Myopic Agency[1]

John Y. Zhu

January 15, 2013

Abstract

I consider a fully rational repeated moral hazard problem in which the agent chooses between hidden "long-term" and "short-term" actions. Relative to the long-term action, the short-term action boosts expected output today at the expense of expected output tomorrow. In general, today's expected output is affected by both today's and yesterday's actions. This is unlike traditional principal-agent models that ignore the action persistence present in many real-life moral hazard problems. The paper shows that ignoring action persistence can significantly distort optimality. For example, the myopic agency optimal contract emphasizes sustained success rather than a more typical index like net success. Compared to its traditional cousins, this optimal contract also has more subdued incentives that escalate after good performance. In addition, the contract provides a fresh perspective on the optimality of upward sloping pay scales and high watermark contracts. Lastly, a common contractual feature that the previous dynamic contracting literature has struggled with emerges as an intuitive optimal arrangement: severance pay.

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

[1]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, Alex Edmans, Mike Fishman, Itay Goldstein, Ben Hermalin, Andy Skyrzypacz and Bilge Yilmaz for comments and advice. Lastly I would like to thank the seminar participants at University of Pennsylvania, 2012 NBER/NSF/CEME Conference on Mathematical Economics and General Equilibrium Theory, and Northwestern Finance for helpful feedback.

1 Introduction

I investigate a novel moral hazard problem I call myopic agency in a dynamic principal-agent setting. The agent takes hidden actions that have persistent effects on firm performance. There are two actions: long-term and short-term. The long-term action maintains a certain benchmark level of expected output. The short-term action causes current expected output to rise and future expected output to drop. The drop is assumed to be sufficiently large relative to the rise so that the principal prefers the long-term action. This paper explicitly characterizes the optimal contract that always induces the long-term action from the agent.

A natural setting for myopic agency is a firm's R&D department. Channeling a portion of revenues to R&D helps increase the long-term profitability of the firm. The short-term action of not investing in R&D will 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 capital investment decisions are vulnerable to myopic agency if the quality of the decisions cannot be well-monitored or understood. For decades, Kodak employed a razor-blade business model with inexpensive cameras and a very profitable film business. Naturally, most of the investment went into film. In the 90s as the digital revolution was in full swing, Kodak faced a difficult decision between doubling down on film or switching to digital. In Kodak's case, the problem wasn't that the firm underinvested like in the R&D example - the company was still developing patents. The problem was that Kodak chose to stick with film. Throughout the 90s, the company suffered very little from this decision. Even in the late 90s, the firm sold about 1 billion rolls of film each year. However, by 2011, the bottom had fallen out and sales had sunk to 20 million.

Financial markets are also rife with myopic agency situations. Subprime lending, if done correctly, is a perfectly legitimate long-term action. However, given the nature of the counterparty, responsible subprime lending requires effort to carefully vet the borrowers and model and weigh the complex associated risks. Without the proper incentives, an agent may easily deviate to the short-term action of indiscriminate lending under terms overly favorable to the borrower. Hedge funds are also susceptible to myopic agency. Here, the long-term action involves the manager exerting effort to turn his innate skill into generating α. The short-term action can simply be writing a bunch of puts, exposing the fund to significant future tail risk. Or it can involve portfolio pumping through, say, thinly traded stocks, exposing the fund to significant future regression to the mean.

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 only affect the outcome in date t. In a traditional, non-persistent moral hazard setting, rewarding high output, as a general rule, helps alleviate the agency problem. However, in a myopic agency setting, rewarding high output can exacerbate just as much as it can alleviate the agency problem. Divorced from the past and the future, a stochastic output today communicates very little information about the agent's decisions. As a result, the optimal dynamic contract under myopic agency differs in significant ways from many of those that arise under traditional, non-persistent moral hazard.

The first difference is what constitutes "good" performance. 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. Traditional dynamic contracts typically have a measure that resembles counting total high output dates net low output dates. In contrast, the measure in my optimal dynamic contract tracks the number of consecutive high output dates leading up to today. Because the contract values sustained high output, a low output after a long history of high output can have a seemingly disproportionate effect on the agent's standing with the principal. This creates a cliff-like arrangement the I show is a natural way of dealing with the double edged sword of rewarding high output in a myopic agency setting. The optimality of the cliff arrangement also helps explain why, given the particular returns characteristics of hedge funds, high watermark contracts are popular.

