes a contract good or bad. Traditional contract theory's basic tenet is: if you want to get the agent to do the right thing, you have to provide him with enough incentives. This value reflects the traditional IC-constraint, which is a lower bound on incentives. The implication being, a bad contract is one whose incentives are too small. The myopic agency IC-constraint is, instead, concerned with balancing incentives across time: Lemma 1 shows that a contract is incentive-compatible only if tomorrow's incentive level is larger than an increasing function of today's incentive level. In the myopic agency world, when a contract is not incentive-compatible today, it is true one can say tomorrow's incentives are too small relative to today's, but it is also equally valid to say that today's incentives are too large relative to tomorrow's. Thus, a bad contract isn't necessarily one that has small incentives, but rather one that has temporally unbalanced incentives.

Contracts in the myopic agency setting are naturally recursive over two state variables: the incentive level $\Delta_t$ and the agent's promised value $W_t$. The stochastic parameter $\Delta_t$ is the sensitivity of the agent's date $t$ ex-post promised value to date $t$ output. I then show that the optimal value function can, in fact, be expressed only as a function of $\Delta_t$. This is because in the optimal contract, $W_t$ is implicitly determined by $\Delta_t$. The value function is the unique solution to a Bellman equation which I solve (see Equation (3)).

I also deduce the Markov law of motion for $\Delta_t$. I show that the incentive level of the optimal contract rises as the agent continues to produce high output. Thus, the optimal contract features incentive escalation. This feature distinguishes the myopic agency optimal contract from optimal contracts in traditional moral hazard models. When low output is generated, the incentive level drops all the way down to zero.

In the paper, I formally compare the myopic agency optimal contract and the contract the principal would use if he decided to abstract away from the persistent nature of the actions. I call the latter contract the traditional contract. It is the optimal contract of the appropriate non-persistent version of the model. I find that the traditional contract isn't terrible in the sense that if it is used in the actual myopic agency setting in lieu of the optimal contract, the agent will still take the long-term action. However, it always over-incentivizes the agent relative to the optimal contract and is therefore inefficient.

Also, the traditional contract has a stationary incentive level instead of having escalating incentives like in the myopic agency optimal contract. Another type of incentive escalation is documented in Gibbons and Murphy (1992), which provides an alternative justification that rests on career concerns and requires risk aversion. The incentive escalation result of this paper is compatible with - but not dependent on - risk aversion. Another difference is that the incentive escalation in my optimal contract exhibits history-dependent dynamics, escalating only after high output.

While this paper explores how to induce the long-term action through contracts, there is a related literature that focuses on why managers oftentimes take a variety of short-term actions in equilibrium. See, for example, Stein (1988, 1989). Another related literature deals with innovation. The process generates a dynamic not unlike the one produced by the long-term action. Manso (2012) embeds such a two-date innovation problem within a principal-agent framework. Edmans, Gabaix, Sadzik and Sannikov (2012) considers a model of dynamic manipulation that allows the agent to trade off, on a state-by-state basis, future and present performance. Their optimal contract can be implemented using a "dynamic incentive account" - a type of deferred reward system that does not wipe out old rewards if the agent does not perform today. Varas (2013) considers a model of project creation where the principal faces a compensation problem similar to the one in this paper. By rewarding the agent for the timely completion of a good project whose quality is hard to verify, the principal might inadvertently induce the agent to cheat and quickly produce a bad project.

This paper is most closely related to Holmstrom and Milgrom (1991). Recall, they observe that if the agent has two tasks $A$ and $B$, the incentives of $A$ may exert a negative externality on that of $B$. In my model, one can think of the task of managing the firm today as task $A$ and managing the firm tomorrow as task $B$. And just as in Holmstrom and Milgrom (1991), if incentives today are too strong relative to those of tomorrow, the agent will take the short-term action, which favors the firm today and neglects the firm tomorrow. Now, Holmstrom and Milgrom use this to explain why contracts often have much lower-powered incentives than what the standard single-task theory might predict. In my paper things are further complicated by the dynamic nature of the model. Specifically, today's task $B$ will become tomorrow's task $A$. Each date's task is both task $A$ and task $B$ depending on the

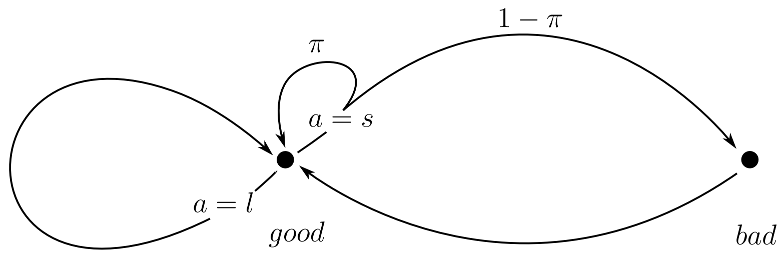


Figure 1: Transition function between the good and bad states.

frame of reference. Therefore, the conclusion in my model is not that incentives should be low-powered, but that incentives start low and optimally escalate over time.

My paper is also part of a small literature on persistent moral hazard. An early treatment by Fernandes and Phelan (2000) provides a recursive approach to computing optimal contracts in repeated moral hazard models with effort persistence. Jarque (2010) considers a class of repeated persistent moral hazard problems that admit a particularly nice recursive formulation: those with actions that have exponential lagged effects. She shows that under a change of variables, models in this class translate into traditional non-persistent repeated moral hazard models. Her work can be interpreted as a justification for the widely used modeling choice of ignoring effort persistence in dynamic agency models.

My paper considers, in some sense, the opposite type of persistence to that of Jarque (2010). Here, the ranking of the actions is flipped over time. Today: short-term > long-term; future: long-term > short-term. With this type of persistence, results become noticeably different from those of the non-persistent class.

Sannikov (2013) considers a Brownian model of persistent moral hazard. He focuses on two tractable cases: a large firm case in which noise goes to infinity and cost of effort goes to zero at comparable rates, and the exponential lagged effects cases considered in discrete time by Jarque (2010). The continuous-time analog of $\Delta$ also appears as a state variable.

The rest of the paper is organized as follows: Section 2 introduces the basic repeated myopic agency model. I recursively characterize and solve for the optimal contract. Section 3 interprets the optimal contract. The novel performance measure and the cliff arrangement emerge. Comparisons with bonus bank contracts and non-persistent optimal contracts are made.

# 2 Repeated Myopic Agency

A principal contracts an agent to manage a firm at dates $t = 0, 1, 2 \ldots$ At each date $t$, the firm can be in one of two states: $\sigma_t = good$ or $bad$. If $\sigma_t = good$, then the agent can apply one of two hidden actions: a long-term action $a_t = l$ or a short-term action $a_t = s$. If the agent applies the long-term action, then the firm remains in the good state: $\sigma_{t+1} = good$. If the agent applies the short-term action, then $\sigma_{t+1} = good$ with probability $\pi < 1$, and $\sigma_{t+1} = bad$ with probability $1 - \pi$. If $\sigma_t = bad$, then there is no action choice and the state reverts back to *good* at the next date. See Figure 1.

Actions and states are hidden from the principal, who can only observe output. At each date $t$, the firm produces either high output $X_t = 1$ or low output $X_t = 0$. If $\sigma_t = good$ and $a_t = l$ then the probability that the firm produces high output at date $t$ is $p < 1$. If $\sigma_t = good$ and $a_t = s$ then the firm produces high output for sure at date $t$. If $\sigma_t = bad$ then the firm produces low output for sure at date $t$. I assume that $\sigma_0 = good$.

Notice, if the agent always takes the long-term action, then the firm is always in the good state and there is always a probability $p$ of high output. A deviation today to the short-term action boosts expected output today by $1-p$ and lowers expected output tomorrow by $Q := (1-\pi)p$. I assume that $1-p < \beta Q$ where $\beta \in (0,1)$ is the intertemporal discount factor. This assumption says that the gain today from taking the short-term action is outweighed by the present discounted loss tomorrow.

*Definition of a Contract*

At each date $t$, the principal may make a monetary transfer $w_t \geq 0$ to the agent. Note, each $w_t$ can depend on the history of outputs up through date $t$. However, $w_t$ cannot depend on the unobservable action nor the state. At each date $t$, the principal may also recommend an action $a_t$ to be taken provided $\sigma_t = good$. A contract is a complete transfer and action plan $w = \{w_t\}, a = \{a_t\}$. The principal's utility is $\mathbf{E}_a\big[\sum_{t=0}^{\infty} \beta^t (X_t - w_t)\big]$ and the agent's utility is $\mathbf{E}_a\big[\sum_{t=0}^{\infty} \beta^t (w_t + c1_{a_t = s})\big]$. I assume selecting the short-term action saves the agent effort cost $c > 0$. I also assume that the agent has an outside option worth $K$.

