Myopic Agency

John Y. Zhu[1]

November 6, 2014

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. It features a cliff-like arrangement that rewards high output today based on the streak of consecutive high outputs the agent has generated leading up to today: The longer the streak, the larger the reward. The optimal contract can be implemented as a bonus bank. Bonus banks feature prominently in the recent debate on bonus reform. I shed light on the opposing arguments driving this debate by formally comparing the myopic agency optimal contract and optimal contracts that arise from traditional effort-shirking agency models.

JEL Codes: C73, D86, J33, J41, M52
Keywords: principal-agent, dynamic contract, myopia, robust contract, investment, bonus deferral, bonus bank, long-term, short-term, persistent moral hazard

[1]zhuyiran@wharton.upenn.edu, The Wharton School, University of Pennsylvania. I thank Sugato Battacharya and Yuliy Sannikov for suggesting a number of key improvements to the paper and Richard Kihlstrom for providing valuable guidance throughout the development of the paper. I thank Daniel Neuhann for excellent research assistance. I am grateful to Robert Anderson, Alex Edmans, Lily Fang, 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, 2013 China International Conference in Finance and INSEAD 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 tying the agent's compensation today to the streak of *consecutive* high outputs produced leading up to today. The longer the streak, the larger the compensation with compensation size leveling off when the streak becomes sufficiently long. Theorem 1 explicitly characterizes this contract. In Section 3.1, I implement this contract as a bonus bank contract: Each date, 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. As the agent continues to produce high output date after date, 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 bonuses he received on previous dates that have not yet fully vested. Eventually, the rate at which old bonuses fully vest and new bonuses come in balance out. At this point, each date's total bonus pay levels off. Finally, if at any date the agent fails to produce high output, then his high output streak is broken. Not only does the agent not receive a bonus this date, but also all of the remaining unpaid portions of previously issued bonuses get wiped out.

Why does the bonus bank arrangement work? When the agent's streak of consecutive high output production is short, most of his compensation is backloaded because there are not a lot of bonuses that can contribute to his pay today. Remember, at this point, even if the agent has had a long history of producing high output, the many bonuses he received for that performance have either fully vested or have been wiped out by a recent low output that "cleaned the slate." In order to accrue (or re-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 over and over again. Thus, at least initially, the agent is motivated to take the long-term action. As the agent continues to build his streak and receive new bonuses, his compensation eventually reaches the large promised level. But this means high output rewards are no longer backloaded. At this point, interestingly, the agent is still motivated to 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. If this happens, all remaining bonus payments are wiped out and the agent has 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 specifically to combat managerial myopia. When will a principal face a repeated myopic agency problem like in the model? A natural situation relates to R&D investment. 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, inflating net asset value today but exposing the fund to significant future tail risk.

A common property in these examples is that the productivity ranking of the agent's hidden actions is reversed over time: In terms of output today, the short-term action is the "good" action; in terms of output in the future, the long-term action is the "good" action. This type of *twisted* moral hazard cannot be properly modeled when actions have non-persistent effects or have persistent effects that are faded copies of today's effect. As a result, the design of the optimal contract is 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 the next date (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 the moral hazard is not twisted, rewarding high output, as a general rule, helps alleviate the agency problem. However, as I just argued in the previous paragraph, rewarding high output in a myopic agency setting can exacerbate the agency problem just as much as alleviate it. A high output can signal the agent took the long-term action the previous date or it could mean the agent is taking the short-term action today. 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* *output* across time. The structure of the myopic agency optimal contract, with its emphasis on high output streaks, reflects this requirement.

1.1 A Comparative Discussion

The optimal dynamic contract under myopic agency differs in significant ways from many of those that arise under traditional effort-shirking moral hazard. Unlike the streaks-based measure of good performance used by the myopic agency optimal contract, traditional optimal dynamic contracts typically have a measure that resembles total or average output. The aggregation is insensitive to the specific sequence of output. This type of order-independent measure of performance and its close approximations have dominated dynamic contracting theory since its inception. Seminal examples include the cumulated performance S of Radner (1981, 1985) and the linear aggregator of Holmstrom and Milgrom (1987).

In Section 3.2, I consider the traditional effort-shirking version of the myopic agency model and solve for its optimal contract, which I call the *traditional contract*. I formally compare the traditional and myopic agency optimal contracts. The traditional contract's incentive level Δ is constant at all times just like many optimal dynamic contracts that sustain effort.[2] However, the myopic agency optimal contract's incentive level Δ_t rises and falls with the length of the high output streak. In fact, Δ_t is the key state variable of the model. Despite its movement, Δ_t is always strictly smaller than Δ, which implies that the traditional contract always gives larger bonuses than the myopic agency optimal contract. See Section 3.2, in particular, the remark and Lemma 3.

I leverage this comparative analysis to shed light on the recent debate on bonus pay reform. The key starting point is to realize: 1. Many of the firms that are in the middle of the debate face myopic agency problems. 2. Given the historical prominence of traditional effort-shirking moral hazard, these firms may think in the traditional effort-shirking framework when designing their incentive contracts. By examining what would happen if the principal in the myopic agency model mistakenly uses the traditional contract on the agent, I can reconcile a number of the opposing views that continue to drive the bonus pay debate. See the second half of Section 3.2 for details.

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 frame of reference. Therefore, the conclusion in my model is not that incentives should be

[2] The incentive level is a standard quantity in the contracting literature. In general, it is the sensitivity of promised value to fundamentals. Specifically, in my model, it is the difference between the agent's total payoff following high and low output. See Definition 1.

low-powered, but that incentives should start low and optimally escalate as the high output streak increases.

Another related paper is Hoffmann and Pfeil (2013), which explores a multi-task extension of the DeMarzo and Sannikov (2006) dynamic agency model where the agent can both secretly steal cash and use cash for investment purposes. The principal does not want the agent to steal but does want the agent to invest. Unfortunately for the principal, investment is unobservable. This creates the main tension of the model: When the principal sees low cash flow, he can't tell if it's because of stealing or because of investment. The paper solves for the optimal contract and studies second-best investment distortions.

While my 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 (2011) 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.

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 and linear cost. 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 when the moral hazard is not twisted. Sannikov (2014) considers a Brownian model of persistent moral hazard that in some ways generalizes the setting of Jarque (2010). He focuses on two 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 case. His paper derives a clever representation of the incentive level and uses it to explicitly characterize the optimal contract's compensation in the large-firm case. In his analysis, the continuous-time analog of Δ_t plays an important role and is also a state variable.

2 Repeated Myopic Agency

A principal contracts an agent to manage a firm at dates $t = 0, 1, 2 \dots$ At each date t, the firm can be in one of two states: $\sigma_t = \textit{good}$ or *bad*. If $\sigma_t = \textit{good}$, then the agent can apply one of two hidden actions: a long-term action $a_t = l$ or a short-term action $a_t = s$. The

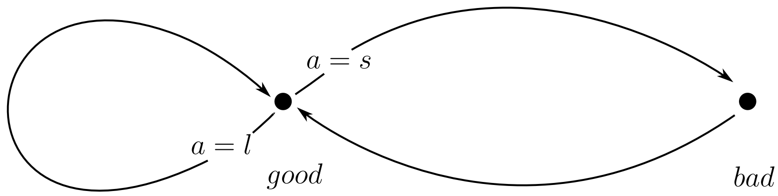

Figure 1: Transition function between the good and bad states.

long-term action requires an extra effort cost $c > 0$ compared to the short-term action. I normalize the effort cost of the long-term action to zero and assume that the short-term action provides "private benefit" worth c. This is without loss of generality and is done purely because certain expressions are simpler under the private benefit model. Everything goes through if I instead assume that the agent suffers $-c$ taking the long-term action. 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 the firm moves to the bad state: $\sigma_{t+1} = bad$. Finally, 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 $X_t = X$ 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 $p \in (0, 1)$. 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 $(1 - p)X$ and lowers expected output tomorrow by pX. I assume that

$$(1 - p)X + c < \beta pX \tag{1}$$

where $\beta \in (0, 1)$ is the discount factor. This assumption says that the gain today from taking the short-term action is outweighed by the present discounted loss tomorrow factoring in private benefits. It is the condition for the long-term action to be first-best optimal.

What is Myopic Agency?

Consider an agent who is supposed to invest in quality projects to help the firm innovate. If taking this long-term action only required cash then the principal could simply command the agent to throw money at R&D and there would be no agency problem. However, investing in projects alone is not enough to generate innovation. The agent needs to find quality projects and if this requires hidden effort then the principal faces a nontrivial agency problem. Without proper incentives, the agent will not exert the necessary effort to find a quality project and may instead deviate to the short-term action of financing safe projects that he knows can produce some quick returns without generating real innovation.

But this is not yet myopic agency. If each project's cash flow is distinguishable, then inducing effort is fairly straightforward because the agent's compensation can be tied to individual project performance: The principal simply needs to avoid prematurely rewarding a high project output early in the project's life since this suggests the short-term deviation. Unfortunately, in practice, a firm's investments may not be so neatly separated into individual projects and it's revenue into individual project output streams. It may be that all the principal has to contract on is some aggregation of the firm's various sources of output. In this case, high output today can either signal the long-term action yesterday or the short-term action today. Rewarding high output today can either encourage the long-term action yesterday or the short-term action today.

This is myopic agency. And trying to induce the long-term action *both* yesterday and today is the challenge of contracting under myopic agency.

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, if the value of the firm is sufficiently large compared to the cost of the optimal contract without imposing Assumption A, then Assumption A will be automatically satisfied. In particular, by making X sufficiently large or c sufficiently small, assuming the principal always wants to induce the agent to take the long-term action is without loss of generality. I will revisit this assumption in Section 3.2 where I discuss the practical importance of the optimal contract under Assumption A when Assumption A is without loss of generality. See Lemma 3 and the surrounding discussion. More generally, the solution to the sustained long-term action case will serve as an important benchmark for future analyses of the unconstrained optimal contracting problem.

Definition of a Contract

At each date t, the principal makes a monetary transfer w_t to the agent. The agent is protected by limited liability, so $w_t \geq 0$. 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 $\sigma_t = \textit{good}$. 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 $\mathbf{E}_a\big[\sum_{t=0}^{\infty} \beta^t (w_t + c1_{a_t=s})\big]$.

My definition of a contract implicitly assumes that the agent cannot report his hidden information σ_t to the principal. This is without loss of generality but is special to my setup. It is due to the fact that the firm is always in the *good* state on the equilibrium path. For the same reason a contract doesn't need to ask the agent to report what hidden action he has taken. Later, in Section 4.3, I will consider more general settings in which many different states can occur on the equilibrium path and where restricting communication is with loss of generality. However, I will show how the results I'm about to derive can be adapted to those more general settings as a form of robust contracting.

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

This is a standard assumption (see for example Debreu, 1986, and Innes, 1990). It means that if the output is high, the agent has the ability to secretly throw it away and make the principal think that output is low. 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:

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

s.t. agent does not want to dispose output nor take the short-term action.

I now derive the mathematical conditions that capture these incentive constraints.

Incentive-Compatibility

Let h_t denote the history of firm outputs up through date t. It is a binary sequence of length $t+1$. Define $h_{-1} := \emptyset$. In general, the agent's promised value W_t depends on the history of outputs up through yesterday as well as today's state. So for each h_{t-1} and state σ_t, define $W_t(h_{t-1}, \sigma_t) := \mathbf{E}_a \left[\sum_{i=t}^{\infty} \beta^{i-t} (w_i + c 1_{a_i = s}) \mid h_{t-1}, \sigma_t \right]$.

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

$$W_t(h_{t-1}) := \mathbf{E}_{\{a_t=l\}_{t=0}^{\infty}} \left[\sum_{i=t}^{\infty} \beta^{i-t} w_i \;\middle|\; h_{t-1} \right]$$

which only depends on the publicly observable h_{t-1} and is known to the principal.

Definition 1. *For each history h_t, define the date t ex-post promised value to be $w_t(h_t) + \beta W_{t+1}(h_t)$. For each h_{t-1}, define the agent's date t incentive level to be the difference between the two date t ex-post promised values that can follow from h_{t-1}:*

$$\Delta_t(h_{t-1}) := [w_t(h_{t-1}X) + \beta W_{t+1}(h_{t-1}X)] - [w_t(h_{t-1}0) + \beta W_{t+1}(h_{t-1}0)]$$

Δ is the key state variable of the model and the IC-constraints for inducing the long-term action all involve Δ. I split the derivation of the IC-constraints into three cases: when deviation is in action only; when deviation is in output disposal only; when deviation is in both action and output disposal.

Lemma 1. *Taking the long-term action is better than taking the short-term action if and only if at each date t and after each history h_{t-1},*

$$\Delta_{t+1}(h_{t-1}X) \geq \varepsilon(\Delta_t(h_{t-1})) := \frac{(1-p)\Delta_t(h_{t-1}) + c}{\beta p} \tag{2}$$

Proof. See appendix. □

The IC-constraint is a lower bound for the incentive level tomorrow as a function of the incentive level today. Or equivalently, it is an upper bound on the incentive level today as a function of the incentive level 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 incentive level is today, the more tempting it is to take the short-term action today. Therefore, the incentive level tomorrow must also keep pace to ensure that the agent properly internalizes the future downside of taking the short-term action today.