The paradigm shift in defining good performance also has theoretical implications. In most traditional dynamic contracts, the state variable is the agent's promised value. Under myopic agency, the optimal contract's state variable is the incentive level: the *sensitivity* of the agent's promised value to performance. Therefore, the optimal contract of this paper exemplifies a new species of arrangements where promised value plays a genuinely secondary role.

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.

The traditional contract has a stationary incentive level. In contrast, the optimal contract's expected incentive level increases in a concave, convergent way over time. This type of behavior is also 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. Specifically, incentives escalate only after high output.

Lastly, I consider optimal contracting when the agent is terminated for poor performance. The resulting optimal contract with termination exhibits an upward sloping pay scale and severance pay. The upward sloping pay scale, which appears in many actual employment contracts, is a direct result of the aforementioned incentive escalation property. The myopic agency driven intuition for upward sloping pay scales differs from previous arguments that relied on combining the standard backloading intuition with firm-side incentive problems or agent risk aversion. See, for example, Lazear (1981).

Severance pay is typically difficult to produce as an optimal arrangement in traditional moral hazard models. However, the incentive escalation result of this paper provides a simple justification for severance pay - even large, mostly lump sum, severance pay. When the principal terminates the agent for poor performance, he may still want to pay the agent a potentially large amount simply to dull incentives today. By dulling incentives today, the principal would not have had to escalate incentives as much tomorrow had the agent generated high output today and avoided termination. This is a good thing because incentive escalation, which is needed to combat myopic agency, is expensive.

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. Reasons put forth include costly signaling in the presence of takeover threats (Stein, 1988), lack of takeover protection (Laffont and Tirole, 1988), myopic market expectations (Stein, 1989) and herding due to reputation concerns (Zwiebel, 1995).

Another related literature deals with innovation. The process is frequently modeled using bandit problems, which generates a dynamic not unlike the one produced by the long-term action. Specifically, by exploring an unknown arm at date 0 the player sacrifices some expected output. However, the player gains extra knowledge about which arm is the most productive. This will increase expected output at date 1. Manso (2012) embeds such a two-date bandit problem within a principal-agent framework and studies how to motivate innovation at the first date. The model shares some common features with the one-shot version of the myopic agency model.

Edmans, Gabaix, Sadzik and Sannikov (2012) consider a model of dynamic manipulation that allows the agent to trade off state-by-state future and present performance. The agency problem can be seen as a particularly virulent strain of the myopic agency considered in this paper. Their optimal contract also exhibits a form of incentive escalation.

The incentive escalation result of this paper has a flavor of Holmstrom and Milgrom (1991). Recall, that paper observes 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 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.

The present 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 persistent moral hazard, results become noticeably different from those of the non-persistent class.

The rest of the paper is organized as follows: Section 2 introduces the basic repeated

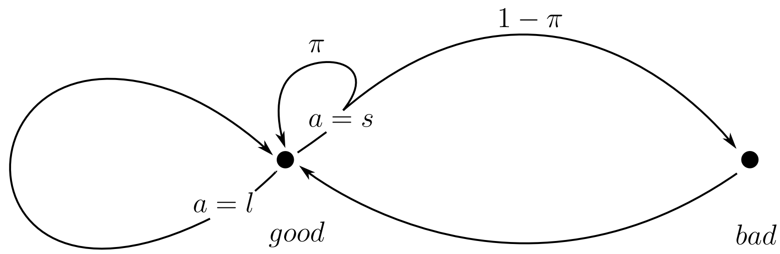

Figure 1: Transition function between the good and bad states.

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 high watermark contracts and non-persistent optimal contracts are made. Section 4 deals with incentive escalation and the option to terminate. Upward sloping pay scales and severance pay emerge as optimal arrangements. Section 5 concludes.

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: $\sigma_t = \mathit{good}$ or *bad*. If $\sigma_t = \mathit{good}$, 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: $\sigma_{t+1} = \mathit{good}$. If the agent applies the short-term action, then $\sigma_{t+1} = \mathit{good}$ with probability $\pi < 1$, and $\sigma_{t+1} = \mathit{bad}$ with probability $1 - \pi$. If $\sigma_t = \mathit{bad}$, 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 $\sigma_t = \mathit{good}$ and $a_t = l$ then the probability that the firm produces high output at date t is $p < 1$. If $\sigma_t = \mathit{good}$ and $a_t = s$ then the firm produces high output for sure at date t. If $\sigma_t = \mathit{bad}$ then the firm produces low output for sure at date t. I assume that $\sigma_0 = \mathit{good}$.