**Assumption** (A)**.** *The principal always wants to induce the agent to take the long-term action.*

The action sequence taken by the agent should, in principle, be determined as part of the optimal contracting problem. However, I show that requiring the agent to take the long-term action is without loss of generality under certain parameterizations of the model. See Corollary to Theorem 1. Moreover, the solution to the sustained long-term action case will serve as an important benchmark for future analyses of the unconstrained optimal contracting problem.

**Assumption** (B)**.** *The agent can freely dispose of output before the principal observes the net output.*

This assumption means that if the output $X_t = 1$, the agent has the ability to secretly throw it away and make the principal think $X_t = 0$. The agent will dispose of output whenever the contract he is facing promises more total expected utility after low output. By the revelation principle, it suffices to add a no output disposal incentive-constraint to the contracting problem and assume the agent does not dispose of output. The optimal contracting problem is the following constrained maximization:

$$\max_{\{w_t \geq 0\}_{t=0}^{\infty}} \mathbf{E}_{\{a_t = l\}_{t=0}^{\infty}} \left[ \sum_{t=0}^{\infty} \beta^t (X_t - w_t) \right]$$

$$
\begin{aligned}
\text{s.t.} \quad & \{a_t = l\}_{t=0}^{\infty} \in \arg\max_{\{a_t\}} \mathbf{E}_{\{a_t\}} \left[ \sum_{t=0}^{\infty} \beta^t (w_t + c1_{a_t=s}) \right] \\
\text{s.t.} \quad & \mathbf{E}_{\{a_t=l\}_{t=0}^{\infty}} \left[ \sum_{t=0}^{\infty} \beta^t w_t \right] \geq K \\
\text{s.t.} \quad & \text{The agent does not dispose output.}
\end{aligned}
$$

*Costly Long-Term Actions and the Short-Termism Problem*

The assumption that the long-term action is more costly to the agent than the short-term action is important. The assumption is rooted in the fact that many real-life long-term actions require more effort to implement than their short-term alternatives. For example, consider the act of fostering a positive workplace culture. Such an action requires financial investment now, say, in attracting or retaining popular role-model workers, sponsoring company-wide social events, or establishing a comfortable working environment. On the other hand, the financial benefits from having a positive workplace culture mostly accrue later relative to the investment. Thus, fostering a positive workplace culture is a long-term action.

But implementing this long-term action doesn't just involve a simple financial investment. Successfully fostering a positive workplace culture also requires a significant effort commitment from management. To attract or retain popular role-model workers, management first needs to exert effort to identify them - through a costly external search or costly internal private evaluations. To ensure that company-wide social events serve their intended purpose, management may need to attend the events and be enthusiastic about it. To successfully establish a comfortable working environment, management first needs to figure out what it entails. Thus, fostering a positive workplace culture is a long-term action that requires significant effort costs. Of course, some components of this costly long-term action are contractible (e.g. attending company sponsored events). However, there are less tangible but no less important components of fostering a positive workplace culture that are harder to contract: figuring out what needs to be done to have a comfortable working environment or identifying special workers or being enthusiastic at company events. The dynamic myopic agency model based on hidden costly long-term actions has something to say about inducing these components.

In the introduction, I also provide examples of hidden costly long-term actions in financial markets: A bank can work hard to determine how best to structure loans made out to risky clients instead of lending indiscriminately. A fund manager can work hard to figure out the right trading strategy to generate sustained returns instead of engaging in cheap tricks that artificially pump up asset value today.

The presence of an effort cost associated to the long-term action significantly affects the contracting problem. Suppose instead, the long-term action is no more costly than the short-term action (i.e. $c = 0$). Then the optimal contracting problem is trivial: By giving the agent a flat wage $(1 - \beta)K$, the agent's participation constraint is satisfied and he is indifferent between all actions. The principal can then recommend that the agent choose the long-term action.

Once $c > 0$, any contract that hopes to induce the long-term action must provide incentives by making the agent's welfare sensitive to firm performance. Given that the agent can freely dispose of output, sensitivity to firm performance means that the agent must be better off following high output.[2] This then creates the fundamental tension in the model:

**The Short-Termism Problem.** *To induce the long-term action which is costly, the agent must be given a contract that makes him weakly better off after high output. But many contracts of this type encourage short-termism instead because the action that leads to the highest chance of generating high output today is the short-term action.*

I now solve for the optimal contract and show how it overcomes the short-termism problem.

*Solving the Optimal Contracting Problem*

Since the principal and agent share a common discount factor, the timing of pay is irrelevant. In this paper, I focus on the version of the optimal contract that pays the agent the earliest. This ensures that the optimal contract is approximately robust to small perturbations to the agent's discount factor that make him more impatient than the principal.

Let $\mathcal{H}_t$ denote the set of all binary sequences of length $t+1$. $\mathcal{H}_t$ is the set of histories of firm outputs up through date $t$. Define $\mathcal{H}_{-1} := \emptyset$. In general, the agent's promised value depends on the history of outputs up through yesterday as well as today's state. So for each $h_{t-1} \in \mathcal{H}_{t-1}$ and state $\sigma_t$, define $W_t(h_{t-1}, \sigma_t)$ to be the agent's date $t$ promised value given all the relevant information:

$$W_t(h_{t-1}, \sigma_t) = \mathbf{E}_a\left[\sum_{i=t}^{\infty} \beta^{i-t}(w_i + c1_{a_i=s}) \;\middle|\; h_{t-1}, \sigma_t\right]$$

In general, the promised value is unknown to the principal since states are hidden. However, the paper restricts attention to only those contracts where the agent takes the long-term action all the time and the state is always good. Therefore, on the equilibrium path, it is well-defined to speak of $W_t(h_{t-1})$, which only depends on the publicly observable $h_{t-1}$ and is known to the principal. It is also useful to define the agent's date $t$ *ex-post promised value*:

$$W_t^+(h_t) = w_t(h_t) + \beta W_{t+1}(h_t)$$

In the traditional approach to solving dynamic contracting problems, the agent's promised value is the key state variable and the entire recursive formulation of optimality is built around it. This will not be the case in the myopic agency model. For each history $h_{t-1}$, define $\Delta_t(h_{t-1}) := W_t^+(h_{t-1}1) - W_t^+(h_{t-1}0)$. As a function over $\mathcal{H}_{t-1}$, $\Delta_t$ is a random variable representing the date $t$ incentive level of the contract. It turns out that $\Delta_t$ is the natural state variable of the myopic agency model. Thus, the optimal contract will be recursive over $\Delta_t$, not $W_t$. The centrality of $\Delta_t$ is due to the incentive-compatibility condition:

[2] It is not necessary to assume free disposal of cash. For example, suppose instead, the agent has a second action which is a traditional effort action: Higher traditional effort today leads to higher likelihood of high output today with no persistent effects. Then if the principal wants to induce high traditional effort in addition to the long-term action, then the agent will have to be better off following high output.

**Lemma 1.** *A contract always induces the long-term action if and only if at each date $t$ and after each history $h_{t-1} \in \mathcal{H}_{t-1}$*

$$\Delta_{t+1}(h_{t-1}1) \geq \varepsilon(\Delta_t(h_{t-1})) := \frac{(1-p)\Delta_t(h_{t-1}) + c}{\beta Q} \tag{1}$$

*Proof.* See Appendix. □

Even though incentive-compatibility involves comparing on- and off-equilibrium continuation values, notice that Equation (1) only involves on-equilibrium continuation values. In the Appendix, I explain how this is possible by first showing that the agent's off-equilibrium continuation value can be expressed as a function of on-equilibrium continuation values following a one-shot deviation. Equation (1) is precisely the condition that prevents such one-shot deviations. I then show that checking for one-shot deviations is sufficient.

The IC-constraint is a lower bound for the incentives tomorrow as a function of the incentives today. Or equivalently, it is an upper bound on incentives today as a function of incentives tomorrow. Either way, the absolute levels of incentives do not matter so much, what matter are the relative levels of incentives over time. The greater the incentives are today, the more tempting it is to take the short-term action today. Therefore, the incentives tomorrow must also keep pace to ensure that the agent properly internalizes the future downside of taking the short-term action today. Similarly, if the effort cost $c$ is large, then it is again tempting to take the short-term action today. So the lower bound on tomorrow's incentives is also an increasing function of $c$.