This myopic agency IC-constraint differs from the traditional moral hazard IC-constraint which is typically an absolute lower bound on incentives. The difference affects how one thinks about what makes a contract good or bad. The basic idea of traditional contract theory is if you want the agent to do the right thing, you have to provide him with enough incentives. The key word being *enough* - this is where the absolute lower bound comes in. The implication is that bad contracts are those with incentives that are too small. With myopic agency, balancing incentives takes precedence. A bad contract under myopic agency isn't necessarily one with small incentives - in fact, the optimal contract starts with incentive level zero - but rather one with temporally unbalanced incentives.

On the technical side, notice that (2) only involves on-equilibrium promised values. In principle, incentive-compatibility involves comparing on- and off-equilibrium promised values. In the appendix, I explain how this is possible by first showing that the agent's off-equilibrium promised value following a one-shot deviation can be expressed as a function of on-equilibrium promised values. (2) is precisely the condition that prevents such one-shot deviations. I then show that checking for one-shot deviations is sufficient.

The IC-constraint when deviation is in output disposal only is $\Delta_t \geq 0$. Mirroring Lemma 1, the IC-constraint when deviation is in both action and output disposal is

$$\Delta_{t+1}(h_{t-1}0) \geq \gamma(\Delta_t(h_{t-1})) := \frac{-p\Delta_t(h_{t-1}) + c}{\beta p} \tag{3}$$

(2) and (3) along with the nonnegativity constraint on Δ_t completely characterize incentive-compatibility. From now on, the term contract means incentive-compatible contract.

The Optimal Contracting Problem

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

s.t. for all t

$$\begin{aligned} \Delta_{t+1}(h_{t-1}X) &\geq \varepsilon(\Delta_t(h_{t-1})) \\ \Delta_t(h_{t-1}) &\geq 0 \\ \Delta_{t+1}(h_{t-1}0) &\geq \gamma(\Delta_t(h_{t-1})) \end{aligned}$$

I now recursively characterize the optimal contract. Due to equal discounting, the timing of pay is largely irrelevant and there are many contracts that achieve the optimum. I focus on the optimal contract with the fastest payment path. This contract is robust to small perturbations to the players' discount factors that make the agent less patient than the principal.

Fix a $\Delta \geq 0$ and call any contract with initial incentive level equal to Δ a "Δ-contract." Consider the optimal Δ-contract with fastest payment path and let $C(\Delta)$ denote its cost to the principal. (2) implies that tomorrow's incentive level following high output today must be at least $\varepsilon(\Delta)$. Intuitively, $C(\cdot)$ should be weakly increasing - providing more incentives to the agent should be costly. This conjecture will be proved in the appendix. Given that it is true, tomorrow's incentive level following high output today should be exactly $\varepsilon(\Delta)$ and the continuation contract should be the optimal $\varepsilon(\Delta)$-contract with fastest payment path. Similarly, tomorrow's continuation contract following low output today should be the optimal $\gamma(\Delta) \vee 0$-contract with fastest payment path.

I have now shown that optimal Δ-contracts with fastest payment paths are recursive over themselves with state variable Δ. To complete the recursive characterization, I deduce the Bellman equation for $C(\Delta)$.

In the optimal Δ-contract, either the initial high output payment or the initial low output payment is 0. Otherwise, one can subtract out the common portion to get a cheaper Δ-contract which is a contradiction. If the initial high output payment is zero, then the contract's cost given an initial high output is $\beta C(\varepsilon(\Delta))$. This then implies that the contract's cost given an initial low output payment must be $\beta C(\varepsilon(\Delta)) - \Delta$. Since the probability of high output is p, the total expected cost is $C(\Delta) = pC(\varepsilon(\Delta)) + (1-p)(C(\varepsilon(\Delta)) - \Delta) = \beta C(\varepsilon(\Delta)) - (1-p)\Delta$. Using a similar argument, it is easy to show that if the initial low output payment is zero, then $C(\Delta) = \beta C(\gamma(\Delta) \vee 0) + p\Delta$. Finally, limited liability requires that $C(\Delta)$ equals the larger of the two expressions.

I can now write down the Bellman equation characterizing $C(\Delta)$. Solving the equation formally solves the optimal contracting problem.

Theorem 1. *The optimal cost function $C(\Delta)$ satisfies the following Bellman equation:*

$$C(\Delta) = \max\left\{\beta C(\varepsilon(\Delta)) - (1-p)\Delta,\ \beta C(\gamma(\Delta) \vee 0) + p\Delta\right\} \tag{4}$$

The solution is a weakly increasing two-piece piecewise linear function:

$$C(\Delta) = \begin{cases} \frac{c}{1-\beta^2} = \beta C(\varepsilon(\Delta)) - (1-p)\Delta & \text{if } 0 \leq \Delta \leq \frac{c}{p(1+\beta)} \\ \frac{\beta c}{1-\beta^2} + p\Delta = \beta C(\gamma(\Delta) \vee 0) + p\Delta & \text{if } \Delta \geq \frac{c}{p(1+\beta)} \end{cases}$$

The corresponding Markov law for the incentive level takes Δ to $\varepsilon(\Delta)$ following high output and to $\gamma(\Delta) \vee 0$ following low output.

The optimal contract's incentive level starts at 0 and stays in the set $\{0\} \cup \{\varepsilon^n(0)\}_{n\geq 1}$ where $\varepsilon^n(0) := \varepsilon(\varepsilon^{n-1}(0))$. Specifically,

$$0 \to \varepsilon(0) = \gamma(0) \quad \text{following high or low output} \qquad \varepsilon^n(0) \to \begin{cases} \varepsilon^{n+1}(0) & \text{following high output} \\ 0 & \text{following low output} \end{cases}$$

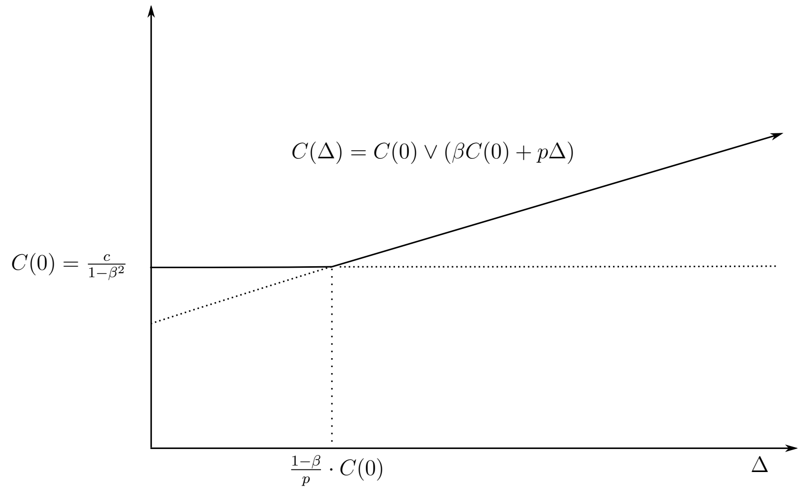

Figure 2: The optimal cost function $C(\Delta)$.

As a function of n, $\varepsilon^n(0)$ is increasing and converges to $\varepsilon^\infty(0) = \frac{c}{\beta p - (1-p)}$.

The agent is not paid following low output or if the incentive level is 0. If today the incentive level is $\varepsilon^n(0)$ and high output is produced, then the agent is rewarded with payment:

$$r(n) := p\left[\varepsilon^n(0) - \frac{c}{p(1+\beta)}\right] = pc\left[\frac{1-\left(\frac{1-p}{\beta p}\right)^n}{\beta p - (1-p)} - \frac{1}{p(1+\beta)}\right] \tag{5}$$

As a function over n, the agent's high output payment is increasing, concave and convergent.

It is straightforward to check that the solution to (4) is indeed the proposed piecewise linear function. The corresponding Markov law for Δ was derived in the discussion leading up to Theorem 1. The characterization of the optimal contract's Markov law follows from the general Markov law and the following two facts: $\varepsilon(0) = \gamma(0)$ and $\gamma(\varepsilon^n(0)) < 0$ for all n. It is straightforward to verify that the agent's high output payment is $r(n)$ given the explicit formula for $C(\Delta)$.

3 The Optimal Contract as a "Cliff"

Formally, the state variable of the optimal contract is the incentive level Δ. In practice it is useful to think of the contract as tracking the number N of *consecutive* high outputs produced leading up to today. As the agent continues to produce high output date after date, the performance indicator N increases one by one. When a low output is finally produced, N drops to zero. The agent is paid today if and only if N is positive and high output is produced. Thus, the agent must have at least a modest history of producing high output before he can be rewarded for producing high output today. The size of the payment is $r(N)$ of (5). Thus, the agent's high output payment today is an increasing, concave, and convergent function of the streak of consecutive high outputs produced leading up to today.

For example, consider the following sequence of output realizations: 1110001. By inspection, the evolution of N is $0, 1, 2, 3, 0, 0, 0$ and the payment sequence is:

$$0, p\left[\varepsilon(0) - \frac{c}{p(1+\beta)}\right], p\left[\varepsilon^2(0) - \frac{c}{p(1+\beta)}\right], 0, 0, 0, 0$$

The optimal contract cares about the pattern of output, valuing streaks of high output over dispersed high output. To see this point starkly, consider another sequence of output realizations: 1010101. It represents the same aggregate output as the previous sequence. However, the payment sequence stays at 0.

The reader may instinctively feel that order dependence is unnatural. After all, the utility function is, up to discounting, order independent. So why should the optimal contract care about order? The reason is because the optimal contract, while maximizing utility, must do so subject to incentive-constraints. And for these incentive-constraints, output pattern does matter. 1010101 looks very much like the agent is taking the short-term action whenever he can. On the other hand 1110001 means the agent must have taken the long-term action on multiple dates and therefore should be rewarded more than 1010101. This order sensitivity of the incentive-constraints is novel. In a traditional effort-shirking world, is 1010101 any more or less incriminating than 1110001?

The next issue to settle is the importance of streaks relative to other potential patterns the optimal contract could care about. To understand why streaks should matter, it helps to re-imagine the optimal contract as a path up a cliff where N measures how high up the agent is on the cliff. Initially, the agent is at the bottom (N starts at 0) and all the big rewards are backloaded toward the top of the cliff ($r(N)$ is large only when N is high). This induces the agent to repeatedly take the long-term action. The reason is that while the short-term action helps the agent produce high output today, it does not help produce high output over and over again, which is what's needed to reach those backloaded rewards.

Eventually, as the agent continues to produce high output and N continues to increase, the agent reaches the top. The rewards stop getting much larger as $r(N)$ levels off and backloading ceases. At this point, without a significantly larger reward on the horizon, what's to prevent the agent from taking the short-term action?

The answer is the cliff itself. *When the agent reaches the top, the fear of falling off becomes an effective substitute to backloading for inducing the long-term action.*

As the cliff analogy demonstrates, the optimal contract motivates the agent in two mutually reinforcing ways. Initially, the backloaded nature of the high output payment 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 payment schedule is mapped against the optimal contract's novel performance indicator N. Because N measures streaks, when it drops, it drops precipitously. This creates the "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.

The cliff-contract differs from optimal-contracts seen in traditional effort-shirking settings. In those settings, optimal dynamic contracts typically care about something that roughly approximates aggregate output. This inattention to the pattern of high output effectively means there's no cliff. I will revisit this point in Section 3.2 where I conduct a formal comparison between the cliff-contract and the optimal contract emerging from the traditional effort-shirking version of my model.

Even if always inducing the long-term action is not optimal, the cliff-contract provides insight regarding what a reasonable course of action might be. The cliff-contract reveals that optimally inducing the long-term action becomes expensive only when the high output streak gets large - that is, when the agent is high up on the cliff. If the agent is unlucky and keeps producing low output, the principal is happy to keep using the cliff-contract: he is getting the long-term action from the agent, and he is paying the agent very little. However, if always inducing the long-term action is not optimal, then at some point, the high output streak gets large enough that continuing to induce the long-term action becomes too costly. At this point, if the principal knows he has to let the agent take the short-term action, then (2) no longer needs to be respected and there is no longer a lower bound constraining the incentive level tomorrow following high output today. The principal might as well reset both the performance indicator N and the corresponding incentive level $\varepsilon^N(0)$ back to zero. This reset figuratively kicks the agent back down to the bottom of the cliff. The agent will of course best-respond by taking the short term action, but the principal expects this anyways and the upside is that the reset makes it cheap again to induce the long-term action.

3.1 Bonus Bank Contracts

The cliff-contract can be implemented as a type of bonus bank contract. Bonus bank arrangements are used in practice to lengthen managerial horizons and better align them with those of the shareholders. The basic idea of a bonus bank arrangement is to "bank" a portion of the bonus received today and pay it out at some future date T conditional on certain benchmarks being reached between now and T. Bonus banks feature prominently in many major companies, notable longstanding examples include Coca-Cola and Temasek - the $215 billion sovereign investment arm of the government of Singapore. Recently, a number of major banks have also adopted various bonus bank arrangements, including Morgan Stanley, Credit Suisse, Barclay's, and UBS. In the case of UBS, a sweeping overhaul of is pay system was enacted in 2013. Central to this overhaul was a new bonus structure for 6500 of its highest earners. Under the new structure, bonuses are issued as "callable" amortizing bonds that mature anywhere between 3 and 5 years. 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 (Schäfer and Shotter, 2013).

In practice, since many of the bonus bank arrangements have only been implemented since the financial crisis, wipeouts have been rare. One notable wipeout occurred at Temasek, one of the few companies that had been using bonus bonds since before the crisis. After portfolio returns missed risk-adjusted hurdles in 2008/2009, bonus bank accounts were wiped out. This is noted in the *Remuneration Philosophy* page of the Temasek website.