Thus, by always taking the long-term action, the agent keeps the firm in the good state and maintains a probability p of high output at each date. Any deviation 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 $\beta Q > 1 - p$ where $\beta < 1$ is the intertemporal discount factor. This assumption implies 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 $\sigma_t = \mathit{good}$. 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 $\mathbf{E}_a\big[\sum_{t=0}^{\infty} \beta^t (w_t + b1_{a_t = s})\big]$. I assume selecting the short-term action provides benefit $b > 0$ to the agent. I also assume that the agent has an outside option worth 0. This assumption eliminates the participation constraint.

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 under certain parameterizations of the model requiring the agent to take the long-term action is without loss of generality. Given Assumption (A), I define the optimal contracting problem to be the following constrained maximization:

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

$$\text{s.t.} \qquad \{a_t = l\}_{t=0}^{\infty} \in \arg\max_{\{a_t\}} \mathbf{E}_{\{a_t\}} \left[\sum_{t=0}^{\infty} \beta^t (w_t + b 1_{a_t = s}) \right]$$

Throughout the paper, if there are multiple optimal arrangements, I focus on the one that pays the agent earlier. 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. In general, the agent's promised value depends on the history of outputs up through today as well as today's state and action. So for each $h_t \in \mathcal{H}_t$, state σ_t and action a_t, define $W_t(h_t, \sigma_t, a_t)$ to be the agent's date t promised value given all the relevant information:

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

In general, the promised value is not part of the principal's information set since states and actions are hidden. However, since the paper restricts attention to only those contracts where the agent takes the long-term action all the time, it is well-defined to speak of $W_t(h_t)$, which only depends on the publicly observable h_t. 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 $\mathcal{H}_{t-1}$, Δ_t is a random variable representing the date t incentive level of the contract.

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

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

Proof. Equation (1) is necessary because it is precisely the condition that prevents one-shot deviations. Moreover, the only times when the agent can take the short term action is when the state is good. Therefore, it is sufficient to only check one-shot deviations. □

The IC-constraint is a lower bound for the incentives tomorrow as a function of the incentives today. The absolute levels of incentives do not matter per se, what matters is 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 private benefit b 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 b.

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 would never occur. The incentive level after low output is immaterial. Therefore, it is not included in the IC-constraint.

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

Now, 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 Assumption (B) implies $\Delta_t \geq 0$ for all t.

Definition. *For any $\Delta \geq 0$, the term Δ-contract will denote a contract whose initial incentive level is Δ. Define $V(\Delta)$ to be the cost to the principal of the optimal Δ-contract.*

A significant technical difference between the myopic agency setting and the traditional, non-persistent setting relates to their respective approaches to the recursive formulation of optimality. In the traditional approach, the optimal value function is defined over the agent's promised value W_t. As a result, the traditional contract is recursive over promised value. In contrast, $V(\Delta)$ will be the "optimal value function" of the myopic agency approach. Thus, the resulting optimal contract is recursive over incentives not promised value.

Lemma 2. $V(\Delta)$ *is convex and weakly increasing. So optimal Δ-contracts are randomization-proof and the optimal 0-contract is the optimal contract.*

Fix an arbitrary Δ-contract. Then the IC-constraint implies that tomorrow's incentives following high output today must be at least $\varepsilon(\Delta)$. Therefore, viewed as a contract in its own right, this date 1 continuation contract is some Δ' contract where $\Delta' \geq \varepsilon(\Delta)$. Now, if the Δ-contract being considered is the optimal Δ-contract then the date 1 continuation contract must also be the optimal Δ'-contract. Moreover, since V is weakly increasing, it must be that $\Delta' = \varepsilon(\Delta)$. Using the same logic, the date 1 continuation contract following low output must be the optimal 0-contract.

Once the continuation contracts are determined, the date 0 high and low output payments are pinned down:

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

These are the minimal payments that ensure today's incentive level is Δ. Together with the continuation contracts, they determine the Bellman equation characterizing $V(\Delta)$. Solving the Bellman equation formally solves the optimal contracting problem.

Proposition 1. *The Markov law for the state variable Δ is*

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

The optimality function $V(\Delta)$ satisfies the following explicitly solvable Bellman equation:

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

The solution $V(\Delta)$ to the Bellman equation is a piecewise linear, weakly increasing, convex function. $V(\Delta)$ not only implies the structure of the optimal contract but more generally, the structure of the optimal Δ-contract for any $\Delta \geq 0$. Viewed as a Markov process, Δ_t has stationary measure π over the discrete set $\{\varepsilon^N(0)\}_{N=0}^{\infty}$. The stationary measure is $\pi(\varepsilon^N(0)) = (1-p)p^N$.