Notice the IC-constraint cares about tomorrow's incentives following high output today but ignores tomorrow's incentives following low output today. This is because if the agent actually deviates and takes the short-term action, low output today never occurs. Since the incentive level after low output is immaterial under a short-term action deviation, it is not included in the IC-constraint.[3]

Recall, if a contract's promised value is decreasing in output, the agent will simply engage in free disposal after high output and mimic low output. Thus, the no free disposal constraint of the optimal contracting problem amounts to a monotonicity of promised value requirement and is equivalent to $\Delta_t \geq 0$ for all $t$.

**Definition.** *The term "$\Delta$-contract" will mean a contract whose initial incentive level is $\Delta$. The term "$(\Delta, W)$-contract" will mean a $\Delta$-contract that delivers promised value $W$ to the agent. For each $\Delta$, define $W(\Delta)$ to the agent's promised value under the optimal $\Delta$-contract.*

[3] What would happen if, more generally, taking the short-term action today only increased the probability of high output today to $p + q \leq 1$? The IC-constraint would take the more general form: $(p+q)\Delta_{t+1}(h_{t-1}1) + (1-p-q)\Delta_{t+1}(h_{t-1}0) \geq \frac{q\Delta_t(h_{t-1})+c}{\beta Q}$. Clearly, the generalized IC-constraint does care about the incentives tomorrow following low output today. However, the weighting of tomorrow's incentives is under the *counterfactual* measure generated by taking the short-term action today. Relative to the actual measure, the counterfactual measure underweights low output today. As a result, the generalized IC-constraint still underemphasizes the incentives tomorrow following low output today. Numerical simulations show that the optimal contract under the generalized IC-constraint is qualitatively similar to the optimal contract under Equation (1).

Fix any contract. The sum of the principal's and agent's utilities is $\frac{p-c}{1-\beta}$, which is a constant of the model. Since any two $(\Delta, W)$-contracts give the agent the same utility (by definition), they must also give the principal the same utility. Moreover, since any contract by assumption always induces the long-term action, any two $(\Delta, W)$-contracts also have the same cost to the principal. Cost is calculated as the expected present discounted value of all payments to the agent. Call this cost quantity $C(\Delta, W)$.

Fix a $(\Delta, W)$-contract which has cost $C(\Delta, W)$. Pick a date 1 continuation contract. This continuation contract is some $\Delta'$-contract delivering some promised value $W'$ to the agent, which has cost $C(\Delta', W')$. So $C$ admits a recursive formulation. Moreover, since all contracts generate the same surplus, $C(\Delta, W) = C(\Delta, W(\Delta)) + (W - W(\Delta))$.

$C(\Delta, W(\Delta))$ is the cost to the principal of the optimal $\Delta$-contract. As an abuse of notation, let $C(\Delta) := C(\Delta, W(\Delta))$. To solve the optimal contracting problem, it suffices to characterize $C(\Delta)$. It turns out that $C(\Delta)$ also admits a recursive formulation. This means without loss of generality, an optimal $\Delta$-contract's continuation contracts are optimal $\Delta'$-contracts for some incentive levels $\Delta'$.

I now sketch an informal derivation of the Bellman equation for $C(\Delta)$ culminating in Theorem 1. A formal proof of the theorem is supplied in the appendix.

Fix an arbitrary $\Delta$ and consider the optimal $\Delta$-contract. Since the date 0 incentive level of this contract is $\Delta$, the IC-constraint implies that date 1's incentive level following high output at date 0 must be at least $\varepsilon(\Delta)$. Therefore, viewed as a contract in its own right, this date 1 continuation contract is some $(\Delta^u, W^u)$-contract where $\Delta^u \geq \varepsilon(\Delta)$ and $W^u \geq W(\Delta^u)$. Similarly, the date 1 continuation contract following low output at date 0 is some $(\Delta^d, W^d)$-contract where $\Delta^d \geq 0$ and $W^d \geq W(\Delta^d)$.

Looking at the function $C(\Delta, W)$, it is clearly increasing in the second argument. Intuitively, it should also be (weakly) increasing in the first argument since providing more incentives to the agent shouldn't be costless. If we take this to be true, then consider the following transformation of our optimal $\Delta$-contract: Replace the $(\Delta^u, W^u)$-continuation contract with the optimal $\varepsilon(\Delta)$-contract and the $(\Delta^d, W^d)$-continuation contract with the optimal 0-contract. Add $W^u - W(\varepsilon(\Delta))$ to $w_0(1)$ and add $W^d - W(0)$ to $w_0(0)$.

The resulting contract respects limited liability, is still incentive-compatible, it's promised value to the agent is the same, and the initial incentive level is still $\Delta$. So it is still an optimal $\Delta$-contract. Thus I have shown $C(\Delta)$ admits a recursive characterization.

Let $w(1)$ and $w(0)$ denote the high and low output payments to the agent at date 0 in the optimal $\Delta$-contract. So far, I have shown that $C(\Delta) = p(w(1) + \beta C(\varepsilon(\Delta))) + (1-p)(w(0) + \beta C(0))$. Once $w(1)$ and $w(0)$ are pinned down, the recursive characterization is complete.

Given the above expression for $C(\Delta)$, the date 0 incentive level is $(w(1) + \beta C(\varepsilon(\Delta))) - (w(0) + \beta C(0))$, which must be $\Delta$ by assumption. A little algebra shows that the minimal $w(1)$ and $w(0)$ satisfying this condition while also respecting the agent's participation constraint are:

$$\begin{aligned} w(1) &= \Big(\Delta - \beta\big[C(\varepsilon(\Delta)) - C(0)\big]\Big)^+ \vee \Big(K - \big[\beta C(\varepsilon(\Delta)) - (1-p)\Delta\big]\Big) \\ w(0) &= \Big(\beta\big[C(\varepsilon(\Delta)) - C(0)\big] - \Delta\Big)^+ \vee \Big(K - \big[\beta C(0) + p\Delta\big]\Big) \end{aligned} \tag{2}$$

I can now express $C(\Delta)$ exclusively as a function of $\Delta$, $C(0)$ and $C(\varepsilon(\Delta))$. This is the Bellman equation characterizing $C(\Delta)$. Solving the equation formally solves the optimal contracting problem.

**Theorem 1.** *The Markov law of the state variable $\Delta$ for optimal $\Delta$-contracts is*

$$\Delta \to \begin{cases} \varepsilon(\Delta) & \textit{following high output} \\ 0 & \textit{following low output} \end{cases}$$

*The optimal cost function $C(\Delta)$ satisfies the following explicitly solvable Bellman equation:*

$$C(\Delta) = \max\{\beta C(\varepsilon(\Delta)) - (1-p)\Delta, \beta C(0) + p\Delta, K\} \tag{3}$$

*The solution is a piecewise linear, weakly increasing, convex function.*

*Proof.* See Appendix. □

**Corollary.** *Fixing all other parameters, for all sufficiently small effort cost $c$, always inducing the long-term action is optimal.*

In the proof of Theorem 1, I show that when $K = 0$, the cost of an optimal $\Delta$-contract is proportional to the size of the effort cost. So as a function, $C(\Delta)$ is homogenous of degree 1 in $c$. After any history, the cost of always inducing the long-term action from now on is bounded above by $C(\varepsilon^\infty(0)) = C(c/(\beta Q - (1-p)))$. Thus, as $c$ tends to zero, so will the upper bound. In particular, the upper bound will be smaller than the surplus generated from taking the long-term action a single time. Finally for the $K > 0$ case, I simply note: If for a set of parameters it is optimal to always induce the long-term action, then clearly it will still be optimal to always induce the long-term action if all the parameters are the same but $K$ increases.

# 3 A Representative Optimal Contract

In this section I explicitly write down a representative solution to the Bellman equation and analyze the implied optimal contract.