I now show how the cliff-contract can be implemented as a bonus bond arrangement. Under the implementation, whenever $N > 0$ and high output is produced, the agent receives a new bonus bond with face value:

$$F = r(1) + \sum_{n=1}^{\infty} \beta^n \left[r(n+1) - r(n) \right] = \frac{pc}{\beta(2p-1)}$$

Notice the more severe the myopic agency problem - as measured by lower p, or lower β, or higher c - the higher the face value F.

The bond's yield y matches the discount factor: $1/(1+y) = \beta$. The amortization table is summarized in the figure below:

Date Since Issue	Payment	Remaining Principal
0	$r(1)$	$\sum_{n=1}^{\infty} \beta^n \left[r(n+1) - r(n) \right]$
1	$r(2) - r(1)$	$\sum_{n=1}^{\infty} \beta^n \left[r(n+2) - r(n+1) \right]$
2	$r(3) - r(2)$	$\sum_{n=1}^{\infty} \beta^n \left[r(n+3) - r(n+2) \right]$
...	...	...
k	$r(k+1) - r(k)$	$\sum_{n=1}^{\infty} \beta^n \left[r(n+k+1) - r(n+k) \right]$
...	...	...

Notice as $k \to \infty$ the bond's principal amortizes completely.

Whenever low output is produced, all existing bonus bonds are called back at zero price and the unpaid portions of the bond principals are wiped out. The agent does not need to give back any payments already received.

Under this bonus bank arrangement, the agent receives zero payment whenever low output is produced - just like the cliff-contract. Thus, to verify that it implements the cliff-contract it suffices to show that following high output, the agent's total payment equals $r(N)$. If $N = 0$, the agent is not issued a bond and has no pre-existing bonds that haven't already been wiped out. So his payment is $0 = r(0)$. If $N \geq 1$ then the agent is issued a new bonus bond. In addition, the agent has $N - 1$ existing bonus bonds issued over the previous $N - 1$ dates. The bond he just received pays him $r(1)$ today. The bond he received the previous date pays him $r(2) - r(1)$ today. In general, the bond he received k dates ago pays him $r(k+1) - r(k)$ for $k < N$. Thus, today's total payment is $r(1) + (r(2) - r(1)) + (r(3) - r(2)) + \ldots + (r(N) - r(N-1)) = r(N)$. This completes the verification.

One notable difference between the cliff-contract's implementation and a real-life bonus bond contract is that unlike real-life bonus bonds, a myopic agency bonus bond does not mature in finite time. However, since $r(N)$ levels off, the later payments of the amortization schedule are vanishingly small. Figure 3 shows an example amortization schedule. This bond

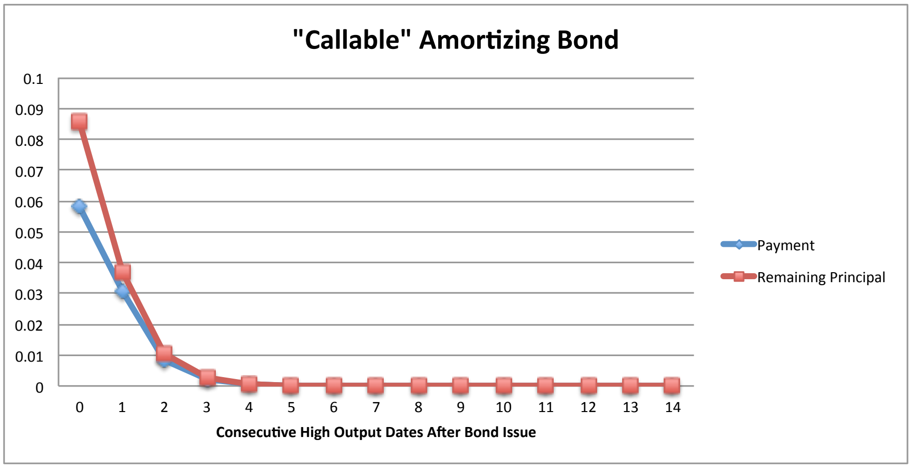

Figure 3: Evolution of the payments and remaining principal of a bonus bond. Parameterization: $(p = .8, \beta = .9, c = .1)$.

can be reasonably approximated using one that matures after three or four payments.

Another notable difference is that in most real-life bonus bond contracts, missing the benchmark does not automatically trigger a full-blown wipeout. That the cliff-contract with its complete wipeouts is optimal is due to the special nature of the model. In Section 4.3, I will show how a *damped* version of the cliff-contract can be robustly adapted to settings that are much more general than the model considered here as a type of restricted optimal contract.

3.2 Incentive Escalation

What if there is no cliff?

The cliff mechanism of the optimal contract works by tying the size of the high output reward today to the size of the high output streak: When the high output streak is long and a low output is produced, it will be a while before the agent is again rewarded generously for high output. Now suppose for the sake of contrast the cliff is eliminated by fixing the high output reward at $r(N = \infty)$. In this case, the high output payment is

$$r(\infty) := p\left[\varepsilon^{\infty}(0) - \frac{c}{p(1+\beta)}\right] = \frac{\varepsilon^{\infty}(0)}{1+\beta}.$$

This cliff-less contract always treats the agent as if he has produced an infinite streak of high output leading up to today. Not surprisingly, it rewards the agent generously: The contract pays the agent whenever the cliff-contract pays the agent, but the payment amount is greater than any payment the agent would ever receive under the cliff-contract. This cliff-less contract has a pay-to-performance sensitivity that is always larger than that of the cliff-contract and a reasonable first-pass assessment of this contract is that it induces the long-term action but at too high a cost. But this is actually giving the cliff-less contract too much credit.

The Markov law for the incentive level of the cliff-less contract is trivial: It stays at $r(\infty)$. A look back at Theorem 1 reveals that increasing all the high output rewards to $r(\infty)$ does not universally increase the incentive level even though it does universally increase pay-to-performance sensitivity. The incentive level is increased after some histories and decreased after others. Unfortunately, since $r(\infty) < \varepsilon^{\infty}(0)$, (2) reveals that the incentive structure is distorted in the worst possible way. After every history the incentive level today is too large compared to tomorrow's and too small compared to yesterday's. As a result, not only does the contract pay the agent too much, the agent always takes the short-term action.

Fixing the high output payment at any level below $r(\infty)$ will result in the same short-termist outcome. Thus, consistently inducing the long-term action without the cliff requires a constant high output reward larger than $r(\infty)$. How much larger must it be? The answer follows from looking at the following issue:

What if the principal models myopic agency as traditional effort-shirking agency?

The primary motivation of this paper is to study the types of hidden actions - both good and bad - that a manager can take beyond well-studied hidden actions of the standard effort-shirking framework. I now show what happens if a principal facing a myopic agency problem replaces the true model with its traditional effort-shirking approximation.

In the traditional version of the model, no action taken today has any effect on output in the future. 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 $1 - \beta p$ of producing high output today. Notice, the principal replaces the short-term action's true multi-period effect on output with its present value. The cost of effort/private benefit of shirking is still c.

Recall in the true model, I assume that the parameters are such that $(1-p)X + c < \beta p X$. That is, the present gain associated with taking the short-term action today is outweighed by the loss tomorrow factoring benefits. In the traditional version of the model, this inequality is precisely the condition that ensures effort is first-best optimal. Therefore, Assumption (A) can still be sensibly applied.

Lemma 2. *In the traditional model, the optimal contract subject to always inducing effort is a stationary contract that pays the agent $\varepsilon^{\infty}(0)$ whenever output is high and nothing whenever output is low. Both the pay-to-performance sensitivity and the incentive level of the contract are always $\varepsilon^{\infty}(0)$.*

Proof. The two date 1 continuation contracts can both be assumed to be the optimal contract. If they weren't, then they would be more expensive, in which case, they could each be replaced with the optimal contract and the surplus payment can be move to date 0. This would not affect incentives nor the cost of the contract (and would make the payment path faster). Finally, given that the two date 1 continuation contracts are the same contract, the optimal thing to do at date 0 is to pay the agent nothing following low output and $e^{\infty}(0)$ following high output. The rest follows by recursion. □

I call this optimal contract of the traditional model the *traditional contract* to distinguish it from the cliff-contract - the optimal contract of the true (myopic agency) model. Just like in the myopic agency setting, always inducing effort in the traditional setting is optimal if the value of the firm is sufficiently large compared to the cost of the traditional contract. Again, this condition can be ensured by assuming X is sufficiently large or c is sufficiently small. (2) implies that the traditional contract is incentive-compatible in the true model. That is, even though the principal sweeps the myopic component 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 cliff-contract. The degree of inefficiency can be usefully quantified. Theorem 1 implies: Relative to the cliff-contract, the traditional contract always over-incentivizes the agent. The traditional contract's over-incentivizing is financed by an inflation in pay compared to the cliff-contract. In the cliff-contract, payments are bounded above by $r(\infty) = \frac{\varepsilon^{\infty}(0)}{1+\beta} < \varepsilon^{\infty}(0)$.

Remark. *The traditional contract is the cheapest contract that induces the long-term action without a cliff. The pay the agent receives after each high output under the traditional contract is strictly higher than the pay he receives, no matter how good his performance, under the cliff-contract.*

This remark also answers the question posed at the end of the previous subsection. To induce the long-term action without a cliff, the best one can do is to import the optimal effort-inducing contract from the effort-shirking version of the model. And this traditional contract pays the agent more, sometimes, a lot more:

The difference in cost to the principal between the cliff and traditional contracts can be substantial. By Theorem 1, the cost of the optimal contract is $c/(1-\beta^2)$, whereas the cost of the traditional contract is $pc/[(1-\beta)(\beta p-(1-p))]$. Thus, the traditional contract is $p(1+\beta)/(\beta p-(1-p))$ times as costly as the optimal contract. Since the only restriction the model imposes on the parameters is that $(1-p)X + c < \beta p X$, this cost ratio can be arbitrarily high.

Moreover, the contract costs, and therefore the cost ratio, are independent of the high output level X. This has implications for the practical importance of the cliff-contract. Recall, always inducing the long-term action is optimal if the firm value is sufficiently large compared to the cost of the cliff-contract. At first, this seems like a weak motivation for using the cliff-contract: When firm value is large relative to contract cost, is it really necessary to employ the absolute optimal contract? Wouldn't something simpler like a traditional cliff-less contract do approximately just as well? The answer is, in general, no. Just because the cliff-contract is relatively cheap does not mean the traditional contract must be as well.

Lemma 3. *Fix any constants $r_1, r_2 \in (0,1)$. There exist parameters of the model such that it is optimal to always induce the long-term action, the optimal contract's cost is less than r_1 times the total value of the firm but the traditional contract's cost is more than r_2 times the total value of the firm.*

Proof. See appendix. □

In particular, by setting r_1 sufficiently low and r_2 sufficiently high, I can simultaneously ensure that always inducing the long-term action is optimal and the cliff-contract is significantly cheaper than the traditional contract. By extension, the cliff-contract will also be significantly cheaper than any other "intuitive" effort-inducing contract imported from the traditional setting, such as periodic review contracts where the agent is rewarded at the end of a review period if and only if the proportion of high output dates matches or exceeds some threshold that approximates p.

The Bonus Reform Debate

The comparative analysis of this section helps put into context some aspects of the recent debate on bonus pay reform. Murphy (2013) discusses the EU proposed cap on bonuses and has a good discussion of and numerous references relating to the bonus reform debate. See also chapter 6 of French et al (2010). In recent years, banks have faced increasing government and popular pressure to change the way they award bonuses. Two of the basic complaints are

1. Bonuses are too large.
2. Bonuses partially contributed to the short-termist behavior that led to the current bad state of the market.

On the other side, the banks counter

3. Bonuses, including potentially large bonuses, are needed to incentivize agents.
4. Capping bonuses can have potentially disastrous consequences due to incentive misalignment.

These contradicting points can be reconciled when understood through the comparative discussion of myopic versus traditional effort-shirking agency.

Given that much of the theory and practice of incentive contract design has been developed within the traditional effort-shirking framework, it is not a stretch to imagine that banks may still think in this way even if the agency problem they face is myopic in nature. If this is true, then the conflicts listed above can be rationalized. For example, the optimal bonus a bank will pay thinking in the traditional framework is $\varepsilon^{\infty}(0)$. The remark from the previous subsection says that this bonus is much too high, but from the bank's perspective it is perfect and lowering the bonus any more will have disastrous consequences, leading to short-termism all the time.

This idea of thinking in the wrong framework can be pushed further. Suppose the bank does not perfectly calculate the severity of the agency problem. It the bank overestimates, then pay is too high and the previous discussion still applies. More interestingly, suppose the bank slightly underestimates the severity of the agency problem and initially sets the bonus to be a little less than $\varepsilon^{\infty}(0)$, in the range $[r(\infty), e^{\infty}(0))$. Given the discussion in "*What if there is no cliff?*" I know that such a bonus is still too high all the time. Moreover, incentives are distorted in the worst possible way so that the agent ends up always taking the short-term action. But when the bank finally realizes that its current bonus system is inducing rampant short-termism, it will not lower bonus payments. Instead, because it is thinking in the wrong framework, the bank will "rationally" conclude that bonus payments must get larger.

This example shows how the banks and their critics can both be right. If the banks could just realize that they're facing a myopic agency problem, then they would lower payments, implement bonus bank systems, and the problem would be solved. On the other hand, if the way bonuses are paid out isn't structurally changed, then the banks are right that further lowering bonuses will just make things worse. Thus, the four points of the aforementioned debate on bonus pay reform are reconciled.