3 A Representative Optimal Contract

For a large set of parameters, including all sets of parameters where $Q \geq 6 - 4\sqrt{2} \approx .343$, $V(\Delta)$ is a two piece piecewise linear function:[2]

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

where

- $V(0) = \frac{1}{1-\beta^2}\frac{pb}{Q}$
- $m(x) = \frac{1-p}{Q}x - (1-p)$
- $x_1 = \frac{1}{1+\beta}\frac{pb}{Q-p(1-p)}$

The optimal contract delivers expected value $V(0)$ to the agent. Since $V(0)$ is proportional to b, it is without loss of generality for the optimal contracting problem to restrict attention to always inducing the long-term action for all sufficiently small b.

[2]The condition on Q is fairly mild. Since I assume in the model $1-p < \beta Q$ and since by definition $Q \leq p$, it must be that $p > \frac{1}{2}$. Therefore, the condition on Q does not further restrict the domain of p. Moreover, the closer Q is to p, the more flexibility one has in setting the values for the other parameters. Formally, $V(\Delta)$ is a two piece piecewise linear function if and only if $\beta \leq \frac{Q-p(1-p)}{(p-Q)(1-p)}$.

As a function of incentives, today's spot contract is

$$w(1) = \begin{cases} 0 & \text{if } \Delta \in [0, x_1] \\ \frac{Q-p(1-p)}{Q}\left[\Delta - x_1\right] & \text{if } \Delta \geq x_1 \end{cases} \tag{5}$$

$$w(0) = \begin{cases} \frac{Q-p(1-p)}{Q}\left[x_1 - \Delta\right] & \text{if } \Delta \in [0, x_1] \\ 0 & \text{if } \Delta \geq x_1 \end{cases} \tag{6}$$

3.1 How and Why the Optimal Contract Works

Formally, the state variable of the optimal contract is Δ. In practice, it suffices to count the number N of consecutive high output dates leading up to today. The performance indicator N completely determines today's spot contract, incentives and continuation contract.

Thus, the optimal contract cares about *sustaining* high output and not simply generating a large number of high output dates.

As a function of N, today's incentive level is $\varepsilon^N(0)$ and is increasing concave. The limiting incentive level is

$$\overline{\Delta} := \varepsilon^\infty(0) = \frac{b}{\beta Q - (1-p)} \tag{7}$$

As a function N, today's spot contract is also determined and its pay-to-performance sensitivity is also increasing concave.

Consider the payment to the agent for high output when his performance indicator is N. Viewed as a function over N, this reward schedule is zero for low values of N and then becomes a 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 increasing rewards. Eventually, rewards for high output level off at a large sum for sufficiently good performance.[3]

The design of the reward schedule 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 task. On the other hand, the principal realizes that a large reward for high output, paradoxically, tempts the agent to deviate 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. 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.

Initially, the agent is at the bottom of the cliff. Performance determines his progress along the path. In particular, every high output moves the agent forward one step. His goal is to consume the big carrots at the top. To reach them, the agent must take the long-term

[3]In this particular class of examples, low values of N means $N = 0$. In general, the initial period of no high output reward may last longer but not less than one date.

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 progressed to 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 likely deprived of a big carrot tomorrow. For example, if $Q = p$, then gorging today ensures no consumption tomorrow. But this is a relatively minor loss compared to the other one looming on the horizon. Recall the agent's fate 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 of the cliff.

3.2 The Cliff Arrangement and High-Watermark Contracts

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? Suppose the principal eliminates the optimal contract's cliff entirely. At each date, the agent is then simply given the limiting spot contract as the performance indicator goes to infinity. It is worth noting, this spot contract's pay-to-performance sensitivity is larger than that of any of the spot contracts used in the optimal contract. Yet, under this arrangement, the agent never chooses to take the long-term action.

The cliff arrangement connects the myopic agency optimal contract to high-watermark contracts used in the hedge fund industry. The myopic agency model captures an important aspect of the agency problem in a hedge fund. The manager can take the long-term action of exerting effort to turn his innate skill into generating alpha. Or, the manager can take the short-term action by engaging in various "information-free activities that amount to manipulation... not based upon the production and deployment of value-relevant information about the underlying assets in the portfolio." (Goetzmann, Ingersoll and Spiegel, 2007).