The mechanics of how the optimal contract provides incentives is best revealed when the agent's participation constraint is not binding. This way, cash payments do not do double duty and are there purely to provide incentives. Also, the payment schedules are simpler because after any history, only the high output or the low output payment is positive, never both. For simplicity, assume $K = 0$, so that the agent's participation constraint is guaranteed not to bind irrespective of the other parameter values. Also assume, for concreteness, $Q = p$. When $Q = p$, the short-term action has maximal negative consequences tomorrow. As a result, the agency problem is the least severe and the optimal cost function $C(\Delta)$ is at its lowest and simplest. When $K = 0$ and $Q = p$, $C(\Delta)$ is a two-piece piece-wise linear function:

$$C(\Delta) = \begin{cases} C(0) & \text{if } \Delta \in [0, x_1] \\ \beta C(0) + p\Delta & \text{if } \Delta \in [x_1, \infty) \end{cases} \tag{4}$$

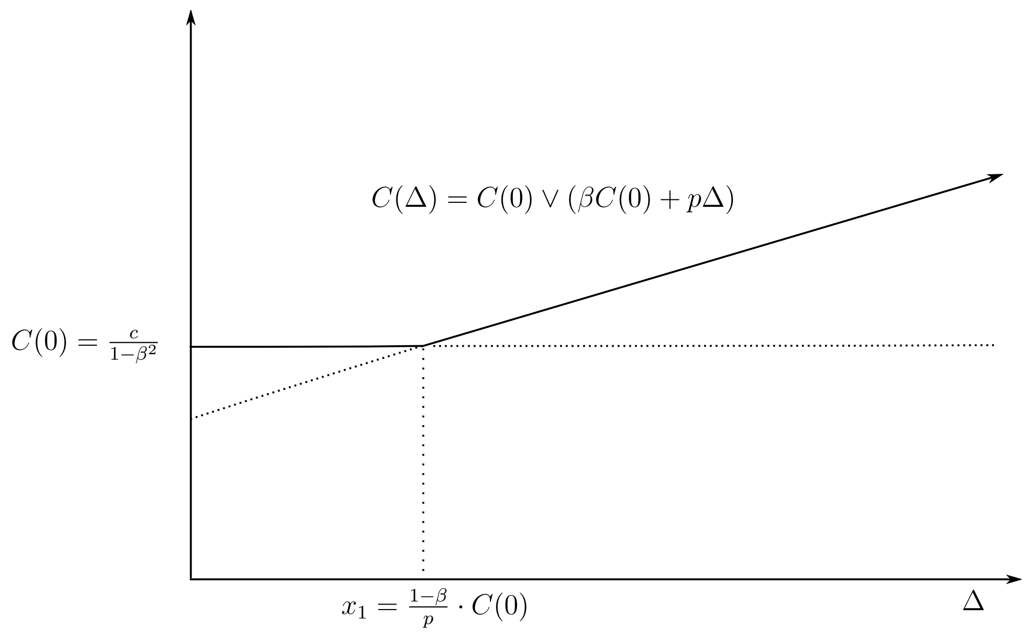

Figure 2: The optimal cost function $C(\Delta)$ when $Q = p$ and $K = 0$.

where $C(0) = \frac{c}{1-\beta^2}$ and $x_1 = \frac{1-\beta}{p} \cdot C(0)$.

The optimal contract is not unique in the $(K = 0, Q = p)$ case because the first piece of $C(\Delta)$ is flat. Thus, the optimal contract is the optimal $\Delta$-contract for any $\Delta \in [0, x_1]$. However, when the agent's participation constraint is not binding, any value of $Q$ other than $p$ results in a strictly increasing $C(\Delta)$, and so the unique robust optimal contract is the optimal 0-contract. In the analysis that follows, I focus on the optimal 0-contract.

The optimal contract delivers expected value $C(0)$ to the agent. As a function of the incentive level $\Delta$, the high output payment is $w(1) = p\,[\Delta - x_1]^+$ and the low output payment is $p\,[x_1 - \Delta]^+$.

## 3.1 How and Why the Optimal Contract Works

In Theorem 1, the Markov law for the optimal contract's incentive level implies that $\Delta$ only takes values in the discrete set $\{\varepsilon^N(0)\}_{N=0,1,\ldots}$. Today's incentive level is $\varepsilon^N(0)$ if and only if the agent has produced high output for $N$ consecutive dates leading up to today. Thus, even though formally the state variable of the optimal contract is $\Delta$, in practice it suffices to count the number $N$ of consecutive high output dates leading up to today. Thus, the optimal contract cares about *sustaining* high output and not simply generating a large number of high output dates.

Figure 3a graphically characterizes $N$ and $\varepsilon^N(0)$. The vertical axis is the incentive level $\Delta$. The tick marks represent today's incentive level as a function of the performance indicator $N$. As $N$ increases, today's incentive level $\varepsilon^N(0)$ increases in a concave way, converging to $\overline{\Delta} := \varepsilon^\infty(0) = \frac{c}{\beta p-(1-p)}$. Each day, if the firm produces high output, then the incentive level increases by a tick mark; if the firm produces low output, then the incentive level drops down to 0.

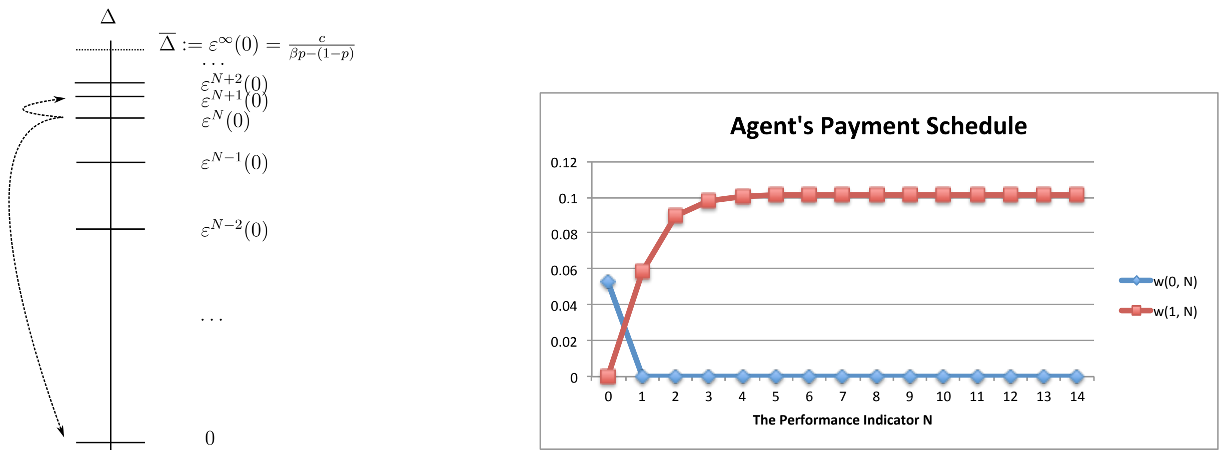


(a) (b)

Figure 3: (a) The law of motion for $N$ and for the optimal contract's incentive level $\varepsilon^N(0)$. (b) Data generated under the parameterization ($Q = p = .8,\ \beta = .9,\ c = .1,\ K = 0$). The vertical axis is the portion of today's output that is paid to the agent.

In addition to the incentive level, the performance indicator $N$ also completely determines today's continuation and spot contracts. The continuation contract is the optimal $\varepsilon^N(0)$-contract. As a function of $N$, today's spot contract is

$$w(1, N) = \begin{cases} 0 & \text{if } N = 0 \\ p \cdot \left\{ \frac{c}{\beta p} \frac{1-\left(\frac{1-p}{\beta p}\right)^N}{1-\frac{1-p}{\beta p}} - \frac{c}{(1+\beta)p} \right\} & \text{if } N > 0 \end{cases} \tag{5}$$

$$w(0, N) = \begin{cases} p \cdot \frac{c}{(1+\beta)p} & \text{if } N = 0 \\ 0 & \text{if } N > 0 \end{cases} \tag{6}$$

Figure 3b plots the high and low output payment schedules $w(1, N)$ and $w(0, N)$ as a function of $N$. Notice, the agent is initially paid for low output. This payment is to ensure that the contract starts with the minimal incentive level, 0.

The main driver of incentives is the contract's "reward" payment schedule for high output. As shown in Figure 3b, $w(1, N)$ is zero for low values of $N$ and then becomes a positive, concave, convergent function. Thus, when the agent's performance is poor, he is not rewarded for high output. Then there is an intermediate stage when the agent is facing small but rapidly increasing rewards. Eventually, rewards for high output level off at a large sum for sufficiently good performance.[4]

The performance indicator and reward schedule arrangement $(N, w(1, N))$ contains the basic intuition for why the optimal contract works: When the principal sets out to write a contract for the agent, he faces a dilemma. On the one hand, he needs to reward the agent for good performance since he is asking the agent to perform a productive yet costly hidden

[4]In this particular example, low values of $N$ means $N \in \{0\}$. In general, the initial period of no high output reward may last longer than one date.

task. On the other hand, the principal realizes that a large reward for high output, paradoxically, tempts the agent to be short-termist since the short-term action guarantees the agent the reward. The optimal contract's performance indicator-reward schedule arrangement is tailored to resolve this dilemma.