In contrast, suppose the principal overestimates the agency problem but understands that the agency problem is myopic in nature. Then the long-term action is still induced but overpayment is not as severe as it would be if the principal had the wrong model in mind. A unit increase in the estimate of c leads to a $1/(1-\beta^2)$ increase in cost under the cliff-contract but a larger $p/[(1-\beta)(\beta p-(1-p))]$ increase in cost under the traditional contract. The contrast is even greater if the overestimation involves the principal getting p wrong. If the principal thinks p is smaller than it actually is, the traditional contract's cost goes up but the cost of the cliff-contract is independent of p. Finally, if the principal initially underestimates the agency problem and then realizes his mistake, then any increase in pay is in line with optimality unlike the thinking-in-the-wrong-framework case.

4 Robustness

4.1 Non-Deterministic Short-Term Actions

In the main model, the short-term action deterministically leads to the bad state tomorrow. I now relax this assumption and show that the main features of the optimal contract, including incentive escalation and the embedded cliff are robust to this extension model. So, suppose the short-term action now causes the bad state to occur tomorrow with probability $\pi \leq 1$. Let $Q := \pi p$. Condition (1) for the first-best optimality of the long-term action now becomes $(1-p)X + c < \beta Q X$.

The IC-constraints (2) and (3) corresponding to deviation in action only and deviation in action and output disposal become

$$\Delta_{t+1}(h_{t-1}X) \geq e(\Delta_t(h_{t-1})) := \frac{(1-p)\Delta_t(h_{t-1}) + c}{\beta Q} \tag{6}$$

and

$$\Delta_{t+1}(h_{t-1}0) \geq g(\Delta_t(h_{t-1})) := \frac{-p\Delta_t(h_{t-1}) + c}{\beta Q} \tag{7}$$

Here, $e(\Delta)$ and $g(\Delta)$ are the $\varepsilon(\Delta)$ and $\gamma(\Delta)$ of the extension model. The IC-constraint corresponding to deviation in output disposal only is still a nonnegativity constraint on the incentive level. Following the same arguments leading up to Theorem 1, I can now state the optimality result for the extension model:

Proposition 1. *The optimal cost function $C(\Delta)$ satisfies the Bellman equation*

$$C(\Delta) = \max\{\beta C(e(\Delta)) - (1-p)\Delta, \beta C(g(\Delta) \vee 0) + p\Delta\}.$$

It is an $m+1$-piece piecewise linear function for some $m \geq 1$ and it is strictly increasing except for the knife-edge case $\pi = 1$ considered in Theorem 1. There exists a threshold Δ^* such that

$$C(\Delta) = \begin{cases} \beta C(e(\Delta)) - (1-p)\Delta & \text{if } 0 \leq \Delta \leq \Delta^* \\ \beta C(0) + p\Delta = \beta C(g(\Delta) \vee 0) + p\Delta & \text{if } \Delta \geq \Delta^* \end{cases}$$

The corresponding Markov law for the incentive level is isomorphic to Theorem 1's. In particular, the optimal contract's incentive level starts at zero and remains in the set $\{0\} \cup \{e^n(0)\}_{n\geq 1}$. When $\Delta = e^n(0)$, it is located in the domain of the $n+1$-th linear piece of $C(\Delta)$ for $n \in \{1, 2, \ldots m-1\}$. For all $n \geq m$, $e^n(0)$ is located on the $m+1$-th linear piece.

Proof. See appendix. □

Corollary. *The optimal contract is still a cliff-contract: The agent is paid for high output if and only if Δ is in the domain of the $m+1$-th linear piece. Over this domain, the agent's high output reward is linear in Δ and therefore, is an increasing, concave and convergent function of the number of consecutive high outputs. This generalizes Theorem 1 where $m = 1$.*

The corollary to Proposition 1 demonstrates that the optimal contract of the extension model is very similar to the special case optimal contract of Theorem 1. The main difference is that, in general, the agent may need to establish a high output streak longer than one before being rewarded for producing high output.

The explicit characterization of the optimal cost function can be more involved than in Theorem 1 because there may be more than two linear pieces. This can make writing down the piecewise linear function notationally cumbersome. Step 1d in the proof of Proposition 1 provides the precise algorithm for how to write down the function.

One notable difference moving from the special case to the general case is that when $C(\Delta)$ involves more than two pieces, the agent receives low output payment when $N \in \{1, 2, \ldots m-1\}$. This corresponds to when the incentive level is in the domain of the middle pieces of $C(\Delta)$; it is empty whenever $C(\Delta)$ has only two pieces. However, it would be incorrect to think of this payment as a reward for low output. When the contract is in this region, the incentive level is strictly positive. This means that the agent's ex-post promised value (recall Definition 1) is strictly smaller following low output even though he receives a payment following low output and nothing following high output.

The low output payment is an artifact of the limited liability assumption interacting with the IC-constraint (2). This IC-constraint puts a cap on today's incentive level relative to tomorrow's. When the limited liability constraint is slack for the high output payment, the principal can always temper today's incentive level by decreasing the high output payment. However, once limited liability binds for the high output payment (which is precisely what happens when $N \in \{1, 2, \ldots m-1\}$) then the only way to temper today's incentive level is to increase the low output payment from zero. Limited liability also binds when $N = 0$, but that is when $\Delta = 0$ and the optimal thing to do is to pay the agent nothing regardless of output and enact the optimal $e(\Delta)$ contract the next date.

Does the existence of payment for low output corrupt the cliff-based intuition for the optimal contract? No. Ultimately, what the agent cares about is his ex-post promised value not spot payments. The reason why I first introduced the intuition for the cliff-mechanism through spot payments is because they comprise a good shadow variable for ex-post promised value - particularly in the special case considered in Theorem 1 - and because spot payments are more concrete objects that reflect our practical conception of incentive pay.

However, the fundamental reason the agent is driven to take the long-term action is because the incentive-structure and not just the spot-pay-structure of the cliff-contract is well-designed. The optimal contract is a cliff-contract not just because the high output payment is a function of the high output streak but also because the incentive level is a function of the high output streak:

Initially, when the high output streak N is low, the incentive level is small and therefore so is the high output ex-post promised value. Most of the large high output ex-post promised values are backloaded and therefore, initially, the agent is induced to take the long-term action to build up N and access those larger high output ex-post promised values. Eventually, as the agent continues to produce high output and N continues to increase, the high output ex-post promised values stop getting much larger and backloading ceases. But this entire time, the increase in the high output ex-post promised value is being driven by the escalation of the incentive level. Thus, by the time the high output ex-post promised values have leveled off and the agent starts to think about taking the short-term action again, the difference between the high output ex-post promised value and the low output ex-post promised value has become very large. Thus, the agent now finds himself on top of an incentive-cliff, and the fear of falling off this cliff becomes the substitute to backloading that continues to induce the long-term action.

4.2 Participation Constraints

Suppose the agent has an outside option worth K that imposes an ex-ante participation constraint. This means the contract must deliver at least promised value K to the agent, but the agent cannot walk away once the contract is signed. If the optimal contract without the participation constraint promises less than K to the agent then simply add a signing bonus to cover the difference. The resulting contract does not affect the incentive structure and so remains incentive-compatible. It is the optimal contract satisfying the participation constraint.

Now suppose K represents an *interim-participation constraint*: The agent can walk away from the contract any time and exercise his outside option worth K. In this case, the optimal contract will change in a more persistent way and the corresponding Bellman equation will be different. The principal must guard against three different participation constraint driven deviations: 1. The agent can simply walk away. 2. The agent can take the short-term action and then walk away the next date. 3. The agent can take the short-term action, dispose output, and then walk away the next date. The type 1 deviation guarantees the agent K. Optimally choosing between type 2 and 3 deviations guarantees the agent $w_t(h_{t-1}X) \vee w_t(h_{t-1}0) + c + \beta K \geq c + \beta K$. Let $C^{pc}(\Delta)$ denote the cost of the optimal Δ-contract in this world with interim participation constraints.

Proposition 2. *The optimal cost function $C^{pc}(\Delta)$ satisfies the following Bellman equation:*

$$C^{pc}(\Delta) = \max\{\beta C^{pc}(\varepsilon(\Delta)) - (1-p)\Delta, \beta C^{pc}(\gamma(\Delta) \vee 0) + p\Delta, c + \beta K, K\}. \tag{8}$$

The solution is a weakly increasing two-piece piecewise linear function:

$$C^{pc}(\Delta) = \begin{cases} C^{pc}(0) & \text{if } 0 \leq \Delta \leq \frac{1-\beta}{p} C^{pc}(0) \\ \beta C^{pc}(0) + p\Delta & \text{if } \Delta \geq \frac{1-\beta}{p} C^{pc}(0) \end{cases}$$

where $C^{pc}(0) = \frac{c}{1-\beta^2} \vee (c + \beta K) \vee K$. The Markov law for the incentive level is identical to that of Theorem 1 and so the optimal contract is still a cliff-contract.

Proof. See main text. □

Consider the relaxed problem where if the agent chooses to take the short-term action today and then walks away tomorrow, he cannot take any of today's cash compensation with him. In the relaxed problem, the type 2 and 3 deviations guarantee the agent precisely $c + \beta K$. Its Bellman equation is naturally (8) and it is straightforward to check that the piecewise linear formula for C^{pc} solves (8).

What Proposition 2 says is that by backloading some of the relaxed optimal contract's compensation in a payoff neutral way, the principal can turn the contract into the optimal contract of the full problem. This payoff-neutral backloading does not change the incentive structure of the contract, and so the optimal contract of the full problem is also a cliff-contract. The precise algorithm for how this shuffling is done is simple. Consider an arbitrary contract of the following form: Today, it delivers value W to the agent; after high output it pays the agent some w_h and after low output it pays the agent some w_l. This contract will induce the agent to not commit the type 2 nor the type 3 deviation if and only if

$$w_h, w_l \leq \overline{w}(W) := W - (c + \beta K)$$

Thus the algorithm for turning the relaxed optimal contract into the optimal contract is as follows: start at date 0 and if a payment exceeds $\overline{w}(C(0))$, take the excess amount and backload it by one date, being sure to scale up the quantity by a factor of $1/\beta$. This keeps the contract's payoff unchanged and ensures that it respects the participation constraint at date 0. Then move to date 1, and for each date 1 continuation contract, compute how much it is worth keeping in mind that some payments may have been backloaded to date 1, and repeat the process etc. Of course, backloading more than the excess over $\overline{w}$ is incentive-compatible as well. But the proposed algorithm preserves the fastest payment path property.

4.3 Taking the Cliff-Contract Beyond the Core Model

In this subsection, I take a more substantial step away from the core model considered in Theorem 1 and Proposition 1, which focused on how the principal can induce the agent to take the long-term action when the state is always *good* on the equilibrium path. I now introduce multiple states, some better than others, all of which can occur on the equilibrium path, and allow the agent the opportunity to take any one of a number of different short-term actions in each of these states. This allows me to look at what happens when the agent is taking the long-term action and through no fault of his own, ends up in a bad state. The fear is if the agent is facing a cliff-contract, he may be tempted to take a short-term action in a desperate attempt to keep the streak going.

I then further extend the model and introduce short-term actions whose effects play out over more than two dates. I show as an example how a three-date short-term action can be decomposed into two two-date short-term actions taken in succession. Any contract that induces the long-term action in the presence of these two two-date short-term actions will also induce the long-term action in the presence of the single three-date short-term action. In general, I show how the techniques and mechanisms coming from the core model can be effectively utilized to induce the long-term action in a rich environment with many states and short-term deviations.

I stress I do not find the true optimal contract in these more general settings where the optimal contracting problem may be intractable. Rather, I highlight how the cliff-contract is naturally adapted to these richer models of myopic agency by showing how it arises as the optimal arrangement under a restricted, robust version of the optimal contracting problem where the state-dependent incentive constraints are replaced with state-independent constraints.

4.3.1 Multiple States

In the first extension, I introduce multiple states, all of which can occur on the equilibrium path, and allow the agent to take the short-term action in any state. I make two restrictions on the contract space and solve for the resulting restricted optimal contract. The first restriction is that the contract does not ask for agent reports. The second restriction imposes some structure on the incentive escalation implied by the short-term deviation IC-constraint. Later on, I also make a third restriction that imposes an upper bound on the promised value. However, this restriction is mostly done for quantitative reasons and does not affect the robustness of the cliff-contract. I will elaborate when I get to Restriction 3.

There is a finite set of states $\{\sigma_i\}$. If today the state is σ, then the probabilities of high output today given the agent's action choice today are $0 \leq p_{\sigma l} \leq p_{\sigma s} \leq 1$. Taking the short-term action in state σ provides benefit c_σ. Given the previous date's action choice a, the probability that today's state is σ is $\mu_{a\sigma}$. Condition (1) for the first-best optimality of the long-term action now becomes:

$$(p_{\sigma s} - p_{\sigma l})X + c_\sigma < \beta \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma_i l} X \qquad \forall \sigma \tag{9}$$

Restriction 1. *Contracts do not ask the agent to send reports.*

This was without loss of generality in the core model but is not without loss of generality here since different states can occur on the equilibrium path. The upside of not asking for reports is that the "restricted" optimal contract I will derive has some attractive robustness properties which I detail later. Moreover, I emphasize that this restriction does not somehow allow me to avoid the motivating cheating-in-a-bad-state problem mentioned earlier. Without the ability to tell the principal that a bad state has unfortunately occurred, the temptation to cheat is even greater. Thus, if anything, it will be more difficult to show that the cliff-contract is an useful contract under Restriction 1.