The mapping of hedge fund returns to the binary output system of my model is a little trickier. Empirical findings suggest that the returns distributions of hedge funds typically exhibit large negative skewness and high excess kurtosis. This means much of the variance is generated by negative tail events. Thus, one reasonable mapping is to let high output denote positive alpha and low output denote the negative tail event.

My paper predicts that the solution to the agency problem is the aforementioned cliff arrangement. In practice, hedge fund managers are compensated using high-watermark contracts. Essentially, a high-watermark contract pays the manager a performance-contingent fee only when cumulative performance exceeds the running maximum, that is, the high-watermark. As the manager consistently generates alpha, more funds flow in and the performance fee typically increases. This fits well with the backloaded reward schedule of the cliff arrangement.

When the negative tail event occurs, the fund's net asset value falls well below the high watermark and performance fees disappear. Compounding the loss, the smaller fund means that the manager takes in less from his management fee. Moreover, since the gap between current net asset value and the high-watermark is very large by definition of a negative tail event, the manager will find it difficult to close the gap quickly. This implies that when things do go south, it may be a while before the manager receives significant compensation again.

It is worth noting that in particularly serious (though not all that uncommon) cases, the manager may find the gap so large that he simply gives up. That is, he closes the fund and starts anew. By doing this, the manager is effectively cushioning his fall off the cliff with the backs of his previous investors. Such a maneuver certainly dilutes the effectiveness of the high watermark's implicit cliff arrangement. That being said, the manager's new fund's size and his new performance fee will likely both be smaller. Therefore, even when a manager abandons his fund, he still experiences something comparable to a fall.

3.3 Comparing with the Non-Persistent Version of the Model

A natural comparison to make is with the non-persistent version of the model. 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.

The long-term action is relabeled as effort. Effort today is modeled to produce a probability p of generating high output today. The short-term action is relabeled shirking. Shirking today is modeled to produce a probability $1 - \beta Q$ of generating high output today. Notice, the principal replaces the short-term action's true multi-period effect on output with its present value. Shirking provides private benefit b.

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 generates 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.

Fortunately for the principal, sweeping the persistence of the moral hazard problem under the rug and using the traditional contract still induces the long-term action from the agent. That is, the traditional contract is incentive-compatible in the true model. Obviously, it is inefficient since it is not isomorphic to the optimal contract. The degree of inefficiency can be usefully quantified once I describe the traditional contract's structure.

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. 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 $\varepsilon^N(0)$ and satisfies $0 = \varepsilon^0(0) < \varepsilon(0) < \varepsilon^2(0) < \ldots < \varepsilon^\infty(0) = \overline{\Delta}$. Thus, the traditional contract over-incentivizes the agent.

Consequently, the pay-to-performance sensitivity of the traditional contract is also too high. There is, however, a small but 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 hold for pay-to-performance sensitivity.

In the traditional contract, the continuation contract is identical across histories leading up through today. Therefore, any level of incentives today must be entirely generated by today's pay-to-performance sensitivity. This is not the case in the optimal contract, where tomorrow's continuation contract is more valuable if high output is generated today. As a result the optimal contract exhibits a certain background level of incentives. Even if pay-to-performance sensitivity were zero today, the incentive level would still be positive. This means today's pay-to-performance sensitivity is responsible for generating only a portion of today's incentive level. In particular, when N tends to infinity, the optimal contract only pays the following quantity for high output:

$$\overline{w} := \overline{\Delta}\frac{1}{1+\beta}\frac{Q - \beta(1-p)(p-Q)}{Q} \tag{8}$$

Thus the optimal contract's rewards for high output are bounded above by a quantity, $\overline{w}$, strictly smaller than the traditional contract's reward for high output, $\overline{\Delta}$. This has important feasibility implications.

Suppose that the agent's private benefit 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. For example, if $Q = p$, then $\overline{w} = \overline{\Delta}/(1+\beta)$. With the optimal contract, the private benefit threshold is $1+\beta$ times larger. Moreover, even if b exceeds this larger threshold, only a fraction of the time is the optimal contract's spot contract rewarding the agent more than he produces.

4 Incentive Escalation and Termination

The incentive structure is a key difference between the optimal and traditional contracts. The incentives of the optimal contract are smaller and have non-stationary dynamics. Specifically, incentives increase after high output. Over time, the expected incentive level $\mathbf{E}\Delta_t$ also increases and converges as N settles into its stationary distribution.