Re-imagine the indicator-reward arrangement as a path up a cliff. It helps to think of Figure 3a as the frontal view of the cliff and the reward schedule $w(1, N)$ in Figure 3b as the profile. At the bottom of the cliff, the path is barren. At the top, the path is lined with big carrots. The path represents the domain of the performance indicator and the carrots comprise the reward schedule. Under the parameterization of Figure 3b, the big carrots amount to about 10% of the firm's output.

Initially ($N = 0$), the agent is at the bottom of the cliff. The performance indicator determines his progress along the path. In particular, every high output moves the agent upward one step. His goal is to consume the big carrots at the top. To reach them, the agent must take the long-term action. Scaling the cliff requires sustained progress. Taking the short-term action today may help the agent progress today but it will not help him progress over and over again.

Once the agent has reached the top of the cliff, he can begin to consume the big carrots. The agent's hidden actions represent two approaches to consumption. He can take sensible bites by continuing to take the long-term action or he can gorge by taking the short-term action.

Will the agent gorge? Suppose he does. Then one consequence of this choice is that the agent is deprived of a big carrot tomorrow. But this is a relatively minor loss compared to the other one looming on the horizon. Recall the agent's progress is governed by a performance indicator that cares only about sustained high output. Therefore, when the likely low output event is realized tomorrow, the streak is broken and the indicator drops down to zero . . .

And the agent falls off the cliff.

## 3.2 The Cliff Arrangement and Bonus Banks

As the cliff analogy demonstrates, the optimal contract motivates the agent in two mutually reinforcing ways. Initially, the backloaded nature of the reward schedule induces the agent to the take the long-term action. This, in and of itself, is unremarkable. Many optimal dynamic contracts have some form of backloading of rents.

The novelty comes when the backloaded reward schedule is mapped against the optimal contract's novel performance indicator $N$. Because $N$ takes precipitous drops, this combination creates a "contractual cliff" that provides the second way to motivate the agent - when performance has already reached a high level and the previously backloaded payments have come to the fore. At this point, the fear of falling off the cliff and starting all over again serves as an effective deterrent to short-termism. This cliff arrangement differs from typical optimal arrangements seen in traditional non-persistent dynamic moral hazard settings. In those settings, optimal contracts typically care about something that roughly approximates high output dates net low output dates. This inattention to the timing of high output effectively means there's no cliff.

How serious is this omission in a myopic agency setting?

**Remark.** *Suppose the principal eliminates the optimal contract's cliff entirely. At each date, the agent is simply given the limiting spot contract as the performance indicator goes to infinity. The limiting spot contract's pay-to-performance sensitivity is larger than that of any of the spot contracts used in the optimal contract. Despite having more "skin in the game," the agent now never chooses to take the long-term action.*

To see this, set $N = \infty$ in Equation (5). A little algebra shows that the limiting spot contract rewards high output with

$$\overline{w} := \overline{\Delta}\frac{1}{1+\beta} \tag{7}$$

Under the proposed arrangement, the incentive level at all times is $\overline{w}$. However, since $\overline{w} < \overline{\Delta}$, the IC-constraint requires tomorrow's incentive level after high output today to be at least $\varepsilon(\overline{w}) > \overline{w}$. Thus, the IC-constraint is always violated.

The cliff arrangement of the myopic agency optimal contract is structurally identical to a type of *bonus bank* arrangement. A bonus bank is essentially a partially deferred bonus that gradually pays out over time and contains a clawback provision.

A number of financial institutions have recently adopted versions of the bonus bank arrangement in order to combat short-termism. In 2008 the banking giant UBS enacted a sweeping overhaul of its pay system. Central to this overhaul was a new bonus structure for 6500 of UBS's highest earners. The new bonuses are issued as "callable" amortizing bonds that mature anywhere between 3 and 5 years. Importantly, these bonds held by the employees can be wiped out (i.e. called back at zero call price) if, in the interim, certain benchmarks are not reached. So for example, by not reaching the predetermined benchmark this year, a member of the UBS executive board will stand to lose up to 4 current and future payments of a bonus received last year.

Both Morgan Stanley and Credit Suisse have also introduced policies that defer bonuses over three years. Just as in the UBS example, these bonuses are subject to partial or full elimination if losses occur.

The cliff-arrangement of the optimal contract admits a natural implementation using bonus banks. For concreteness, I will implement the arrangement as a callable amortizing bond like what is done at UBS:

- When $N = 0$, the agent does not receive a bonus for generating high output.
- Otherwise, whenever the agent generates high output he receives a callable amortizing bond with principal $P_0$ defined to be the following constant:

$$P_0 = \sum_{t=0}^{\infty} \beta^t \left[w(1, t+1) - w(1, t)\right]$$

- The certificate pays an initial amount $w(1, 1) - w(1, 0) = w(1, 1)$.
- The rest of $P_0$ then compounds at the model interest rate $1/\beta - 1$.
- In subsequent dates, the amortization schedule dictates payments $w(1, 2) - w(1, 1)$, $w(1, 3) - w(1, 2)$, ... as long as the agent continues to generate high output.

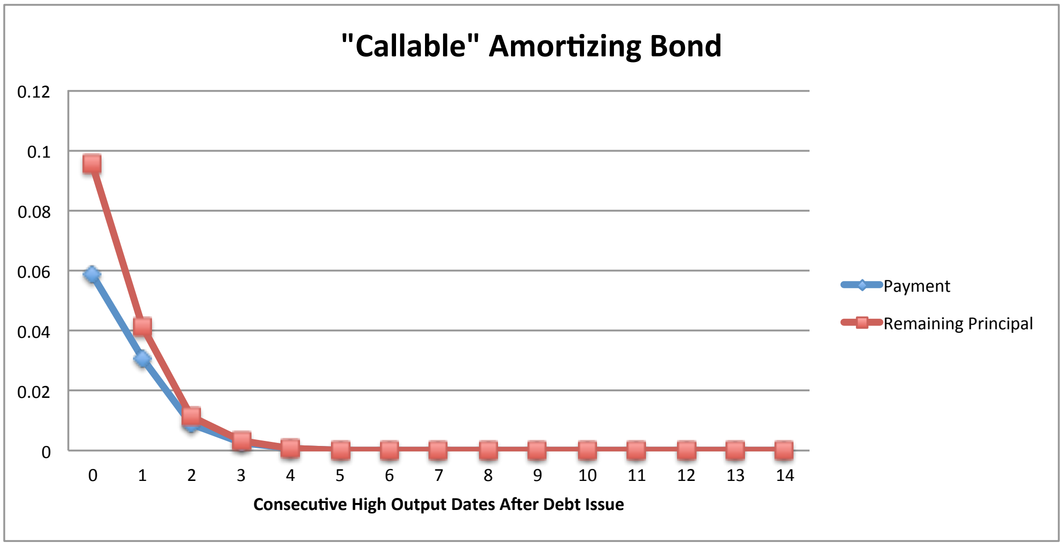


Figure 4: Evolution of the payments and remaining principal of a debt certificate issued to fulfill a bonus obligation.

- After $t$ consecutive dates of high output following the issue date, the remaining principal is

$$P_t = \sum_{N=t}^{\infty} \beta^{N-t} \left[ w(1, N+1) - w(1, N) \right]$$

- If at any point low output is produced, the debt is wiped out.

Unlike the UBS certificate, the optimal contract's bond never matures. However, Figure 4 shows that the payments $w(1, t+1) - w(1, t)$ rapidly decrease over time and therefore the bond can be approximated by one that matures in finite time.

# 4 Incentive Escalation

The agency problems that the myopic agency setting is geared toward analyzing have traditionally been dealt with in settings where actions have no persistent effects on firm performance. In this section I show that there are significant differences between the optimal contracts that arise in these two settings. The key conceptual difference lies with the incentive structure. The myopic agency optimal contract's incentive structure features smaller incentive levels and exhibits *incentive escalation*. The optimal contract in the non-persistent setting exhibits a larger, fixed incentive level.

To highlight these differences, I suppose the true model is the myopic agency model but the principal decides to abstract away from the persistent nature of the moral hazard problem. He replaces the true model with the appropriate, simplified, non-persistent approximation. I solve for the optimal contract in this non-persistent setting and compare it against the myopic agency optimal contract.

In the non-persistent version of the model, the long-term action is relabeled as effort. Effort today is assumed to generate a probability $p$ of producing high output today. The short-term action is relabeled shirking. Shirking today is assumed to generate a probability $1-\beta Q$ of producing high output today. Notice, the principal replaces the short-term action's true multi-period effect on output with its present value. Effort cost is still $c$.