The IC-constraint (2) for deviation in action only now becomes

$$p_{\sigma s}\Delta_{t+1}(h_{t-1}X) + (1-p_{\sigma s})\Delta_{t+1}(h_{t-1}0) \geq \frac{(p_{\sigma s}-p_{\sigma l})\Delta_t(h_{t-1}) + c_\sigma}{\beta\sum_{\sigma_i}(\mu_{l\sigma_i}-\mu_{s\sigma_i})p_{\sigma_i l}} \qquad \forall\sigma \tag{10}$$

and the IC-constraint (3) for deviation in action and output disposal now becomes

$$\Delta_{t+1}(h_{t-1}0) \geq \frac{-p_{\sigma l}\Delta_t(h_{t-1}) + c_\sigma}{\beta\sum_{\sigma_i}(\mu_{l\sigma_i}-\mu_{s\sigma_i})p_{\sigma_i l}} \qquad \forall\sigma \tag{11}$$

The IC-constraint for deviation in output disposal only is still $\Delta_t(h_{t-1}) \geq 0$.

The core model is the special case where $p_{\sigma s} = 1$ for $\sigma = good$, $p_{\sigma s} = 0$ for $\sigma = bad$, and $c_{bad} = 0$. In the core model, (10) reduces to a lower bound on $\Delta_{t+1}(h_{t-1}X)$, and it's optimal to set $\Delta_{t+1}(h_{t-1}0)$ to be as close to zero as (11) will allow. In general, though, there can be equilibrium path σ where $p_{\sigma s}$ is significantly less than one, in which case, it is not optimal to have the incentive escalation from Δ_t to Δ_{t+1} be entirely borne by $\Delta_{t+1}(h_{t-1}X)$. Instead, it will be optimal for $\Delta_{t+1}(h_{t-1}X)$ and $\Delta_{t+1}(h_{t-1}0)$ to share the escalation.

The extent of optimal incentive escalation sharing may fluctuate in complex, history dependent ways that are sensitive to the particular details of the model. To make progress, I make a second restriction that imposes some structure on the sharing:

Restriction 2. *Replace* (10) *with sufficient constraints of the form:*

$$\Delta_{t+1}(h_{t-1}X) \geq E_X(\Delta_t(h_{t-1})) := M\Delta_t(h_{t-1}) + B \tag{12}$$

$$\Delta_{t+1}(h_{t-1}0) \geq E_0(\Delta_t(h_{t-1})) := \theta M\Delta_t(h_{t-1}) + B \tag{13}$$

for some nonnegative constants θ, M, and B.

This restriction replaces the state-dependent, non-fixed-share IC-constraint (10) with a pair of state-independent, fixed-share IC-constraints. It does not restrict the overall degree of escalation sharing in the resulting contract. By shifting θ, the contract designer can set the degree of sharing to be whatever he wants. Rather, the restriction simply asks the designer to pick a sharing rule and stick to it.

An immediate benefit of imposing Restriction 2 is that (11) is now redundant. To see the redundancy, fix a $\Delta_{t+1}(h_{t-1}0)$. (13) implies that $\Delta_{t+1}(h_{t-1}0) \geq B$. Thus, it suffices to show that $B \geq c_\sigma/(\beta\sum_{\sigma_i}(\mu_{l\sigma_i}-\mu_{s\sigma_i})p_{\sigma_i l})$ for all σ. By setting $\Delta_t(h_{t-1}) = 0$, and $\Delta_{t+1}(h_{t-1}0) = \Delta_{t+1}(h_{t-1}X) = B$, the inequality follows from Restriction 2 requiring that (12) and (13) imply (10).

By eliminating (11), the optimal contracting problem which originally was a constrained maximization problem subject to (10), (11), and $\Delta_t(h_{t-1}) \geq 0$, now simplifies to:

The Restricted Optimal Contracting Problem

$$\max_{\{w_t \geq 0\}_{t=0}^{\infty}} \mathbf{E}_{\{a_t=l\}_{t=0}^{\infty}} \left[\sum_{t=0}^{\infty} \beta^t (X_t - w_t) \right]$$
$$\text{s.t. for all } t, \text{ (12), (13), } \Delta_t(h_{t-1}) \geq 0$$

Because this maximization problem is constructed in a state-independent way, the restricted optimal contract ignores the private information of the agent. However, due to the ignorance, the contract does have some appealing robustness properties. Because the IC-constraints are state-independent, the contract induces the long-term action under all states and all histories including off-equilibrium histories. Moreover, the contract induces the long-term action no matter the degree of common knowledge of the underlying state. In the core model I've assumed that it is common knowledge that only the agent knows the state. But the restricted optimal contract will also induce the long-term action, for example, if the agent is sometimes not sure what the state is and the principal is not sure if the agent knows the state.

Implicit in the restricted optimal contracting problem is the problem of optimally choosing θ, M, and B. I will solve the problem in two steps. First, I deduce and solve the Bellman equation for the cost function taking θ, M, and B as given. Then I find the optimal θ, M, and B and discuss the restricted optimal contract.

Lemma 4. *Given θ, M, and B, the cost function $C(\Delta)$ satisfies the following Bellman equation:*

$$C(\Delta) = \max\{\beta C(E_X(\Delta)) - (1-\overline{p})\Delta, \beta C(E_0(\Delta)) + \overline{p}\Delta\}$$

where $\overline{p} = \sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i l}$. When $\beta\theta M < 1$ and $\frac{\beta M - 1}{1-\overline{p}} \leq \frac{1-\beta\theta M}{\overline{p}}$, the solution is a linear function:

$$\begin{aligned} C(\Delta) &= \beta C(E_0(\Delta)) + \overline{p}\Delta \\ &= \frac{\overline{p}}{1-\beta\theta M}\Delta + \frac{\beta}{1-\beta} \cdot \frac{\overline{p}}{1-\beta\theta M} \cdot B \end{aligned} \tag{14}$$

Otherwise, the cost is infinite.

The corresponding Markov law for the incentive level takes Δ to $E_X(\Delta)$ following high output and to $E_0(\Delta)$ following low output.

Proof. See appendix. □

This lemma is the key result of the restricted optimal contracting problem. Even though I haven't yet computed the optimal θ, M, and B, nor have I described the restricted optimal contract, Lemma 4 already settles all the key structural questions. Indeed, qualitatively speaking, Lemma 4 reveals that the solution to the Bellman equation and the contract's basic incentive structure are both largely independent of the choice of θ, M, and B. Thus the remaining problem of computing the optimal θ, M, and B and figuring out exactly how large contract payments is essentially a quantitative issue.

Before addressing the quantitative side and describing the optimal contract, I impose one last restriction:

Restriction 3. *The promised value is bounded above by some $\overline{W} > \overline{p}V/(1-\beta)$ where*

$$V = \max_{\sigma,\sigma'} \frac{c_\sigma}{\beta \sum_{\sigma_i}(\mu_{l\sigma_i} - \mu_{s\sigma_i})p_{\sigma_i l} - (p_{\sigma' s} - p_{\sigma' l})}$$

Obviously $\overline{W}$ cannot be too small, or else no incentive-compatible contract exists. Once $\overline{W}$ exceeds the given lower bound, the restricted optimal contract exists and its structure is qualitatively invariant over the choice of $\overline{W}$. Notice, if the private benefit is constant over all states or, more generally, if the private benefit is weakly positively correlated with the output boost of the short-term action, then (9) implies that $V < X$. This means that a sufficiently large $\overline{W}$ is simply the value of the firm $\overline{p}X/(1-\beta)$.

Capping the promised value in this way is often unnecessary. For example, in the core model, the promised value of the cliff-contract is always strictly below the promised value of the traditional contract which is strictly smaller than the total value of the firm. In general, Restriction 3 is not crucial for the qualitative results of this section - again, as a reminder, Lemma 4 was proved before I imposed Restriction 3. Later when I explain how the restricted optimal contract is a "damped" version of the cliff-contract, I can easily adapt the explanation to the setting without Restriction 3. Thus, Restriction 3 does not affect the main qualitative features of the restricted optimal contract and therefore does not affect the robustness of the cliff-contract.

However, from a quantitative standpoint, something like Restriction 3 that caps the promised value (or ex-post promised value) can potentially have an impact. As we shall see shortly, Lemma 4 implies that the restricted optimal contract tries to be as close to a cliff-contract as possible: It tries to load as much of the incentive escalation to be after high output, or equivalently, it tries to pick as low a θ as possible without causing the cost function to explode to infinity. This means that, without Restriction 3, sometimes the optimal thing to do is to set θ so low that $M > 1$. In these cases, the escalation following high output is unbounded. The agent's payoff still increases as a function of the high output streak - preserving the cliff-mechanism - but it no longer converges as the high output streak becomes long.[3] Restriction 3 rules out unbounded escalation by capping the promised value. I emphasize that, even without Restriction 3, the restricted optimal contract does not always feature $M > 1$. It will become apparent that, for a generic set of model parameters, Restriction 3 is automatically satisfied and the optimal contract is simply a cliff-contract with bounded incentive escalation.

I now turn my attention to finding the optimal θ, M, and B. It is clear that B should be set to be the maximum y-intercept of the family of linear constraints that comprise (10):

$$B := \max_{\sigma} \frac{c_\sigma}{\beta \sum_{\sigma_i}(\mu_{l\sigma_i} - \mu_{s\sigma_i})p_{\sigma_i l}} \tag{15}$$

[3] One might wonder how such an unboundedly escalating contract can have finite cost much less optimal/minimal cost. The reason is that the unbounded escalation is occurring only along the high output history and the probability of a high output streak of length N occurring converges to 0 as N goes to infinity.

Given θ, Restriction 2 implies that it is optimal to set M to be

$$M := \max_{\sigma} \frac{1}{p_{\sigma s} + \theta(1 - p_{\sigma s})} \cdot \frac{p_{\sigma s} - p_{\sigma l}}{\beta \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma_i l}} \tag{16}$$

Lastly, notice by (14), θ factors into the cost function through θM while the cost function is finite. Given (16), θM is proportional to $\theta/(p_{\sigma s} + \theta(1 - p_{\sigma s}))$ for some state σ. This fraction is increasing in θ, and therefore θ should be as small as Restriction 3 allows. The smallest amount of maximal incentive-escalation occurs when $\theta = 1$ and Restriction 3 is lax enough that $\theta = 1$ is strictly feasible. Thus, the optimal θ is strictly smaller than 1.

More generally, once M given θ is optimally chosen, Lemma 4 implies that the range of θ for which the cost function is finite is a closed interval and so the optimal θ without Restriction 3 would be the lower bound of this interval. Restriction 3 imposes a strictly stronger upper bound and a weakly stronger lower bound on θ - the lower bound could be zero in both cases. As a result, the set of θ that respects Restriction 3 given that M is optimally chosen is a closed subinterval and the optimal θ given Restriction 3 is the lower bound of this subinterval. In particular, since the optimal θ given Restriction 3 is smaller than 1, so would the optimal θ without Restriction 3.

I now state the restricted optimality theorem:

Theorem 2. *The optimal B is given by* (15) *and the optimal M given θ is given by* (16)*. Given* (16)*, the set of θ that respects Restriction 3 is an interval $[\underline{\theta}, \overline{\theta}]$ where $0 \le \underline{\theta} < 1 < \overline{\theta}$. The optimal θ is the smallest possible θ subject to Restriction 3 and so is equal to $\underline{\theta}$.* (14) *implies all three quantities can be explicitly computed.*

The optimal contract pays the agent only when high output is produced. The high output reward is

$$\frac{1 - [\theta(1 - \overline{p}) + \overline{p}]\, \beta M}{1 - \beta\theta M} \cdot \Delta \tag{17}$$

When the optimal $\theta = 0$, the restricted optimal contract is a cliff-contract: After the first date, the incentive level stays in the set $\{E_X^n(0)\}_{n\ge1}$ which is bounded above by $E_X^\infty(0) = B/(1-M)$. Thus, to pay the agent it suffices to keep track of the number N of consecutive high outputs produced leading up to today. As a function of N, the agent's high output reward is increasing, concave and convergent.

When the optimal $\theta \in (0, 1)$, the restricted optimal contract is a "damped" cliff-contract because the agent does not start over when a high output streak is broken. The damped cliff-contract can be naturally implemented as a bonus bank where low output does not completely wipe out all of the existing bonus bonds.

By (17), the high output reward dynamic mirrors the incentive level dynamic. So I will just discuss the incentive level dynamic in detail. The discussion is derived from Lemma 4 assuming Restriction 3. Without Restriction 3, the only aspect of the discussion that would change is that under some model parameterizations, the incentive level would never converge as the high output streak gets longer and longer.