Thus, the optimal contract exhibits a form of incentive escalation. It predicts that agents with histories of sustained high output should receive the highest incentives and pay-to-performance sensitivities. The incentive 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 incentive escalation result in my model is fully compatible with a risk averse agent, it is not dependent on risk aversion.

In general, introducing risk aversion tends to escalate pay-to-performance sensitivity in optimal dynamic contracts. But this escalation should not be confused with the result in my paper. Suppose over time the agent's pay level is increasing, either due to accumulating good performance or contractual backloading. Then if the agent is risk averse, wealth effects kick in and may force the principal to escalate pay-to-performance sensitivity simply to induce the same level of incentives as before. This type of escalation result is fundamentally different from the incentive escalation of the myopic agency optimal contract. Since incentives measure *promised-value-to-performance* sensitivity, a wealth effects driven escalation in pay-to-performance sensitivity does not automatically count as any escalation in incentives.

4.1 Optimality with Termination

The incentive escalation property is the main reason I chose to not include in the original model an important aspect of contracting: termination. If the only agency problem is myopic agency, then giving the principal the ability to terminate the current agent and contract a new one produces a perverse optimal termination rule. When the current agent is doing well, the principal wants to terminate him due to the incentive escalation that occurs after high output. And when the agent is doing poorly, the principal paradoxically wants to keep him. This unwillingness to terminate the agent when he is doing poorly is a consequence of a severance pay result that I will discuss shortly. Ultimately, it boils down, again, to the incentive escalation property.

Thus, even though these termination tendencies are perverse, they are rational responses given the nature of this particular second-best setting. Of course, by focusing on myopic agency, the model ignores important facets of the principal-agent relationship such as reputation concerns and power dynamics. These considerations will cause the principal to strongly reconsider how to optimally terminate the agent. Instead of adding these other elements to the model and letting the termination choice arise endogenously, I will, for the sake of simplicity, fix a "natural" benchmark termination rule. This rule will essentially serve as a proxy for those unmodelled elements.

The rule is: terminate the agent if and only if low output is generated. This rule is the simplest one that subscribes to the more intuitive and empirically justified idea that the agent should be fired for poor not good performance. It is also the endogenously optimal termination rule of the non-persistent version of the myopic agency model. And lastly, despite the previous discussion to the contrary, this rule is actually a natural one to consider for the myopic agency setting. To see this, notice that the optimal contract without termination regenerates if and only if the agent produces low output. Thus quite literally, the building block of the entire optimal contract without termination is itself restricted over the domain of histories implied by the proposed termination rule: $\{0, 10, 110, 1110, 11110\ldots\}$. Not surprisingly, this building block is a remarkably good approximation of the actual optimal contract with the imposed termination rule.

The derivation of the optimal contract with termination mirrors the original analysis. The state variable is still Δ, and the Bellman equation characterizing optimality is only slightly changed: $V(\Delta) = \max\{\beta V(\varepsilon(\Delta)) - (1-p)\Delta, p\Delta\}$. The equation can be solved and the optimal contract with termination can be explicitly characterized. Mirroring the previous approach, I now describe a representative optimal contract with termination and analyze it.

Let w_t denote the agent's salary in year t if high output is produced and he is retained. Let S_t denote the agent's severance pay in year t if low output is produced and he is terminated. Suppose parameters of the model satisfy $\beta \leq \frac{Q-(1-p)p}{pQ}$. Then the optimal contract's salary, severance pay and incentive level are as follows:

- Salary: $w_0 = 0$. For $t > 0$,

$$0 < w_1 < \ldots < w_t = \varepsilon^t(0)p\left[\frac{Q-(1-p)p}{Q}\right] - \frac{bp^2}{Q} < \ldots w_\infty = \overline{\Delta}p(1-p\beta)$$

- Severance pay: $S_0 = \frac{pb}{Q}$. For $t > 0$, $S_t = 0$.

- Incentive level: $0 = \Delta_0 < \ldots < \Delta_t = \varepsilon^t(0) < \ldots \Delta_\infty = \overline{\Delta}$.

4.2 Upward Sloping Pay Scale

When the agent is young, his salary is zero. After the initial stage, salary becomes nonzero and increases with tenure. Eventually, as the agent's tenure becomes long, his salary approaches the steady state w_∞. In contrast, the traditional contract with termination has a fixed salary $\overline{w} = \overline{\Delta}(1-p\beta)$ which is $1/p$ times as large as w_∞. The salary comparison between the optimal and traditional contracts with termination closely mirrors the previous pay comparison between the no termination contracts. I will not say anything more about this.