Recall in the true model, I assume that the parameters are such that $\beta Q > 1-p$. That is, the loss tomorrow associated with taking the short-term action today outweighs the present gain. In the non-persistent model, this is precisely the condition that ensures shirking produces high output with lower probability compared to effort. Therefore, Assumption (A) can still be sensibly applied and the principal can solve for the optimal contract (subject to always inducing effort). I will call this contract the *traditional contract* to distinguish it from the optimal contract of the true model.

The traditional contract is stationary. The spot contract used at every date is: pay the agent $\overline{\Delta}$ if output is high and nothing if output is low. Thus the traditional contract simply repeats the optimal contract of the one shot version of the non-persistent model. The pay-to-performance sensitivity of the contract is trivially always $\overline{\Delta}$. In addition, the incentive level is also always $\overline{\Delta}$. In contrast, the optimal contract's history dependent incentive level is $\varepsilon^N(0)$ which is always below $\overline{\Delta}$, satisfying $0=\varepsilon^0(0)<\varepsilon(0)<\varepsilon^2(0)<\ldots<\varepsilon^\infty(0)=\overline{\Delta}$.

Lemma 1 implies that the traditional contract is incentive-compatible. That is, even though the principal sweeps the persistence of the moral hazard problem under the rug and uses the traditional contract, he still induces the long-term action from the agent. Obviously, the traditional contract is inefficient since it is not isomorphic to the optimal contract. The degree of inefficiency can be usefully quantified:

**Remark.** *Relative to the optimal contract, the traditional contract always over-incentivizes the agent.*

The incentives of the optimal contract are always smaller since $\varepsilon^N(0)<\varepsilon^\infty(0)$ for all $N$. Also, the incentive level of the optimal contract is non-stationary, always increasing after high output and dropping to zero after low output.

Not surprisingly, the pay-to-performance sensitivity of the traditional contract is also too high relative to the optimal contract. However, there is a meaningful difference between comparing incentive levels and comparing pay-to-performance sensitivities. Notice as the agent's performance increases, the optimal contract's incentive level converges to that of the traditional contract. This convergence does not happen for pay-to-performance sensitivity. Specifically, the limiting spot contract of the optimal contract is not the spot contract of the traditional contract. Recall, Equation (5) implies that limiting reward for high output in the optimal contract is $\overline{w}:=\overline{\Delta}/(1+\beta)$ which is strictly smaller than $\overline{\Delta}$.

The comparison of pay-to-performance sensitivities reveals a significant difference in pay level between the optimal and traditional contracts. As the agent's performance becomes arbitrarily good, the reward he receives for producing high output in the optimal contract is still strictly smaller than the reward he would receive each time he produced high output in the traditional contract. This has important feasibility implications.

For example, suppose that the agent's cost $c$ exceeds the threshold $\beta Q-(1-p)$. Then $\overline{\Delta}>1$ and the traditional contract pays the agent more than he produces. This need not be the case if the principal uses the optimal contract. Since $\overline{w} < \overline{\Delta}$, there are values of $c$ where $\overline{\Delta} > 1$ but $\overline{w} < 1$. Moreover, even if $\overline{w} > 1$, the backloaded nature of the reward schedule implies that only a fraction of the time is the optimal contract's spot contract rewarding the agent more than he produces.

The incentive escalation feature of the optimal contract predicts that agents with histories of sustained high output should receive the highest incentives and pay-to-performance sensitivities. This escalation result complements and provides an alternative theory of wage dynamics to Gibbons and Murphy (1992). In that paper, the authors observe that the optimal contract should emphasize total reward-to-performance sensitivity, which should factor in implicit career concerns as well as explicit pay. Escalation in explicit pay-to-performance sensitivity is then driven by the gradual disappearance of career concerns over time. Since career concerns disappear regardless of performance history, the escalation of pay-to-performance sensitivity is history independent. This contrasts with the myopic agency optimal contract that concentrates all of the escalation behind high output histories. Another difference is with the agent's risk attitude. While the escalation result in my model is fully compatible with a risk averse agent, it is not dependent on risk aversion.

On a related note, introducing risk aversion to the agent's utility function can produce a wealth effect that escalates a contract's pay-to-performance sensitivity without affecting incentives. Such a dynamic should not be confused with the incentive escalation result of my paper. Incentives and pay-to-performance sensitivity are used somewhat interchangeably in the literature but they are theoretically distinct objects. Recall, incentives measure promised-value-to-performance sensitivity, which directly drives agent behavior through the IC-constraint. Fluctuations in pay-to-performance sensitivity may reflect an underlying change in incentives but may also be a symptom of other changes such as in wealth effects or career concerns.

# 5 Conclusion

Short-termism is a major component of many managerial agency problems. This paper investigates optimal contracting when a manager can take hidden short-term actions that hurt the future health of the firm. Like in many real-life settings, the short-term action in this model boosts performance today. This temporarily masks the inferiority of the short-term action and creates a tricky contracting setting where simply rewarding high output is no longer guaranteed to eliminate the agency problem. In this setting, I derive the optimal contract that always induces the long-term action and show that it exhibits a cliff-arrangement that can be implemented using bonus banks. The contract draws clear parallels with the new bonus-pay arrangements introduced at a number large financial institutions to combat short-termism in the wake of the 2008 crisis. The optimal contract differs in significant ways from traditional dynamic optimal contracts. The myopic agency optimal contract values sustained high output instead of aggregate or average output. Its incentive level is non-stationary, escalating over time and after high output while remaining strictly below the incentive level of the traditional contract. More generally, the paper establishes a framework that can be used to model a variety of agency problems where the assumption that actions have no persistent effects is flawed.

# 6 Appendix

*Proof of Lemma 1.* Fix a contract that calls for the agent to always take the long-term action. Consider the diagram below representing today's pay and tomorrow's "ex-post" promised values following a history $h$ leading up through yesterday. Promised values are calculated with respect to the measure generated by always taking the long-term action.

$$\begin{array}{ll} & W^+(h11) \\ w(h1) & \\ & W^+(h10) \end{array}$$

$$\begin{array}{ll} & W^+(h01) \\ w(h0) & \\ & W^+(h00) \end{array}$$

Suppose the agent decides to commit a one-shot deviation to the short-term action today. Then his payoff is $c+w(h1)+\beta(\pi(pW^+(h11)+(1-p)W^+(h10))+(1-\pi)W^+(h10))$. Letting $W(h1)$ and $W(h0)$ denote the agent's promised values tomorrow following high and low output today (again, calculated under the measure generated by always taking the long-term action), the utility from deviation can be rewritten as $c + w(h1) + \beta W(h1) - \beta((1-\pi)p(W^+(h11) - W^+(h10))) = c + w(h1) + \beta W(h1) - \beta Q\Delta(h1) = c + W^+(h1) - \beta Q\Delta(h1)$. Incentive compatibility requires that

$$pW^+(h1) + (1-p)W^+(h0) \geq c + W^+(h1) - \beta Q\Delta(h1)$$

which, upon rearrangement, is equivalent to

$$\Delta(h1) \geq \frac{(1-p)\Delta(h) + c}{\beta Q}$$

Thus the proposed IC-constraint is a necessary condition ensuring that one-shot deviations from always taking the long-term action are suboptimal. I now show sufficiency by proving that if one-shot deviations are suboptimal then all deviations are suboptimal.

First note, if after some history the agent is better off employing a deviation strategy, then he is better employing a deviation strategy that only involves deviating in a finite number of dates. This is due to discounting. This observation allows me to prove sufficiency using induction.

So fix a contract that always calls for the long-term action and satisfies the proposed IC-constraint. Suppose there are no profitable $T$-length deviations. Now, suppose on the contrary, there exists a history $h$ such that following $h$ there exists a profitable $T+1$-length deviation. If this deviation does not involve deviating right away, then it is in fact, a $T$-length deviation. Contradiction. So suppose the deviation does involve deviating right away. Then the agent's payoff following $h$ is $w(h1)+\beta(\pi U(D(h1))+(1-\pi)(w(h10)+\beta U(D(h10)))$ where $U(D(h1))$, $U(D(h10))$ are the payoffs from employing the continuations $D(h1)$, $D(h10)$ of the deviation strategy after histories $h1$ and $h10$. By induction, this payoff is weakly less than $w(h1) + \beta(\pi(pW^+(h11) + (1-p)W^+(h10)) + (1-\pi)(w(h10) + \beta(pW^+(h101) + (1-p)W^+(h100))) = w(h1) + \beta(\pi(pW^+(h11) + (1-p)W^+(h10)) + (1-\pi)W^+(h10)) \leq pW^+(h1) + (1-p)W^+(h0)$.