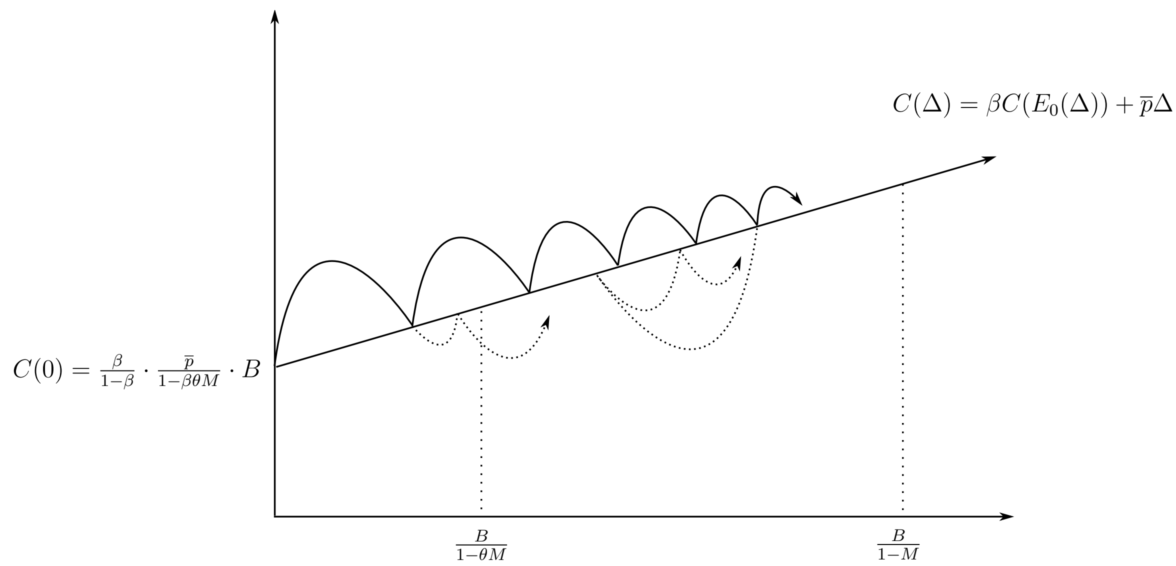

Figure 4: Pictured is the restricted optimal cost function $C(\Delta)$ along with some sample paths of the restricted optimal contract's incentive level. The increasing linear function is $C(\Delta)$. The solid squiggly path represents the location of the restricted optimal contract on the optimal cost function when high output is produced date after date. The path is converging to the point $(B/(1-M), C(B/(1-M)))$. The first diverging dotted path shows that alternate history where after the first high output, a low output is produced followed by high output again. The second diverging dotted path shows the alternate history where after the fifth high output, a low output is produced followed by high output again. Notice the damped drop off the cliff captured by the second dotted path.

The optimal contract's incentive level starts at 0. In the initial phase, the incentive level escalates no matter the output, but high output leads to greater escalation. This does not mean the agent is paid no matter the output. It means that, in the initial phase, if the agent produced high output at date t, then low output at date $t+1$, then high output again at date $t+2$, his date $t+2$ reward is still higher than his date t reward - although not as high as it would have been had he produced high output at date $t+1$ as well. When $\theta = 0$ so that the restricted optimal contract is a cliff-contract, the initial phase is trivial, lasting only the first date.

Once the incentive level reaches or exceeds the fixed point $B/(1-\theta M)$ of E_0, the contract transitions to its steady state dynamic. Now the incentive level forever stays in $[B/(1-\theta M), B/(1-M))$ and moves around in a damped-cliff fashion. High output continues to escalate the incentive level. As the high output streak tends to infinity, the incentive level increases in a concave, convergent way, converging to $B/(1-M)$. Any break in the high output streak causes the incentive level to drop. The higher the incentive level, the bigger the drop. However, as long as $\theta > 0$, the drop is damped - following a low output, the incentive level decreases but doesn't decrease all the way to $B/(1-\theta M)$. Of course, a couple of successive low outputs will drop the incentive level to approximately $B/(1-\theta M)$. See Figure 4 for some sample paths.

Theorem 2 implies that when $\theta = 0$ doesn't violate Restriction 3 then it is optimal to set $\theta = 0$ and the optimal contract is a cliff-contract. It $\theta = 0$ does violate Restriction 3, then it is optimal to get as close to 0 and to load as much of the incentive-escalation on $\Delta_{t+1}(h_{t-1}X)$ as possible. The restricted optimal contract strives to be a cliff-contract, but some damping is necessary.

The reason for the damping is driven by the motivating problem mentioned at the beginning of this subsection: When an agent winds up in a "bad" state - perhaps through no fault of his own - and he is facing a cliff-contract, there may be a strong temptation to cheat and take the short-term action. To see the connection between the temptation to cheat in a "bad" state and the need to damp, I will compare two versions of the multi-state model. The two models are identical except that in some state(s) σ, $p_{\sigma s}$ is higher in the second model. Such a state σ is relatively worse in the second model compared to the first model. Thus, if there were a positive link between cheating in a bad state and damping, one would expect that the restricted optimal contract of the second model would be more damped.

Indeed this is exactly what happens. For any given θ, the optimal M given θ is weakly larger in the second model. Thus, Restriction 3 imposes a weakly higher lower bound on θ in the second model. As a result, the optimal θ is weakly larger in the second model which proves the claim.

4.3.2 Multiple Short-Term Actions

The restricted optimal contracting approach can be easily extended to include multiple short-term deviations $s \in \{s_j\}$. Each short-term deviation s produces high output today with probability $p_{\sigma s}$ and has private benefit $c_{\sigma s}$ in state σ and produces a measure $\{\mu_{s\sigma_i}\}$ over the set of states $\{\sigma_i\}$. The restricted optimal contracting problem is the same except that in (9), (10), and (11), a $\forall s$ needs to be added. Also, in (15) and (16), the maximizations are now over s and σ.

Lemma 4 and Theorem 2 clearly still hold and the restricted optimal contract is still a (possibly damped) cliff-contract.

4.3.3 Longer Short-Term Actions

The above result for inducing the long-term action in the presence of multiple short-term actions is useful because it allows me to further extend the restricted optimal contracting approach and the damped cliff-contract to models with short-term deviations whose effects play out over more than two dates. Here, a K-date short-term action is one that boosts output for the first m dates for some $m < K$ and then lowers output for the last $K - m$ dates. The basic idea is that a K-date short-term action can be be decomposed into a sequence of $K - 1$ two-date short-term actions taken in succession. The restricted optimal damped cliff-contract that induces the long-term action in the presence of the $K-1$ two-date short-term actions also induces the long-term action in the presence of the K-date short-term action.

I will demonstrate how the decomposition works in two examples of a model with a single three-date short-term action. In the first example, this short-term action boosts output for two dates and then lowers output on the third date. In the second example, the short-term action boosts output today and then lowers output on the second and third dates. By working through these two representative examples, I demonstrate the two broadly applicable techniques needed to carry out the decomposition strategy: *homogenization* and *shrinking*. Longer short-term actions can be dealt with similarly.

For simplicity, I will only check for the two deviations of first-order importance: one-shot deviation in action only and deviation in output disposal only. In general, if there are multi-shot deviations or deviations involving some combination of output disposal and action that need to be checked, then for each such deviation, decompose the deviation into a sequence of two-date short-term deviations combined with some output disposal. Take the union of all these short-term actions and find a restricted-optimal contract that induces the long-term action in the presence of this union.

Example 1.
Consider a short-term action that boosts output today and tomorrow and hurts output the day after tomorrow. The expected present discounted production of the firm over today, tomorrow and the day after tomorrow following a one-shot deviation to the short-term action today in state σ is:

$$\left[p_{\sigma s} + \beta \sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl} + \beta^2 \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l}\right] \cdot X$$

where $p_{\sigma s} \geq p_{\sigma l}$ and $p_{\sigma sl} \geq p_{\sigma l}$ for all σ. For simplicity, assume that the benefit of s is the same over all states: $c_\sigma = c$ for all σ.

(9) now becomes

$$(p_{\sigma s} - p_{\sigma l})X + \beta \sum_{\sigma_i} \mu_{l\sigma_i}(p_{\sigma_i sl} - p_{\sigma_i l})X + c < \beta^2 \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma_i l} X \qquad \forall \sigma \tag{18}$$

Let $q_0 := \max_\sigma p_{\sigma s} - p_{\sigma l}$, $q_1 := \sum_{\sigma_i} \mu_{l\sigma_i}(p_{\sigma_i sl} - p_{\sigma_i l})$ and $Q_2 := \sum_{\sigma_i}(\mu_{l\sigma_i} - \mu_{s\sigma_i})p_{\sigma_i l}$. The incentive-constraint for deviation in action only, (10), now becomes

$$\begin{aligned}
&\beta^2 Q_2 \Bigg[p_{\sigma s} \left(\sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl} \Delta_{t+2}(h_{t-1}XX) + (1 - \sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl}) \Delta_{t+2}(h_{t-1}X0) \right) + \\
&\qquad (1 - p_{\sigma s}) \left(\sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl} \Delta_{t+2}(h_{t-1}0X) + (1 - \sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl}) \Delta_{t+2}(h_{t-1}00) \right) \Bigg] \\
\geq &\qquad c + (p_{\sigma s} - p_{\sigma l})\Delta_t(h_{t-1}) + \beta q_1 \left(p_{\sigma s} \Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s}) \Delta_{t+1}(h_{t-1}0) \right)
\end{aligned} \tag{19}$$

I now construct a new model with two two-date short-term actions such that a restricted optimal contract of the new model will also induce the long-term action if used in the original model.

The goal is to find two two-date short-term actions of the same "level of myopia," such that when taken in succession, the net effect is the three-date short-term action. Here, the level of myopia of a short-term action is defined to be the ratio of the present benefit over the present value of the future loss associated with the action. This strategy of decomposing a longer short-term action into two-date short-term actions that are of the same level of myopia is what I call *homogenization*:

Let Q_1 be the unique positive solution to:

$$\frac{q_0}{\beta Q_1} = \frac{Q_1 + q_1}{\beta Q_2} \tag{20}$$

Solving for Q_1:

$$Q_1 = \frac{-q_1 + \sqrt{q_1^2 + 4q_0 Q_2}}{2}$$

Roughly speaking, the new model will have two two-date short-term actions where the boost and drop in the high output probability for the first (second) short-term action are q_0 and Q_1 ($Q_1 + q_1$ and Q_2). By taking these two short-term actions in succession, the net expected change in the high output probability over the next three dates is then $(q_0, -Q_1 + Q_1 + q_1 = q_1, -Q_2)$ which basically matches that of the three-date short-term action.

Here is an exact construction: The new model has the same states and same long-term action as the original model. There are two two-date short-term actions: s_a and s_b with state independent benefits:

$$c_{s_a} = \frac{Q_1}{Q_1 + \beta Q_2} c \qquad\qquad c_{s_b} = \frac{Q_2}{Q_1 + \beta Q_2} c \tag{21}$$

For every state σ define $p_{\sigma s_a} = p_{\sigma s}$. Define $\{\mu_{s_a \sigma}\}$ be a measure over the states so that $\sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s_a \sigma_i}) p_{\sigma_i l} = Q_1$. That this is possible is implied by $Q_1 < Q_2$. Let $\sigma^{\min}$ be the state with the smallest $p_{\sigma l}$. Define $p_{\sigma^{\min} s_b} = p_{\sigma^{\min} l} + Q_1 + q_1$. This is possible because $Q_1 + q_1 < Q_2$ and $1 - p_{\sigma^{\min} l} \geq \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma l} = Q_2$. For all other states assume that $p_{\sigma s_b} - p_{\sigma l} \leq Q_1 + q_1$. Finally, define $\mu_{s_b \sigma} = \mu_{s\sigma}$.

In order to apply restricted optimal contracting approach to the new model, I must now verify that (9) holds for both short-term actions. (20) and (21) imply that s_a and s_b have, loosely speaking, the same level of myopia. Formally, for both short-term actions s,

$$\max_{\sigma} \frac{(p_{\sigma s} - p_{\sigma l})X + c_s}{\beta \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma_i l} X} \leq \frac{q_0 X + c_{s_a}}{\beta Q_1 X} = \frac{(Q_1 + q_1)X + c_{s_b}}{\beta Q_2 X} \tag{22}$$

Thus, it suffices to show that the common ratio of the RHS of (22) is less than 1. Suppose not. Then

$$\begin{aligned}(q_0 + \beta q_1 - \beta^2 Q_2)X + c &= q_0 X + c_{s_a} - \beta Q_1 X + \beta \left[(Q_1 + q_1)X + c_{s_b} - \beta Q_2 X\right] \\ &\geq 0\end{aligned}$$

This contradicts (18). I can now use the techniques of the previous sections to find the restricted optimal contract of the new model which is a (possibly damped) cliff-contract.