The upward sloping pay scale w_t faced by the agent in the optimal contract with termination is a direct result of the incentive escalation property. It is an important feature of many real life employment contracts and this paper is certainly not the first to have reproduced it in a theoretical setting. The standard incentives-based arguments for upward sloping pay scales begin with the backloading idea: In general, it is not optimal to tie pay, on a date-by-date basis, to marginal product or the flow payoff of the outside option. By delaying a payment, the principal can efficiently reuse the threat of withholding that payment to motivate the agent across multiple dates. This is particularly useful when the agent's participation constraint is a binding constraint.[4] What the idea suggests is that the principal should pay the agent only on the last date. Technically speaking, loading all the payments onto the last date does produce an upward sloping pay scale, albeit an extreme one. To get a more realistic shape, a counterbalance is typically introduced.

For example, Lazear (1981) argues that the principal also has incentive problems. In particular, there is a risk that he may renege on payments promised to the agent. In this situation, loading everything onto the last date is far too risky and so the agent's pay increases more gradually over time. Another argument used by many authors appeals to risk aversion. With a sufficiently risk averse agent, it is simply far too expensive to exclusively use the last date's pay to generate the necessary incentives that motivate the agent throughout the entire contract. Again, a more smoothly upward sloping pay scale results.

In my model, the principal is fully committed and the agent is risk neutral. Therefore, myopic agency and incentive escalation combine to provide a fully independent reason for having an upward sloping pay scale.

4.3 Severance Pay

The incentive escalation property provides a simple intuition for severance pay - potentially even large, mostly lump sum, severance pay. Suppose after some low output event it is optimal to terminate the agent. Facing such a situation, the principal may still want to pay the agent a nontrivial amount simply to dull incentives today. Why might the principal want to dull incentives today? By dulling incentives the principal need not escalate incentives as much tomorrow following high output. This is a good thing, because incentive escalation, which is an important tool used to combat the agency problem, is expensive.

Therefore, severance pay for poor performance is a mechanism to temper today's incentives so as to preserve the future viability of the principal-agent relationship should it continue.

The nontrivial S_0 of the optimal contract with termination demonstrates that the severance pay intuition works in practice. If the principal deletes S_0 then one of two things will happen: either the agent takes the short-term action or the principal must increase the incentives at all dates $t > 0$ to compensate for the increased date 0 incentives. Neither is palatable.

The severance pay S_0 is proportional to the private benefit b. Therefore, it's value relative to the surplus generated by the firm is ambiguous. What is not ambiguous is it's value relative to the contract. Since the contract's initial incentive level is zero, the severance pay S_0 is equal in value to all of the agent's future expected earnings had he generated high output today and avoided termination. As this example demonstrates, even relatively massive, lump sum severance pay can be supported as optimal arrangements.

[4] Notice though this is not the case for the myopic agency model.

That being said, the example should not be interpreted as a blanket justification of golden parachutes. The size of the severance pay today is proportional to the degree to which today's incentive level can be compressed. In the optimal contract with termination, there is no bound on incentive compression at date 0. But if in addition, there was, say, a non-persistent moral hazard problem or some other constraint on the degree of incentive compression, the severance pay would be smaller. In general, the intuition is compatible with severance packages both large and small, immediate or vested, lump sum or performance-contingent.

One curious feature of the optimal contract with termination is that severance pay is decreasing with tenure. In fact, if tenure is sufficiently long, the optimal contract with termination gives no severance pay.[5] This unusual dynamic is an artifact of a simplifying assumption of the model: the agent has a worthless outside option. If, instead, the agent's outside option increases with tenure, which seems likely if the labor pool is heterogeneous with respect to ability, then it's possible for severance pay to be also increasing.

In contrast, the traditional contract with termination exhibits no severance pay. The downside of increasing today's incentives that drives the severance pay intuition in the myopic agency setting is simply not applicable in settings with non-persistent moral hazard. The IC constraint in non-persistent settings is a lower bound on each date's incentives. Deleting severance pay for low output only further ensures that the IC constraint is not violated. This intuition explains why the dynamic contracting literature traditionally has had difficulty explaining severance pay.