The key step in the proof is to realize that the payoffs from employing the continuations of the deviation strategy are the same regardless of the initial deviation at history $h$. This is what allows the inductive step to go through. □

*Proof of Theorem 1.*
*Step 1. $C(\Delta)$ is convex and the optimal contract is randomization-proof.*

Fix $\lambda \in (0,1)$ and $\Delta_1$- and $\Delta_2$-contracts with costs $C_1$ and $C_2$ and payments $w_{\Delta_1}(h)$ and $w_{\Delta_2}(h)$ after each history $h$. Then the contract that pays $\lambda w_{\Delta_1}(h) + (1-\lambda)w_{\Delta_2}(h)$ is a $\lambda\Delta_1 + (1-\lambda)\Delta_2$-contract with cost $\lambda C_1 + (1-\lambda)C_2$.

To show that the optimal contract is randomization-proof, consider a contract with possibly random payments to the agent based on a sequence of random signals $\theta_0, \theta_1 \dots$ where the payments at date $t$ can depend on $\{\theta_0, \dots \theta_t\}$ in addition to $h_t$. Assume the action at date $t$ is chosen before the realization of $\theta_t$. Define $W_t(h_{t-1}, \theta_0, \dots \theta_{t-1})$ as the appropriate generalization of $W_t(h_{t-1})$ where the expectation is also taken over all possible present and future signal realizations $\theta_t, \theta_{t+1} \dots$ in addition to all possible future histories. We can then define $\Delta_t(h_{t-1}, \theta_0, \dots \theta_{t-1})$ as the appropriate generalization of $\Delta_t(h_{t-1})$. A slight generalization of the proof of Lemma 1 shows that the IC-constraint with public randomization is:

$$\mathbf{E}_{\theta_t}\Delta(h_{t-1}1, \theta_0, \dots \theta_t) \geq \frac{(1-p)\Delta(h_{t-1}, \theta_0, \dots \theta_{t-1}) + c}{\beta Q}$$

If the optimal contract were not randomization-proof, then there would exist some contract with randomization only in a finite number of dates that was less costly. Thus, it suffices to show for every $n$, given any contract with randomization only in the first $n$ dates and initial incentive level $\Delta$, there exists a non-random $\Delta$-contract with equal cost. Here, "non-random" only means it is not a function of the public randomization variables.

The proof is by induction. The $n = 1$ case: Consider a contract that only randomizes over $\theta_0$ and has initial incentive level $\Delta$. Let $w_0(h_0, \theta_0)$ be the random date 0 payments of the contract. Define $w_0(h_0) := \mathbf{E}_{\theta_0} w_0(h_0, \theta_0)$. Given a realization $(h_0, \theta_0)$, the date 1 continuation contract is by assumption some non-random $\Delta(h_0, \theta_0)$-contract with some cost $C(\Delta(h_0, \theta_0))$. Define $\Delta(h_0) := \mathbf{E}_{\theta_0}\Delta(h_0, \theta_0)$ and $C(\Delta(h_0)) := \mathbf{E}_{\theta_0} C(\Delta(h_0, \theta_0))$. The proof of the convexity of $V$ shows that there exists a non-random $\Delta(h_0)$-contract with cost $C(\Delta(h_0))$.

Consider the following non-random contract: pay $w_0(h_0)$ at date 0 after history $h_0$, followed by the $\Delta(h_0)$-continuation contract. By construction, it has the same cost as the original contract with randomization and is incentive-compatible for all dates $t \geq 1$. All that is left to show is incentive-compatibility at date 0:

$$\Delta(1) = \mathbf{E}_{\theta_0}\Delta(1,\theta_0) \geq \frac{(1-p)\Delta + c}{\beta Q}$$

So the $n = 1$ case is proved.

Now suppose it is proved for $n = N$. Consider a contract with randomization only in the first $N+1$ dates. For every realization $(h_0, \theta_0)$, consider the date 1 continuation contract with randomization in at most $N$ dates. By assumption, it can be replaced with a non-random $\Delta(h_0, \theta_0)$-contract. And now we are back to the $n = 1$ case.

*Step 2.* $\Delta^* < \varepsilon(0)$ *where* $\Delta^*$ *is the smallest element of* $\arg\min C(\Delta)$.

Suppose not. Then consider the following contract: at date 0, pay the agent $(K - \beta C(\Delta^*))^+$; at date 1 give the agent the optimal $\Delta^*$-contract. Since $D^* \geq \varepsilon(0)$, this contract is incentive-compatible. It is a 0-contract by construction and its cost is $\beta C(\Delta^*) \vee K$. If the cost is $\beta C(\Delta^*)$, then it is strictly less costly than $C(\Delta^*) = \min C(\Delta)$. Contradiction. Otherwise it has cost $K = \min C(\Delta)$ and $\Delta^* = 0$. Contradiction.

*Step 3.* $C(\Delta) = \max\{\beta C(\varepsilon(\Delta)) - (1-p)\Delta,\ \beta C(\Delta^*) + p\Delta\}$.

Consider the following sequence of panels:

$$\begin{matrix} w(1) & C_1 & & w(1) + \beta(C_1 - C(\varepsilon(\Delta))) & C(\varepsilon(\Delta)) & \\ & & \longrightarrow & & & \longrightarrow \\ w(0) & C_2 & & w(0) + \beta(C_2 - C(\Delta^*)) & C(\Delta^*) & \end{matrix}$$

The first panel depicts a generic $\Delta$-contract: it's date 0 high and low output payments as well it's date 1 continuation contracts' costs. The second panel depicts the first step of a two step cost-decreasing transformation of the $\Delta$-contract. The continuation contracts are changed to be the least costly ones subject to the date 0 incentive constraint and the resulting distortion in promised value is balanced out by increasing the date 0 payments by appropriate amounts. That the date 1 continuation contract following high output is the optimal $\varepsilon(\Delta)$-contract follows from Step 1 and Step 2. Notice, the resulting contract in the second panel is also a $\Delta$-contract; it has the same cost as the first panel's $\Delta$-contract; and it has weakly larger date 0 payments.

$$\begin{matrix} \Big(\Delta - \beta\big[C(\varepsilon(\Delta)) - C(\Delta^*)\big]\Big)^+ \vee \Big(K - \big[\beta C(\varepsilon(\Delta)) - (1-p)\Delta\big]\Big) & C(\varepsilon(\Delta)) & \\ & & \longrightarrow \\ \Big(\beta\big[C(\varepsilon(\Delta)) - C(\Delta^*)\big] - \Delta\Big)^+ \vee \Big(K - \big[\beta C(\Delta^*) + p\Delta\big]\Big) & C(\Delta^*) & \end{matrix}$$

The third panel depicts the second step of the two step cost-decreasing transformation. As much of the common portion of the date 0 high and low output payments from the previous panel's $\Delta$-contract that can be factored out has been factored out. Note, the presence of a participation constraint means that in some cases, it is not possible to factor out the entire common portion. This transformation weakly decreases the cost of the contract while maintaining its property of being a $\Delta$-contract. The date 0 high and low output payments are now the smallest possible ones for a $\Delta$-contract. Thus, the contract represented by the third panel is in fact the optimal $\Delta$-contract which, by definition, has cost $C(\Delta)$. This cost can be explicitly computed using the diagram, yielding the following Bellman equation:

$$C(\Delta) = \max\{\beta C(\varepsilon(\Delta)) - (1-p)\Delta,\ \beta C(\Delta^*) + p\Delta,\ K\}$$

Note this Bellman equation differs slightly from Equation (3): 0 is replaced with $\Delta^*$. For the rest of the proof, I will call this Bellman equation the *weak Bellman equation* and Equation (3) the *strong Bellman equation*. The weak Bellman equation becomes the strong Bellman equation once we show that $C(\Delta)$ is weakly increasing.

*Step 4. $C(\Delta)$ is a weakly increasing, piecewise linear function.*

Let $\mathscr{C}$ be the space of all weakly convex functions $f(x)$ defined on $[0, c/(\beta Q - (1-p)]$ satisfying $0 \leq f(x) \leq pc/[(1-\beta)Q] \vee K + px$ and $-Q(1-p)/(Q-(1-p)) \leq f'^{-}(x) \leq f'^{+}(x) \leq p$ for all $x$. Under the $L^\infty$ norm, $\mathscr{C}$ is a compact Banach space.