Reviewing the previous steps, it is apparent that the key task is to find the right future downside for s_a and present upside for s_b so that (9) holds for both short-term actions and the net effect of taking the two actions in succession matches that of the longer short-term action. Homogenization systematically accomplishes this.[4]

I now show that the restricted optimal contract of the new model, which is a possibly damped cliff-contract, induces the long-term action in the original model. Since the IC-constraint corresponding to deviation in output disposal only is still a nonnegativity constraint on the incentive level, it suffices to verify that the restricted optimal contract of the new model satisfies (19):

$$\begin{aligned}\text{LHS of (19)} & \\ = \; & \beta^2 Q_2 \Bigg[p_{\sigma s}\left(M\Delta_{t+1}(h_{t-1}X)\left(\sum_{\sigma_i}\mu_{l\sigma_i}p_{\sigma_i sl} + \theta(1-\sum_{\sigma_i}\mu_{l\sigma_i}p_{\sigma_i sl})\right) + B\right) + \\ & (1-p_{\sigma s})\left(M\Delta_{t+1}(h_{t-1}0)\left(\sum_{\sigma_i}\mu_{l\sigma_i}p_{\sigma_i sl} + \theta(1-\sum_{\sigma_i}\mu_{l\sigma_i}p_{\sigma_i sl})\right) + B\right)\Bigg] \end{aligned} \tag{24}$$

[4]More generally, if the short-term action produces a sequence of boosts $q_0, q_1, \dots q_m$ followed by a sequence of drops $Q_{m+1}, Q_{m+2}, \dots Q_{m+n}$, then homogenization uniquely determines quantities $Q_1, \dots Q_m$, $q_{m+1}, \dots q_{m+n-1}$ via the equations:

$$\frac{q_0}{\beta Q_1} = \frac{Q_1+q_1}{\beta Q_2} = \dots = \frac{Q_m+q_m}{\beta\left(Q_{m+1}+q_{m+1}\right)} = \frac{q_{m+1}}{\beta\left(Q_{m+1}+q_{m+2}\right)} = \dots = \frac{q_{m+n-1}}{\beta Q_{m+n}}$$

For simplicity, assume $X=1$. Following the example analysis, there are $m+n$ two-date short-term actions $\{s_j\}_{j=1,2,\dots m+n}$. Roughly speaking, the j-th short-term action has a boost today equal to the numerator of the j-th fraction and a drop tomorrow whose present value is equal to the denominator of the j-th fraction. The short-term action benefits c_{s_j} are uniquely determined by the equations $\sum_{t=0}^{m+n-1}\beta^t c_{s_{t+1}} = c$ and

$$\frac{q_0+c_{s_1}}{\beta Q_1} = \frac{Q_1+q_1+c_{s_2}}{\beta Q_2} = \dots = \frac{Q_m+q_m+c_{s_{m+1}}}{\beta\left(Q_{m+1}+q_{m+1}\right)} = \frac{q_{m+1}+c_{s_{m+2}}}{\beta\left(Q_{m+1}+q_{m+2}\right)} = \dots = \frac{q_{m+n-1}+c_{s_{m+n}}}{\beta Q_{m+n}} \tag{23}$$

Finally, to prove that the common ratio of (23) is smaller than 1 it suffices to use proof by contradiction. Suppose not. Then

$$\begin{aligned} & \sum_{t=0}^{m}\beta^t q_t - \sum_{t=m+1}^{m+n}\beta^t Q_t + c \\ = \; & (q_0 - \beta Q_1 + c_{s_1}) + \sum_{t=1}^{m-1}\beta^t(Q_t + q_t - \beta Q_{t+1} + c_{s_{t+1}}) + \beta^m(Q_m + q_m - \beta(Q_{m+1}+q_{m+1}) + c_{s_{m+1}}) + \\ & \sum_{t=m+1}^{m+n-2}\beta^t(q_t - \beta(Q_t + q_{t+1}) + c_{s_{t+1}}) + \beta^{m+n-1}(q_{m+n-1} - \beta Q_{m+n} + c_{s_{m+n}}) \geq 0 \end{aligned}$$

Contradiction. The only concern is that homogenization may produce some short-term actions with infeasible drops. This is not an issue in Example 1 but in the general case, some of the later short-term actions with drops of the form $Q_j + q_{j+1}$ may violate feasibility: $Q_j + q_{j+1} > 1$. This issue is also encountered in Example 2 and is resolved by the second technique: *shrinking*.

By setting $s = s_b$ and $\sigma = \sigma^{\min}$, it is evident

$$M = \max_{s,\sigma} \frac{p_{\sigma s} - p_{\sigma l}}{(p_{\sigma s} + \theta(1 - p_{\sigma s}))\beta \sum_{\sigma_i}(\mu_{l\sigma_i} - \mu_{s\sigma_i})p_{\sigma_i l}}$$
$$\geq \frac{Q_1 + q_1}{(p_{\sigma^{\min} s_b} + \theta(1 - p_{\sigma^{\min} s_b}))\beta Q_2}$$

and

$$B \geq \frac{c_b}{\beta Q_2}$$

Next, $\sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl} = \sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i l} + Q_1$ and $\sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i l} = \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l} + Q_2 \geq p_{\sigma^{\min} l} + q_1$ together imply $\sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl} \geq p_{\sigma^{\min} s_b}$. Since $\theta < 1$,

$$M \geq \frac{Q_1 + q_1}{(\sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl} + \theta(1 - \sum_{\sigma_i} \mu_{l\sigma_i} p_{\sigma_i sl}))\beta Q_2}$$

This inequality and the previous inequality for B imply

$$\text{RHS of (24)} \geq \beta c_b + \beta(Q_1 + q_1)\left(p_{\sigma s}\Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s})\Delta_{t+1}(h_{t-1}0)\right)$$

To complete the verification of (19), it suffices to show that

$$\begin{aligned} &\beta Q_1 \left(p_{\sigma s}\Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s})\Delta_{t+1}(h_{t-1}0)\right) \\ = \beta \sum_{\sigma_i} &(\mu_{l\sigma_i} - \mu_{s_a\sigma_i})p_{\sigma_i l} \left(p_{\sigma s_a}\Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s_a})\Delta_{t+1}(h_{t-1}0)\right) \\ &\geq c - \beta c_b + (p_{\sigma s} - p_{\sigma l})\Delta_t(h_{t-1}) = c_a + (p_{\sigma s_a} - p_{\sigma l})\Delta_t(h_{t-1}) \end{aligned}$$

But this is just the incentive-constraint for s_a which is assumed to be satisfied by the restricted optimal contract.

Example 2.
Consider a short-term action that boosts output today and hurts output tomorrow and the day after tomorrow. The expected present discounted production of the firm over today, tomorrow and the day after tomorrow following a one-shot deviation to the short-term action today in state σ is:

$$\left[p_{\sigma s} + \beta \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l} + \beta^2 \sum_{\sigma_i} \mu_{sl\sigma_i} p_{\sigma_i l}\right] \cdot X$$

where $p_{\sigma s} \geq p_{\sigma l}$, and $\sum_{\sigma_i}(\mu_{l\sigma_i} - \mu_{s\sigma_i})p_{\sigma_i l} > 0$, $\sum_{\sigma_i}(\mu_{l\sigma_i} - \mu_{sl\sigma_i})p_{\sigma_i l} > 0$. For simplicity, assume that the benefit of s is the same over all states: $c_\sigma = c$ for all σ.

(9) now becomes

$$(p_{\sigma s} - p_{\sigma l})X + c < \beta \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma_i l} X + \beta^2 \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{sl\sigma_i}) p_{\sigma_i l} X \qquad \forall \sigma$$

Let $q_0 := \max_\sigma p_{\sigma s} - p_{\sigma l}$, $Q_1 := \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{s\sigma_i}) p_{\sigma_i l}$ and $Q_2 := \sum_{\sigma_i} (\mu_{l\sigma_i} - \mu_{sl\sigma_i}) p_{\sigma_i l}$. Let q_1 be the unique positive solution to:

$$\frac{q_0}{\beta(Q_1 + q_1)} = \frac{q_1}{\beta Q_2}$$

Solving for q_1:

$$q_1 = \frac{-Q_1 + \sqrt{Q_1^2 + 4 q_0 Q_2}}{2}$$

Before I construct the new model with two two-date short-term actions, I first construct an intermediate model with a single three-date short-term action which is basically a scaled down version of the original short-term action. I call this *shrinking*. The reason I need to perform shrinking is because $Q_1 + q_1$ may be larger 1, in which case it is impossible to create a short-term action whose drop in output tomorrow is $(Q_1 + q_1)X$. Shrinking overcomes this problem. Then I decompose the shrunk model's short-term action into two two-date short-term actions and show that the resulting restricted optimal contract still induces the long-term action in the original model.

Fix a sufficiently small $\lambda \in (0, 1)$. In this model, there is a new set of states $\{\hat{\sigma}\}$ such that $p_{\hat{\sigma}s} - p_{\hat{\sigma}l} = \lambda(p_{\sigma s} - p_{\sigma l})$, for every $p_{\sigma s}$ there exists a state $\hat{\sigma}$ such that $p_{\hat{\sigma}s} = p_{\sigma s}$, $\sum_{\hat{\sigma}_i} (\mu_{l\hat{\sigma}_i} - \mu_{s\hat{\sigma}_i}) p_{\hat{\sigma}_i l} = \lambda Q_1$, $\sum_{\hat{\sigma}_i} \mu_{s\hat{\sigma}_i} p_{\hat{\sigma}_i l} \leq \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l}$, and $\sum_{\hat{\sigma}_i} (\mu_{l\hat{\sigma}_i} - \mu_{sl\hat{\sigma}_i}) p_{\hat{\sigma}_i l} = \lambda Q_2$. Also assume that $\sum_{\hat{\sigma}_i} \mu_{l\hat{\sigma}_i} p_{\hat{\sigma}_i l} - \min_{\hat{\sigma}} p_{\hat{\sigma} l} > \lambda(Q_1 + q_1)$. Such a model is possible for any λ chosen to be sufficiently small. Finally, $\hat{c} = \lambda c$.

Now I construct the new model with two two-date short-term actions: $\hat{s}_a$ and $\hat{s}_b$ with state independent benefits

$$\hat{c}_a = \frac{Q_1 + q_1}{Q_1 + q_1 + \beta Q_2} \hat{c} \qquad\qquad \hat{c}_b = \frac{Q_2}{Q_1 + q_1 + \beta Q_2} \hat{c}$$

For every state $\hat{\sigma}$ define $p_{\hat{\sigma}s_a} = p_{\hat{\sigma}s}$. Let $\{\mu_{s_a\hat{\sigma}}\}$ be a measure over the states so that $\sum_{\hat{\sigma}_i} (\mu_{l\hat{\sigma}_i} - \mu_{s_a\hat{\sigma}_i}) p_{\hat{\sigma}_i l} = \lambda(Q_1 + q_1)$. Let $\hat{\sigma}^{\min}$ be the state with the smallest $p_{\hat{\sigma}l}$. Define $p_{\hat{\sigma}^{\min} s_b} = p_{\hat{\sigma}^{\min} l} + \lambda q_1$. For all other states assume that $p_{\hat{\sigma}s_b} - p_{\hat{\sigma}l} \leq \lambda q_1$. Finally, define $\mu_{s_b\hat{\sigma}} = \mu_{sl\hat{\sigma}}$.

Verifying the restricted optimal contract of the new model induces the long-term action in the intermediate model is similar to Example 1. So I will just verify that this contract also induces the long-term action in the original model.

The incentive constraint for deviation in action only is

$$\begin{aligned}\beta^2 Q_2 \Bigg[p_{\sigma s} \Bigg(\sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l} \Delta_{t+2}(h_{t-1}XX) + (1 - \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l}) \Delta_{t+2}(h_{t-1}X0) \Bigg) + \\ (1 - p_{\sigma s}) \Bigg(\sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l} \Delta_{t+2}(h_{t-1}0X) + (1 - \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l}) \Delta_{t+2}(h_{t-1}00) \Bigg) \Bigg] + \\ \beta Q_1 \left(p_{\sigma s} \Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s}) \Delta_{t+1}(h_{t-1}0) \right) \\ \geq c + (p_{\sigma s} - p_{\sigma l}) \Delta_t(h_{t-1}) + \end{aligned} \tag{25}$$

It suffices to verify (25),

$$\begin{aligned}\lambda \cdot (\text{LHS}) \text{ of } (25) = & \\ & \lambda\beta^2 Q_2 \Bigg[p_{\sigma s} \Bigg(M\Delta_{t+1}(h_{t-1}X) \Bigg(\sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l} + \theta(1 - \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l}) \Bigg) + B \Bigg) + \\ & (1 - p_{\sigma s}) \Bigg(M\Delta_{t+1}(h_{t-1}0) \Bigg(\sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l} + \theta(1 - \sum_{\sigma_i} \mu_{s\sigma_i} p_{\sigma_i l}) \Bigg) + B \Bigg) \Bigg] + \\ & \lambda\beta Q_1 \left(p_{\sigma s} \Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s}) \Delta_{t+1}(h_{t-1}0) \right) \\ \geq \quad & \beta \hat{c}_b + \lambda\beta(Q_1 + q_1) \left(p_{\sigma s} \Delta_{t+1}(h_{t-1}X) + (1 - p_{\sigma s}) \Delta_{t+1}(h_{t-1}0) \right) \\ \geq \quad & \hat{c}_a + \beta \hat{c}_b + \lambda(p_{\sigma s} - p_{\sigma l}) \Delta_t(h_{t-1}) \\ \geq \quad & \lambda \cdot (\text{RHS}) \text{ of } (25)\end{aligned}$$

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. The twist is that the short-term action 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. This cliff-contract draws clear parallels with the new bonus-pay arrangements introduced at a number of large financial institutions to combat short-termism in the wake of the financial crisis. The cliff-contract differs in significant ways from traditional dynamic optimal contracts. The cliff-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 myopia-related agency problems and presents a robust contract to induce the long-term action in settings with such agency problems.

6 Appendix

Proof of Lemma 1. Let $W_t^+(h_t) := w_t(h_t) + \beta W_{t+1}(h_t)$ denote the agent's date t ex-post promised value given history h_t. 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.

$$\begin{array}{ll} & W^+(hXX) \\ w(hX) & \\ & W^+(hX0) \end{array}$$

$$\begin{array}{ll} & W^+(h0X) \\ w(h0) & \\ & W^+(h00) \end{array}$$

Suppose the agent decides to commit a one-shot deviation to the short-term action today. Then his payoff is $c + w(hX) + \beta W^+(hX0)$. Letting $W(hX)$ 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 payoff from deviation can be rewritten as $c + w(hX) + \beta W(hX) - \beta(p(W^+(hXX) - W^+(hX0))) = c + w(hX) + \beta W(hX) - \beta p\Delta(hX) = c + W^+(hX) - \beta p\Delta(hX)$. Incentive compatibility requires that

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

which, upon rearrangement, is equivalent to

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

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 $w(hX) + \beta((w(hX0) + \beta U(D(hX0)))$ where $U(D(hX0))$ is the payoffs from employing the continuation $D(hX0)$ of the deviation strategy after history $hX0$. By induction, this payoff is weakly less than $w(hX) + \beta(w(hX0) + \beta(pW^+(hX0X) + (1-p)W^+(hX00))) = w(hX) + \beta W^+(hX0) \leq pW^+(hX) + (1-p)W^+(h0)$.