The severance pay result in this paper is not driven simply by the fact that actions have persistent effects. It is worth distinguishing between the intuition described in this paper coming from myopic agency in particular and the rudimentary motivation for severance pay coming from action persistence in general. The latter idea simply argues that since there is persistence, actions taken by the agent continue to have effects on the firm post-termination. Therefore, the principal may want to have part of the agent's compensation postponed to after termination - hence severance pay. This is clearly not the same idea as the one relying on myopic agency and incentive escalation. The intuition for severance pay in this paper, for example, can also explain lump sum severance pay, which is something the action persistence intuition cannot.

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 choice and creates a tricky contracting setting where simply rewarding high output is no longer guaranteed to eliminate the agency problem. I show that the optimal contract in this setting differs in significant ways from traditional dynamic optimal contracts. The

[5] In this example, sufficiently long means $t > 0$. In general, the initial period with severance pay may last longer but not less than one date.

myopic agency optimal contract values sustained high output. Its incentive level is nonstationary, escalating over time and after high output while remaining strictly below the incentive level of the traditional contract. Also, severance pay may accompany termination for low output. In addition, the myopic agency optimal contract provides new perspectives on high watermark contracts and upward sloping pay scales. 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 2 and Proposition 1. Let $\mathscr{C}$ be the space of all nonnegative convex functions defined on $[0, b/(\beta Q - (1-p)]$ with slope at most p and bounded above by $M := pb/[(1-\beta)(\beta Q - (1-p)]$. Under the L^∞ norm, $\mathscr{C}$ is a compact Banach space. Define the operator $Tf(x) := \max\{\beta f(\varepsilon(x)) - (1-p)x, px + \beta \min\{f(x)\}\}$. It is easy to check that T is well-defined endomorphism of $\mathscr{C}$. Note this operator differs slightly from the implied operator in Equation (3).

Let f, $g \in \mathscr{C}$. For any x, $|Tf(x) - Tg(x)| \leq \max\{|T_a f(x) - T_a g(x)|, |T_b f(x) - T_b g(x)|\}$ where $T_a f(x) := \beta f(\varepsilon(x)) - (1-p)x$ and $T_b f(x) := px + \beta \min\{f(x)\}$. 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.

Consider two model realizations of Q: $\hat{Q} < \tilde{Q}$, with corresponding spaces $\hat{\mathscr{C}}$, $\tilde{\mathscr{C}}$ and operators $\hat{T}$, $\tilde{T}$. Let $\hat{f} \in \hat{\mathscr{C}}$ and $\tilde{f} \in \tilde{\mathscr{C}}$. Then I say $\hat{f} > \tilde{f}$ if and only if on the domain of $\tilde{\mathscr{C}}$, $\hat{f} - \tilde{f}$ is weakly increasing, nonnegative and not identically zero.

It is easy to show that if $\hat{f}$ and $\tilde{f}$ are both nondecreasing and $\hat{f} \geq \tilde{f}$ then $\hat{T}\hat{f} > \tilde{T}\tilde{f}$. Consider any model realization with $Q := p$, space $\mathscr{C}_p$ and operator T_p. It is easy to show that $T_p^n 0$ is nondecreasing for all n. Then $T^n 0$ is nondecreasing for all n in general. Therefore, the unique fixed point $T^\infty 0$ is nondecreasing.

Fix a $\Delta \in (0, b/(\beta Q - (1-p)))$. 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 v piece by piece starting from the left:

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

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

As a function of Δ, $d(\Delta) := v(0) - v(\varepsilon(0))$ is strictly increasing. For Δ sufficiently close to 0, $d(\Delta) < 0$. For Δ sufficiently large, $d(\Delta) > 0$.[6] Let Δ^* satisfy $d(\Delta^*) = 0$. Then the

[6] In fact, it suffices for Δ to be large enough so that $m^{N-1}(p) \leq 0$.

corresponding function v^*, when restricted to $\mathscr{C}$, is the unique fixed point of T. Since the fixed point is nondecreasing, v^* is the unique solution to the Bellman equation (3). $\square$

References

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

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

[3] 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

[4] 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

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

[6] Lazear, E. (1981) "Agency, Earning Profiles, Productivity, and Hours Restrictions," *The American Economic Review* Vol. 71, pp. 606 - 620

[7] Laffont, J.-J. and J. Tirole (1988) "Repeated Auctions of Incentive Contracts, Investment, and Bidding Parity with an Application to Takeovers," *RAND Journal of Economics* Vol. 19, pp. 516-537

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

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

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

[11] Zwiebel, J. (1995) "Corporate Conservatism and Relative Compensation," *Journal of Political Economy* Vol. 103, pp. 1-25