Define the operator $Tf(x) := \max\{\beta f(\varepsilon(x)) - (1-p)x, \beta \min\{f(x)\} + px, K\}$ based on the weak Bellman equation.

*Step 4a. $T$ is a contraction of $\mathscr{C}$.*

We first need to show that $T$ is an endomorphism of $\mathscr{C}$. Define operators $T_a$ and $T_b$ where $T_a f(x) := \beta f(\varepsilon(x)) - (1-p)x$ and $T_b f(x) := \beta \min\{f(x)\} + px$.

If $f$ is weakly convex, then clearly so are $T_a f$, $T_b f$. Then $Tf = \max\{T_a f, T_b f, K\}$ is weakly convex.

If $f$ is nonnegative, then clearly so is $Tf$.

If $f'^{-}(x) \geq -Q(1-p)/(Q-(1-p))$ for all $x$, then $T_a f'^{-}(x) = \beta f'^{-}(\varepsilon(x)) - (1-p) \geq \beta(-Q(1-p)/(Q-(1-p)))(1-p)/(\beta Q) - (1-p) = -Q(1-p)/(Q-(1-p))$ for all $x$. Also $T_b f'^{-}(x) = p$. Therefore $Tf'^{-}(x) \geq -Q(1-p)/(Q-(1-p))$ for all $x$. A similar argument shows that if $f'^{+}(x) \leq p$ for all $x$ then $Tf'^{+}(x) \leq p$ for all $x$.

Lastly, suppose $f(x) \leq pb/[(1-\beta)Q] \vee K + px$ for all $x$. I show $Tf(x) \leq pb/[(1-\beta)Q] + px$ for all $x$ by breaking the analysis into two cases: when $pb/[(1-\beta)Q] \geq K$ and when $K > pb/[(1-\beta)Q]$.

First, suppose $pb/[(1-\beta)Q] \geq K$. Then $T_a f(x) = \beta f(\varepsilon(x)) - (1-p)x \leq \beta pb/[(1-\beta)Q] + \beta p[(1-p)x + b]/(\beta Q) - (1-p)x = (\beta pb)/Q \cdot [1/(1-\beta) + 1/\beta] + [p(1-p)/Q - (1-p)]x \leq pb/[(1-\beta)Q] + px$.

Now suppose $K > pb/[(1-\beta)Q]$. Then $T_a f(x) = \beta f(\varepsilon(x)) - (1-p)x = \beta K + \beta p[(1-p)x + b]/(\beta Q) - (1-p)x = \beta K + pb/Q + [p(1-p)/Q - (1-p)]x \leq K + px$.

Also, in both cases, $T_b f(x) < pb/[(1-\beta)Q] \vee K + px$ for all $x$. Thus $Tf(x) \leq pb/[(1-\beta)Q] \vee K + px$ for all $x$. I have proved $T(\mathscr{C}) \subset \mathscr{C}$.

Let $f$, $g \in \mathscr{C}$. For any $x$, $|Tf(x) - Tg(x)| \leq \max\{|T_a f(x) - T_a g(x)|, |T_b f(x) - T_b g(x)|\}$. Since both $T_a$ and $T_b$ are contractions on $\mathscr{C}$, $T$ is a contraction on $\mathscr{C}$. Therefore, a unique fixed point of $\mathscr{C}$ under $T$ is assured.

*Step 4b. $V(\Delta)$ is weakly increasing and therefore, satisfies the strong Bellman equation.*

In this sub-step, I will affix the superscript $Q$ to $\mathscr{C}$ and $T$ to emphasize that these objects change as I change the model parameter $Q$.

Consider a realization of the model with $Q = p$, space $\mathscr{C}^p$ and operator $T^p$. It is easy to show that for all $n$, $(T^p)^n K$ is a convex two-piece piecewise linear function $\geq K$ where the first piece is flat and the second piece has slope $p$. In particular, $(T^p)^n K$ is nondecreasing for all $n$ and therefore, so is the unique fixed point of $T^p$.

Now consider any model realization with $Q < p$ with corresponding $\mathscr{C}^Q$ and $T^Q$. Let $f^Q \in \mathscr{C}^Q$ and $f^p \in \mathscr{C}^p$. Then I say $f^Q > f^p$ if and only if on the smaller domain of $\mathscr{C}^p$ functions, $f^Q - f^p$ is nondecreasing, nonnegative and not identically zero.

It is easy to show that if $f^Q$ and $f^p$ are both nondecreasing and $f^Q \geq f^p$, then $T^Q f^Q > T^p f^p$. Thus $(T^Q)^n K$ is strictly increasing for all $n > 1$ and therefore, so is the unique fixed point of $T^Q$.

*Step 4c. $C(\Delta)$ is a piecewise linear function.*

I provide an explicit algorithm for computing $C(\Delta)$ assuming the participation constraint is not binding (e.g. $K = 0$). A similar but more complicated algorithm works when $K > 0$.

Fix a $\Delta \in (0, c/(\beta Q - (1-p)))$. Define $N \geq 1$ to be the unique integer satisfying $\varepsilon^{-N}(\Delta) \leq 0 < \varepsilon^{-N+1}(\Delta)$. Define $m(x) := (1-p)x/Q - (1-p)$. Construct the following piecewise linear function $f$ piece by piece starting from the left:

$$f(x) = \begin{cases} f(0) + m^N(p)x & x \in (0, \varepsilon^{-N+1}(\Delta)] \\ f(\varepsilon^{-N+1}(\Delta)) + m^{N-1}(p)(x - f(\varepsilon^{-N+1}(\Delta))) & x \in (\varepsilon^{-N+1}(\Delta), \varepsilon^{-N+2}(\Delta)] \\ \dots & \\ f(\varepsilon^{-1}(\Delta)) + m(p)(x - f(\varepsilon^{-1}(\Delta))) & x \in (\varepsilon^{-1}(\Delta), \Delta] \\ f(\Delta) + p(x - f(\Delta)) & x > \Delta \end{cases}$$

where $f(0)$ is (uniquely) defined to satisfy $f(\Delta) = p\Delta + \beta \min\{f\}$. Notice, it must be that $f(0) > 0$.

As a function of $\Delta$, $d(\Delta) := f(0) - f(\varepsilon(0))$ is strictly increasing. For $\Delta$ sufficiently close to 0, $d(\Delta) < 0$. For $\Delta$ sufficiently large - it suffices for $\Delta$ to be large enough so that $m^{N-1}(p) \leq 0$, $d(\Delta) > 0$. Let $\Delta^*$ satisfy $d(\Delta^*) = 0$. Then the corresponding function $f^*$, when restricted to $\mathscr{C}$, is the unique fixed point of $T$. So $C \equiv f^*$. $\square$

# References


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

[2] *Executive Remuneration and Employee Performance-Related Pay: A Transatlantic Perspective,* ed. by T. Boeri, C. Lucifora and K. J. Murphy. Oxford University Press, USA 2013

[3] Edmans, A., X. Gabaix, T. Sadzik and Y. Sannikov (2012) "Dynamic CEO Compensation," *The Journal of Finance* Vol. 67, pp. 1603-1647

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

[5] Gibbons, R. and K. J. Murphy (1992) "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy* Vol. 100, pp. 468-505

[6] Green, E. (1987) "Lending and the Smoothing of Uninsurable Income," in *Contractual Arrangements for Intertemporal Trade,* ed. by E.C. Prescott and N. Wallace. Minneapolis: University of Minnesota Press pp. 3 - 25

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

[8] Holmstrom, B. and P. Milgrom (1991) "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, & Organization* Vol. 7, pp. 24-52

[9] Jarque, A. (2010) "Repeated Moral Hazard with Effort Persistence," *Journal of Economic Theory* Vol. 145, pp. 2412-2423

[10] Manso, G. (2012) "Motivating Innovation," forthcoming in *The Journal of Finance*

[11] Radner, R. (1981) "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica* Vol. 49, pp. 1127-1148

[12] Radner, R. (1985) "Repeated Principal-Agent Games with Discounting," *Econometrica* Vol. 53, pp. 1173-1198

[13] Sannikov, Y. (2013) "Moral Hazard and Long-Run Incentives," working paper

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

[15] Stein, J. (1988) "Takeover Threats and Managerial Myopia," *Journal of Political Economy* Vol. 96, pp. 61 - 80

[16] Stein, J. (1989) "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics* Vol. 104, pp. 655 - 69

[17] Varas, F. (2013) "Contracting Timely Delivery with Hard to Verify Quality," working paper