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 Lemma 3. The ratio of the optimal contract's and traditional contract's costs to the total value of the firm are

$$\frac{c/X}{p(1+\beta)} \qquad \text{and} \qquad \frac{c/X}{\beta p - (1-p)}$$

respectively. A sufficient condition for the optimality of always inducing the long-term action is to assume that the total cost of the optimal contract is less expensive than the loss in production coming from even a single deviation to the short-term action:

$$\frac{c}{1-\beta^2} \leq [\beta p - (1-p)]\, X$$

So to prove them lemma, set $c = (1-\beta^2)(\beta p - (1-p))X$ so that it is optimal to always induce the action. Then the aforementioned ratios simply to

$$\frac{(1-\beta)(\beta p - (1-p))}{p} \qquad \text{and} \qquad (1-\beta)^2$$

It is clear that one can make the first ratio sufficiently small and while the second ratio sufficiently close to 1 by setting β to be sufficiently low and p to be sufficiently high while keeping $\beta p - (1-p)$ sufficiently low. □

Proof of Proposition 1.
Step 1. Solving the relaxed problem.

Consider the relaxed optimal contracting problem without constraint (7). Let $C^r(\Delta)$ denote the cost of the optimal relaxed Δ-contract.

Step 1a. $C^r(\Delta)$ is weakly convex.

Fix $\lambda \in (0,1)$ and optimal relaxed Δ_1- and Δ_2-contracts with 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 relaxed $\lambda\Delta_1 + (1-\lambda)\Delta_2$-contract with cost $\lambda C^r(\Delta_1) + (1-\lambda)C^r(\Delta_2)$ which by definition must be $\geq C^r(\lambda\Delta_1 + (1-\lambda)\Delta_2)$.

Step 1b. If $\Delta > e(0)$ then $C^r(\Delta) > C^r(e(0))$.

Suppose not. Pick a $\Delta^* > e(0)$ satisfying $C^r(\Delta^*) \leq C^r(e(0))$. Then consider the following contract: at date 0, pay the agent nothing; at date 1 give the agent the optimal Δ^*-contract. Since $D^* \geq e(0)$, this contract is incentive-compatible. It is a 0-contract by construction and its cost is $\beta C(\Delta^*)$. So I have shown that $C^r(0) \leq \beta C^r(\Delta^*) < C^r(e(0)) \geq C^r(\Delta^*)$. This contradicts the weak convexity of C^r.

Step 1c. $C^r(\Delta) = \max\{\beta C^r(e(\Delta)) - (1-p)\Delta, \min_{\Delta_2} \beta C^r(\Delta_2) + p\Delta_2\}$.

Following the arguments leading up to Theorem 1, I can claim that for an optimal Δ-contract the continuation contract following high output must be the cheapest $\Delta(h)$-contract subject to $\Delta(h) \geq e(\Delta)$. Step 1b. implies that this must the optimal $e(\Delta)$-contract. Similarly, the continuation contract following low output must be the cheapest $\Delta(l)$-contract where $\Delta(l)$ is unconstrained since by assumption the relaxed problem disregards constraint (7). This implies that C^r satisfies the following Bellman equation (again using the arguments leading up to Theorem 1)

$$C^r(\Delta) = \max\{\beta C^r(e(\Delta)) - (1-p)\Delta, \min_{\Delta_2} \beta C^r(\Delta_2) + p\Delta_2\} \tag{26}$$

If C^r is weakly increasing, then (26) becomes

$$C^r(\Delta) = \max\{\beta C^r(e(\Delta)) - (1-p)\Delta, \beta C^r(0) + p\Delta\}. \tag{27}$$

To give an explicit characterization of the optimal relaxed contract, it suffices to explicitly characterize the solution to (26). To explicitly characterize the solution to (26) it suffices to explicitly characterize the solution to (27) and then verify that it is weakly increasing. This is what I will do. From now on, C^r is defined to be the solution to (27).

Step 1d. $C^r(\Delta)$ is a piecewise linear function that can be explicitly characterized.

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

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

where $f^{\Delta}(0)$ is defined so that $f^{\Delta}(\Delta) = \beta f^{\Delta}(0) + p\Delta$.

Step 1d.1 $f^{\Delta}(0)$ is unique and nonnegative.

Since $f^{\Delta}(\Delta)$ is a linear function of $f^{\Delta}(0)$ with slope 1 and $\beta f^{\Delta}(0) + p\Delta$ is a linear function of $f^{\Delta}(0)$ with slope β, they must intersect at a unique $f^{\Delta}(0)$. To prove nonnegativity, first set $f^{\Delta}(0) = 0$. By the definition of f^{Δ}, it follows that $f^{\Delta}(\Delta) \leq p\Delta = 0 + p\Delta = \beta f^{\Delta}(0) + p\Delta$. Therefore, one must weakly increase $f^{\Delta}(0)$ for the equality to hold.

Next, define $d(\Delta) = f^{\Delta}(0) - \beta f^{\Delta}(e(0))$.

Step 1d.2 $d(\Delta)$ is strictly increasing in Δ when $\Delta \in (0, e^{\infty}(0) = c/(\beta Q - (1-p)))$.

First consider the case when $e^{-N}(\Delta) < 0$. Increase Δ by a small amount dx. Then the change in $f^{\Delta}(\Delta)$ is $df^{\Delta}(0) + m^N(p)dx$; and the change in $\beta f^{\Delta}(0) + p\Delta$ is $\beta df^{\Delta}(0) + pdx$. Since $f^{\Delta}(\Delta) = \beta f^{\Delta}(0) + p\Delta$, therefore $df^{\Delta}(0) + m^N(p)dx = \beta df^{\Delta}(0) + pdx \Rightarrow (1-\beta)df^{\Delta}(0) = (p - m^N(p))dx > 0$. Thus, the change in $d(\Delta)$ is $df^{\Delta}(0) - \beta(df^{\Delta}(0) + m^N(p)dx - m^{N-1}(p)dx) = (1-\beta)df^{\Delta}(0) + \beta(m^{N-1}(p) - m^N(p))dx > 0$.

Next consider the knife-edge case when $e^{-N}(\Delta) = 0$. Then $f^{\Delta}(\Delta) = \beta f^{\Delta}(0) + p\Delta$ implies $df^{\Delta}(0) + m^{N+1}(p)dx = \beta df^{\Delta}(0) + pdx \Rightarrow (1-\beta)df^{\Delta}(0) = (p - m^{N+1}(p))dx > 0$. Then change in $d(\Delta)$ is $df^{\Delta}(0) - \beta(df^{\Delta}(0) + m^{N+1}(p)dx - m^N(p)dx) = (1-\beta)df^{\Delta}(0) + \beta(m^N(p) - m^{N+1}(p))dx > 0$.

When $\Delta = 0$, $d(\Delta) = 0 - \beta pe(\Delta) < 0$. For Δ sufficiently large, $d(\Delta) > 0$. A sufficient condition is for Δ to be large enough so that $m^{N-1}(p) \leq 0$. Let Δ^* be the unique value satisfying $d(\Delta^*) = 0$. Then the corresponding function f^{Δ^*}, solves (27). So $C^r \equiv f^{\Delta^*}$.

Step 1e. C^r is weakly increasing.

In the explicit construction of C^r, one cannot definitively conclude that it must be weakly increasing since for N large enough, m^N is negative. However, when $\pi = 1$, the explicit construction of C^r matches the solution C to (27) in Theorem 1 which is weakly increasing. When $\pi < 1$, I will prove the following stronger result about C^r:

Step 1e.1 C^r is strictly increasing if $\pi < 1$.

(27) implies $C(0) = \beta C(e(0))$ and $C(e(0)) \geq \beta C(0) + pe(0) = \beta^2 C(e(0)) + pe(0)$. So

$$C(e(0)) \geq \frac{pe(0)}{1-\beta^2} \tag{28}$$

Next, let $C' := d^- C^r(e(0))$. Since C^r is weakly convex, it must be that $C(e(0)) - C'e(0) \leq C(0) = \beta C(e(0))$. So

$$C(e(0)) \leq \frac{C'e(0)}{1-\beta} \tag{29}$$

Equations (28) and (29) together imply

$$C' \geq \frac{p}{1+\beta}$$

To prove C^r is strictly increasing it suffices to prove $m(C')$ is positive which is equivalent to proving $C' > Q$. The proof is by contradiction. So suppose instead, $C' \leq Q$. This means Q satisfies the following conditions

$\frac{p}{1+\beta} \leq Q < p$ and $\beta Q > 1-p$. I claim that these conditions imply $\frac{(1-p)p}{Q} - (1-p) < \frac{p}{1+\beta}$. To prove my claim, I prove the following stronger result: if $\frac{p}{2} \leq Q \leq p$ and $Q \geq 1-p$ then $\frac{(1-p)p}{Q} - (1-p) \leq \frac{p}{2}$. First suppose $p \in [\frac{1}{2}, \frac{2}{3}]$. The conditions on Q imply that $Q \geq 1-p$. Therefore $\frac{(1-p)p}{Q} - (1-p) \leq p - (1-p) = 2p - 1 \leq \frac{p}{2}$. Second suppose $p \in [\frac{2}{3}, 1]$. The conditions on Q imply that $Q \geq \frac{p}{2}$. Therefore $\frac{(1-p)p}{Q} - (1-p) \leq 2(1-p) - (1-p) = 1-p \leq \frac{p}{2}$. I have now shown $m(p) = \frac{(1-p)p}{Q} - (1-p) < \frac{p}{1+\beta} \leq C' < p$. But the characterization of C^r shows that its slope is always of the form $m^k(p)$ for some k. Contradiction. C^r is strictly increasing.

Step 2. Solving the full problem.

The same arguments showing $C^r(\Delta)$ is weakly convex also show that $C(\Delta)$ is weakly convex. Since C^r is the solution to the relaxed problem, it must be that $C \geq C^r$. Suppose it is true that $C(0) = C^r(0)$. Then C must also be strictly increasing if $\pi > 1$ and weakly increasing if $\pi = 1$. I now prove that $C(0) = C^r(0)$ by explicitly showing how to transform the optimal relaxed contract into the optimal contract. The rest of Proposition 1 follows immediately.

Take the optimal relaxed contract and at every date when the incentive level is 0, replace the payment and continuation contract following low output with the payment and continuation contract following high output. Notice this does not change the value of the contract: since the incentive level is 0, by definition the payment and continuation contract following low output has the same value as the payment and continuation contract following high output. The incentive level of the contract still stays in the set $\{0, e(0), e^2(0), \ldots\}$. The Markov law for the incentive level is unchanged except when the incentive level is 0. Before, 0 would go to $e(0)$ or 0 depending on the output. Now it deterministically goes to $e(0)$ just like in the optimal contract in Theorem 1. This ensures that constraint (7), which was ignored in the relaxed problem, is now respected. As a result the modified optimal relaxed contract becomes fully incentive-compatible and is therefore the optimal contract. □

Proof of Lemma 4. The derivation of the Bellman equation is the same as before. When $\beta\theta M \leq \frac{\theta}{\overline{p}+\theta(1-\overline{p})}$, the operator $\beta C(E_X(\Delta)) - (1-\overline{p})\Delta$ takes (14) to a function with the same y-intercept but weakly smaller slope. When $\beta\theta M > \frac{\theta}{\overline{p}+\theta(1-\overline{p})}$, then starting with (14), repeated iterations of the operator $\beta C(E_X(\Delta)) - (1-\overline{p})\Delta$ takes the function to infinity. □

References

[1] Debreu, G. *Mathematical Economics: Twenty Papers by Gerard Debreu,* intr. by Hildenbrand, W. Cambridge University Press, 1986

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

[3] *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

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

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

[6] French, K.R., M. N. Baily, J. Y. Campbell, J. H. Cochrane, D. W. Diamond, D. Duffie, A. K Kashyap, F. S. Mishkin, R. G. Rajan, D. S. Scharfstein, R. J. Shiller, H. S. Shin, M. J. Slaughter, J. C. Stein, and R. M. Stulz. *The Squam Lake Report: Fixing the Financial System,* Princeton University Press, 2010

[7] Hoffmann, F. and S. Pfeil (2013) "Delegated Investment in a Dynamic Agency Model," working paper

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

[9] 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

[10] Innes, R. D. (1990) "Limited Liability and Incentive Contracting with Ex-Ante Action Choices," *Journal of Economic Theory* Vol. 52, pp. 45-67

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

[12] Manso, G. (2011) "Motivating Innovation," *The Journal of Finance* Vol. 66, pp. 1823-1869

[13] Murphy, K.J. (2013) "Regulating Banking Bonuses in the European Union: A Case Study in Unintended Consequences," *European Financial Management* Vol. 19, pp. 631-657

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

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

[16] Sannikov, Y. (2014) "Moral Hazard and Long-Run Incentives," working paper

[17] Schäfer, D. and J. Shotter "UBS leads way with bonuses shake-up," *Financial Times* February 5, 2013

[18] Schäfer, D. and C. Hall "Citigroup bankers to be paid at least half their bonuses in cash," *Financial Times* January 26, 2014

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

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

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